THE UNIVERSITY OF CHICAGO

TRACKING FEATURE-SPECIFIC MEMORY REPRESENTATIONS WITH OSCILLATORY BRAIN ACTIVITY

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

DAVID WILLIAM SUTTERER

CHICAGO, ILLINOIS

JUNE 2018

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF PUBLICATIONS

Sutterer, D, J Foster, J Serences, E Vogel, and E Awh. 2018. “Alpha-Band Oscillations Track the Retrieval of Precise Spatial Representations from Long-Term Memory.” *bioRxiv*, 1–26. (Included as Chapter 4)

van Moorselaar, Dirk, Joshua J. Foster, David W. Sutterer, J Theeuwes, Christian N. L. Olivers, and Edward Awh. 2017. “Spatially Selective Alpha Oscillations Reveal Moment-by-Moment Trade-Offs between Working Memory and Attention.” *Journal of Cognitive Neuroscience*, 1–11.

Foster, Joshua J., David W. Sutterer, John T. Serences, Edward K. Vogel, and Edward Awh. 2017. “Alpha-Band Oscillations Enable Spatially and Temporally Resolved Tracking of Covert Spatial Attention.” *Psychological Science* 28 (7): 929–41.

Oberauer, Klaus, Edward Awh, David W Sutterer. 2016. “The Role of Long-Term Memory in a Test of Visual Working Memory: Proactive Facilitation but No Proactive Interference.”

Sutterer, David W., and Edward Awh. 2016. “Retrieval Practice Enhances the Accessibility but Not the Quality of Memory.” *Psychonomic Bulletin & Review* 23 (3): 831–41. (Included as Chapter 3)

Ester, E. F., D. W. Sutterer, J. T. Serences, and E. Awh. “Feature-Selective Attentional Modulations in Human Frontoparietal Cortex.” *Journal of Neuroscience* 36, no. 31 (2016): 8188–99. doi:10.1523/JNEUROSCI.3935-15.2016.

Emrich, Stephen M, Jeffrey S Johnson, David W Sutterer, and Bradley R Postle. 2016. “Comparing the Effects of 10-Hz Repetitive TMS on Tasks of Visual STM and Attention.” *Journal of Cognitive Neuroscience* 29 (2): 286–97.

Foster, Joshua J, David W Sutterer, John T Serences, Edward K Vogel, and Edward Awh. 2016. “The Topography of Alpha-Band Activity Tracks the Content of Spatial Working Memory” 115: 168–77. doi:10.1152/jn.00860.2015.

Postle, Bradley R, Edward Awh, John T Serences, David W Sutterer, and Mark D’Esposito. 2013. “The Positional-Specificity Effect Reveals a Passive-Trace Contribution to Visual Short-Term Memory.” *PloS One* 8 (12): e83483.

Kundu, Bornali, David W Sutterer, Stephen M Emrich, and Bradley R Postle. 2013. “Strengthened Effective Connectivity Underlies Transfer of Working Memory Training to Tests of Short-Term Memory and Attention.” *Journal of Neuroscience* 33 (20): 8705–15.

Johnson, Jeffrey S., David W. Sutterer, Daniel J. Acheson, Jarrod A. Lewis-Peacock, and Bradley R. Postle. 2011. “Increased Alpha-Band Power during the Retention of Shapes.” *Frontiers in Psychology* 2: 1–9.

ACKNOWLEDGEMENTS

First, I would like to thank my advisor, Ed Awh, for his tireless support over the years. His unbounded enthusiasm for experimental design, new discoveries, and engaging in rigorous scientific debate is inspiring and contagious. Of the many lessons he has taught me over the years one in particular will always stick with me: Never let my current skill set determine the types of questions I'm willing to ask. I am also grateful to Ed for the countless hours he has spent with me fine tuning experiments, correcting manuscripts, and pushing me to tackle new challenges. I could not have asked for a more hard-working and dedicated mentor. My gratitude also goes out to Ed Vogel for his guidance throughout my doctoral program. Discussing my projects with Ed has not only improved my individual projects but also my approach to research and science communication. I would also like to thank my committee members David Gallo and David Freedman for the feedback and advice on my dissertation, and for being so actively involved in welcoming to me, and our group, during our transition to the University of Chicago. Special thanks to Dave Gallo for including me in his lab meetings; I have benefited greatly from discussions with him and his lab about episodic memory.

This work would not have been possible without the advice, help, and support of multiple individuals at both the University of Chicago and the University of Oregon. I would like to thank our fantastic lab managers Richard Matullo, Brendan Colson, and Ariana Gale for keeping the lab running smoothly under their care. Thanks as well to Mike, Janet, Sharon, James, and Scott at the IMB for going above and beyond the call of duty day in and day out to support our work. Special thanks to Mimi, Smiley, and Ray. In particular, we would have been totally lost in the transition to U Chicago without Mimi's support. Thanks also to the incredible staff at the University of Oregon, especially Roger Woods, and Lori Olsen.

I owe a great deal of thanks to my lab mates for their unwavering support. I count myself very lucky to have started in the lab with Kirsten Adam and soon after Josh Foster. Moving the lab cross country and rebuilding wouldn't have been nearly as enjoyable without their friendship and support. I am continually humbled by both of their research accomplishments and their generosity of time and knowledge. Thanks to the rest of the incredible group of Awh and Vogel Lab graduate students and post-docs that I've enjoyed working with over the years: Gisella Diaz, Megan deBettencourt, Brittany Dungan, Eren Gunseli, Keisuke Fukuda, Nicole Hakim, Irida Mance, Tobias Feldmann-Wüstefeld, and Colin Quirk. You've each contributed to a supportive lab environment in your own unique ways. I will sorely miss being able to turn one desk in any direction to workshop a new argument, get excellent feedback on a manuscript, or learn how to run a new analysis. Not only have you all been incredible to work with, but you've made earning my doctorate a blast outside of lab as well. I'd also like to express my thanks to the excellent collaborators who I have had the privilege of sharing a lab space and projects with over the years. Thanks to Dirk van Moorselaar, Michel Failing, and Will Ngiam for the fun we had working together and for encouraging me to join you in getting out and exploring Chicago. Thanks as well to my collaborators at UC San Diego. John Serences has managed to provide thoughtful and instructive responses to every email I've ever sent him, usually within a few hours. Tommy, Vy, Sirawaj, Eddie, Steph, Nuttida, Anna, Mary, Maggie, Chaipat, Rosanne, you all have been incredible to work with. A special thanks to the undergraduate research assistants who have helped me over the years: Nikkie Snow, Anubuv Guhpta, Nicholas Diaz, Jared Evans, Dylan Seitz, and Ricardo Fernandez. None of this work would have been possible without your assistance.

I would never have started down this path if not for the incredible people at UW Madison. I'd first like to thank Bryan Hendricks who opened my eyes to the world of experimental psychology. His dedication to his students serves as a constant reminder to be encouraging, patient, and generous with my time. Thanks also to Chris Rozek and David Rozek who taught me that research could be both intellectually gratifying and fun. Of course, many thanks are due to Brad Postle for welcoming an enthusiastic Sconnie kid to his lab, and somehow always giving me just enough rope for the task at hand but not quite enough to hang myself with. Thanks as well to the amazing team I worked with in the Postle Lab; especially Dan Acheson for bringing me in to the lab and Jeff Johnson for taking me under his wing during my first summer in the lab and all the subsequent years I've worked in the field. Thanks also to Steve Emrich for teaching me countless important lessons about data collection, analysis, and which important scientific discoveries and musical contributions were made by Canadians. Big thanks also to Josh LaRocque, Adam Riggall, Jarrod Lewis-Peacock, Yelena Guller, Bornali Kundu, Alex Shackman, and Mike Starrett. You have each aided my scientific progress in more ways than you know. Go Badgers!

Finally, I would like to extend a special thanks to my friends and family. Mom and Dad, thank you for always putting us first and encouraging me to follow my dreams no matter where they take me. Your love and support mean the world to me. Mom, thank you for being my first and best teacher; those summers in Mom school have really paid off. Dad, thank you for diving headfirst with me into any activity I decided to take on, whether that was soccer, mock trail, or scouts. My Grandma Betty deserves a lot of credit for sparking my interest in science at a young age. She spend many hours entertaining me with stories about Grandpa Sutterer's experiments working with the G-force centrifuge and computers the size of our living room. She also ensured

that I made it this far by, saving me from “near starvation” at Disney world when I was a child. To my brother Matt, thank you for all the years of support. Including helping me decide on a research lab that fit my interests (despite our inability to agree on which topics in neuroscience are most interesting), editing my applications for grad school, and letting me tag along with you and your friends both when we were growing up and as an adult at conferences. You are a world class big brother and scientist. Special thanks as well to all of my friends for the late-night conversations, soul searching, and weekend excursions that made the past years so enjoyable. Last and certainly not least, I would like to thank my partner Siena. You have been with me through all of this, from participating in my senior thesis study, to moving across the country three times so I can continue to pursue this crazy career path. Thank you for listening to me ramble on endlessly about testing effects, inverted encoding models, alpha-band activity, and for always encouraging and believing in me. Most importantly thank you for helping me find joy each day along the way.

ABSTRACT

Our ability to maintain and later retrieve detailed visual memories is an everyday experience, and completing almost any task requires us to shuttle this information between working and long-term memory. However, the mechanisms by which the brain supports detailed, online memory representations and the temporal dynamics of these representations remain poorly understood. Another open question is the degree to which these representations evolve from short to long delays. Across three experiments, we tested how and when the brain represents feature specific memories –such as the color or location of an object– during both working memory and retrieval from long-term memory. These experiments employed a combination of continuous report measures and modeling of multivariate patterns of electroencephalography (EEG) activity to identify and characterize the patterns of oscillatory activity that encode the online maintenance of feature-specific memory representations.

In Experiment1, we explored which frequencies of EEG oscillations support online spatial working memories, and then use this activity to test the long-standing question of whether or not multiple items can be simultaneously maintained in working memory. In Experiment 2, we leveraged continuous report measures to examine how retrieval practice affects the probability that precise color memories can be retrieved and the precision with which that information can be retrieved from long-term memory. Finally, in Experiment 3, we combine these two approaches in order to test whether EEG activity can be used to obtain a time-resolved measure of long-term memory retrieval and to explore the degree to which patterns of rhythmic EEG activity at encoding are reinstated at retrieval. Taken together these experiments highlight the broad potential for using continuous report measures, EEG, and multivariate analysis to characterize memory function.

CHAPTER 1. INTRODUCTION

A hallmark of episodic memory is the phenomenon of mentally re-experiencing the details of past events. Imagine that you are on a hike and have reached an unfamiliar fork in the path and need to consult a map. One way to determine where you need to go is to consult the map in your back pack, and to then briefly hold the landmarks from the map in mind while you determine how the trail ahead best fits the direction you want to go. This operation of holding recently observed information in mind to guide behavior is an example of working memory (WM). However, if you forgot your map in the car, you would also likely be able to call up a memory of the map you studied at the trailhead before embarking on your hike, and use that mental representation to guide your decision. This is an example of relying on long-term memory (LTM) to call up a detailed visual memory that has not been in mind for minutes or hours in order to accomplish your goal of successfully navigating your hike. Despite being a commonplace operation, both how the brain accomplishes rapid shuttling of information between an active state (WM) and a passive state (LTM) and how representations change across delays are poorly understood. Over the course of my doctoral program I have focused on developing and using computational analyses to track online memories, and in my dissertation I will describe three sets of experiments in which we leverage precise measurements of basic visual features and their neural correlates to understand how and when the brain represents feature specific memories.

Features are the basic building blocks of visual memory

The example from the preceding paragraph relies on a complex mental representation; however, this representation can be broken down into individual features that make up the whole.

In this way, we can think of individual features – such the color, location, and orientation of objects– as the basic building blocks of visual memory. Crucial steps towards understanding memory function are mapping out the behavioral limits of memory performance for basic features and assessing how the brain supports these representations. Continuous report tasks are a useful tool for testing the behavioral limits of feature memory. In such tasks, observers memorize and reproduce the precise value of a single memorized feature like the color of an object. Beyond testing the resolution of memory for basic visual features, adoption and adaptation of this continuous report approach has proven a useful for answering broad questions about memory and cognition.

The adoption of this approach has been particularly fruitful for advancing the debate about whether the limited nature of working memory is the result of a limit in the total number of items that can be stored concurrently (Luck & Vogel, 2014; Wilken & Ma, 2004; Zhang & Luck, 2008). Specifically, application of this approach provided behavioral evidence that capacity limits in memory are not a unitary measure, instead they reflect both the quality and quantity of remembered information. In a seminal study, Zhang and Luck (2008) had observers remember arrays of colored squares varying the number squares to be memorized (1, 2, 3, or 6) across trials. After a delay, observers reported the precise color one of the squares by choosing the target color from continuous color space. The authors reasoned that on any trial observers, either would remember the color of the probed square and report the correct color with some small amount of error, or observers would not remember the color and guess randomly producing a uniform distribution of guesses across trials. A mixture model was applied to the response error distribution for each subject. This mixture model included parameters to quantify the proportion of random guesses (uniform responses) and the precision of correct responses for each observer.

Results of this analysis revealed that observers could remember the color of 3-4 squares with relatively high memory precision, but retained no information about the additional items when an array exceeded this capacity. Additionally, they found that memory precision decreased as memory load increased monotonically from one to three items until capacity was reached, after which memory precision remained constant. This pattern of results suggests that memory capacity is limited at 3 to 4 objects or discrete slots, and that resources for remembered items may be shared among items.

Alternative accounts argue that this same pattern of behavior can be described by a model in which WM is well described as a continuous resource with no discrete capacity limit (van den Berg, Shin, Chou, George, & Ma, 2012). In this view, memory performance varies randomly trial to trial and memory precision increases in variability as memory load increases. Critically, this account allows for extremely imprecise memories that mimic guessing at high set sizes. Both accounts describe the existing behavior well (van den Berg, Awh, & Ma, 2014), thus the debate continues. One promising approach has been to directly observe brain activity during the delay period of visual working memory tasks. Indeed, observations that neural signals during the delay period of a working memory task increase with memory load and level off when at memory capacity provided support for the concept of discrete capacity limits in WM (Todd & Marois, 2004; Vogel & Machizawa, 2004). However recent work has suggested new models that may also account for these observations (Bays, 2018), so the debate continues.

A related but orthogonal question to this slots vs resources debate revolves around why memory precision decreases as memory load increases, even when memory load is within capacity (1 to 3 items). Evidence from neuroimaging suggests that WM can be decoded from persistent patterns of neuronal activity (Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh,

2009) and that the quality of memory representations decreases with increasing memory load (Emrich, Riggall, Larocque, & Postle, 2013; Sprague, Ester, & Serences, 2014, 2016). However, the existing data can be explained by opposing mechanisms. On the one hand, the predominant view is that all items are concurrently represented via persistent patterns of neural activity, and the decrease in representation quality reflects competition between these items (Bays, 2014; Franconeri, Alvarez, & Cavanagh, 2013). On the other hand, recent "activity-silent" models of working memory (Stokes, 2015) have challenged the view that all items maintained in WM are represented by persistent patterns of neural activity. Under this view, it is possible that not all items are represented concurrently during the delay with some items represented via "passive" changes in synaptic weights. In Chapter 2, we will combine a continuous report task with modeling of rhythmic EEG activity to adjudicate between these competing accounts.

This continuous report approach has quite recently been applied to the study of long-term memory (Brady, Konkle, Gill, Oliva, & Alvarez, 2013) and is already providing new evidence on the nature of feature representations in LTM. For instance, recent evidence suggests that the probability that a feature associated with an item is remembered can be modeled separately from the quality of the remembered representation (Harlow & Donaldson, 2013; Harlow & Yonelinas, 2014), and fMRI work confirms that different areas of the brain correlate with these two factors (Richter, Cooper, Bays, & Simons, 2016). Moreover, Brady and colleagues directly compared the precision with which colors can be maintained between working memory and long-term memory and found that working memory and long-term memory appear to have the same upper limit in the precision with which colors can be remembered. Finally, evidence suggests that probability of retrieval and memory quality are differentially affected by attention at encoding (Fan & Turk-Browne, 2013).

In sum, these findings suggest that carefully testing memory for basic features conveys several advantages. First, such an approach allows for careful modeling of memory performance into component processes to inform cognitive theory. In Chapter 3, we will apply this modeling approach to determine how memory retrieval differentially affects the probability that a memory is retrieved and the precision with which memories are retrieved. Second, this approach reduces ambiguity about the content of retrieved memories, even across longer delays, which is particularly advantageous for interpreting neural correlates of memory activity. Thus, the discovery and study of neuronal activity that tracks these fine-grained memory representations will also provide a powerful platform for understanding memory function. In Chapters 2 and 4 we will combine this behavioral approach with EEG recordings in order to track precise location specific memories in working and long term memory.

Distinguishing between working memory and long-term memory

The idea that there is a fundamental difference between memory for information that was encountered seconds ago and memory that has been stored over long delays has been around since the beginning of the programmatic study of psychology. Indeed, the prolific William James coined these two states respectively as primary and secondary memory (James, 1890). These terms are still used by some researchers today, although the field typically prefers the distinction of working memory and long-term memory. While intuitive constructs, a central question for memory researchers is how these two types of memory are related, and an ongoing debate revolves around if they are separate memory systems.

On one end of the debate, WM and LTM are defined as separate memory systems. According to this view WM is defined as a capacity limited system for maintaining a small set of representations that are behaviorally relevant, and LTM is defined as an unlimited capacity

system for maintaining information (Atkinson & Shiffrin, 1968). Under these multi store models, information enters long-term memory via working memory, and long-term memories are then retrieved back into working memory when they are needed for behavior. The strongest evidence supporting the view that WM and LTM are separate systems comes from neuropsychology. The key observation is that patients with damage to their medial temporal lobe have impaired performance on LTM tasks but normal performance on WM tasks (Baddeley & Warrington, 1970; Scoville & Milner, 1957), while patients with damage to their perisylvian cortex exhibit normal performance on LTM tasks but impaired performance on WM tasks (Shallice & Warrington, 1970). However, recent unit recording evidence in humans suggests that MTL may be more involved in WM (Kamiński et al., 2017; Kornblith, Quiroga, Koch, Fried, & Mormann, 2017) than early studies suggest, leaving this evidence far from conclusive.

On the other end of the debate is the framework that WM and LTM are a part of the same system but that different processes are engaged at various stages of perception and memory formation (Craik Fergus & Lockhart, 1972). Numerous models formalizing this process framework conceptualize working memory as an activated portion of long-term memory rather than as a separate system (Cowan, 1995; McElree, 2006; Oberauer, 2002). This concept maps nicely onto a long standing idea in cognitive neuroscience that WM is supported by persistent neural activity (Funahashi, Bruce, & Goldman-Rakic, 1989; Fuster & Alexander, 1971; Luck & Vogel, 2014), although this assumption has recently been questioned (N. S. Rose et al., 2016; Stokes, 2015; Wolff, Jochim, Akyurek, & Stokes, 2017). I will approach the distinction between working and long-term memory through this lens, and explore how similarities and differences in patterns of recurrent activity across delays map onto existing cognitive models throughout my dissertation.

Evidence for representation overlap in working and long-term memory

An interesting implication of these embedded process models (Cowan, 1999) is that the active representational bases of working memory and long-term memory should be identical (Jonides et al., 2008), and many accounts also suggest that activity during perception is reinstated at retrieval. In line with this account, numerous fMRI studies have found increased BOLD activity in sensory regions during retrieval (Danker & Anderson, 2010; Nyberg, Habib, McIntosh, & Tulving, 2000; Wagner, Shannon, Kahn, & Buckner, 2005) along with evidence that activity in these same regions during encoding is related to subsequent memory. More recently, multivariate imaging approaches have shown that patterns of activity observed during encoding are reinstated during memory retrieval (Kuhl, Rissman, Chun, & Wagner, 2011; Norman, Polyn, Detre, & Haxby, 2006; Polyn, Natu, Cohen, & Norman, 2005).

The hemodynamic response inherent to fMRI obscures the temporal dynamics of such retrieval processes. Electroencephalography (EEG) provides enhanced temporal resolution, capable of measuring the time course of retrieval-related neural activity. However, studies investigating reinstatement with EEG or magnetoencephalography (MEG) have primarily been focused on the reinstatement of categorical information such as the encoding task (Johnson et al., 2015) or the category of a paired associate (Jafarpour, Fuentemilla, Horner, Penny, & Duzel, 2014; Morton et al., 2013; Morton & Polyn, 2017; Waldhauser, Braun, & Hanslmayr, 2016). Interpretation of these category level representations is limited because they only reveal that some process was consistent across items in each category (Larocque, Riggall, Emrich, & Postle, 2017). Thus, an open question is whether the patterns of activity that support encoding and maintenance of feature-specific representations are reinstated at retrieval. To test this question we can simply assign a precise location value to each remembered item. This approach allows

for a direct comparison of the patterns of activity that represent memoranda across time scales. In Chapter 4, we will test the extent to which the oscillatory patterns of activity engaged during encoding and maintenance of precise location memories are reinstated at memory retrieval.

Topography of alpha-band activity tracks the precise location held in working memory

Neuronal oscillations have been suggested as a mechanism by which populations of neurons can be coordinated to represent specific stimulus values in memory (Canolty & Knight, 2010; Fell & Axmacher, 2011; Singer, 1999; Watrous, Fell, Ekstrom, & Axmacher, 2015). In support of this hypothesis, many studies have shown that alpha-band (8–12 Hz) oscillations measured with electroencephalography (EEG) play a critical role in the selection and storage of information in the brain. For example, many studies have demonstrated that alpha-band power at posterior electrodes decreases in electrodes contralateral to the focus of attention (Sauseng et al., 2005; Thut, 2006). Moreover, evidence suggests that the topography of alpha-band power tracks more than the attended hemifield, finding that alpha-band power tracks the precise focus of spatial attention (Foster, Sutterer, Serences, Vogel, & Awh, 2017; Rihs, Michel, & Thut, 2007; Samaha, Sprague, & Postle, 2016; Treder, Bahramisharif, Schmidt, van Gerven, & Blankertz, 2011; Worden, Foxe, Wang, & Simpson, 2000). Finally, Foster et al., (2017) showed that that not only does alpha-band power track the precise attended location, but also tracks the time course of covert orienting in a spatial attention task.

Consistent with work demonstrating functional overlap between attention and WM (Awh & Jonides, 2001; Awh, Vogel, & Oh, 2006; Chun, 2011) studies have demonstrated similar changes in the topography of alpha-band power during working memory. Specifically, the topography of alpha-band power also co-varies with the hemifield maintained in WM (Grimault et al., 2009; Medendorp et al., 2007; van Dijk, van der Werf, Mazaheri, Medendorp, & Jensen,

2010). Foster et al., (2016) used an inverted encoding model (IEM) to provided clear evidence that the topography of alpha-band power tracks the precise location held in memory rather than just the remembered hemisphere. Taken together these results suggest that alpha-band oscillations play a key role in spatial cognition.

In addition to tracking the location of a single to-be-remembered item, alpha power suppression at posterior electrodes reflects the number of items maintained in working memory and levels off between 3 and 4 items (Fukuda, Mance, & Vogel, 2015; Sauseng et al., 2009). This observation suggests that the topography of alpha-band power may prove useful for tracking multiple locations in working memory. In Chapter 2 of the presented work, I test if multiple locations are tracked simultaneously by the topography of alpha-band power. Embedded process models predict that the activity coding for representations during WM maintenance should be reinstated during memory retrieval. We tested this hypothesis in Chapter 4 of the presented work by investigating which frequency bands support the retrieval of precise spatially specific representations and the extent to which patterns of activity observed at encoding are reinstated at retrieval.

Summary and outline

Completing almost any task in everyday life requires us to shuttle information between working and long-term memory. Central questions for the field of memory research are what neuronal representations contribute to active memory representation, when long-term memories are retrieved into an active state, and how neural codes for these representations evolve over long delays. In the chapters that follow, I will review an approach in which we use an inverted encoding model (IEM) and multivariate patterns of electroencephalography activity to track the maintenance of online spatial representations. By combining this approach with

continuous measures of memory performance, we will ask three open questions about the nature of online memory representations. In Chapter 2, we will attempt to answer the long-standing question: how many items can be held in working memory simultaneously? In Chapter 3, we will examine how taking a test affects the probability that information is retrieved relative to the probability that it is later retrieved. Finally, in Chapter 4, we will test the extent to which the frequencies and patterns of information that represent information during memory encoding are reinstated during memory retrieval.

CHAPTER 2. DELAY-PERIOD ACTIVITY ENCODES MULTIPLE ITEMS STORED IN WORKING MEMORY

Introduction

Working memory (WM) is an "online" memory system that maintains information in a readily accessible state. The term "online" refers both to the accessibility of working memory representations for behavior and also to the putative neural mechanisms by which representations are maintained. For example, the content of WM can be decoded from patterns of persistent neural activity in both human neuroimaging (Harrison & Tong, 2009; Serences et al., 2009) and unit recordings in non-human primates(Funahashi et al., 1989; Funahashi, Chafee, & Goldman-Rakic, 1993). Furthermore, in line with the known declines in behavioral performance as WM load increases (Fukuda, Awh, & Vogel, 2010; Luck & Vogel, 1997), the selectivity of representations decoded from persistent activity also declines with increased memory load (Buschman, Siegel, Roy, & Miller, 2011; Emrich et al., 2013; Matsushima & Tanaka, 2014; Sprague et al., 2014, 2016). For example, Sprague et al. (2014) used an inverted encoding model (IEM) to reconstruct spatial representations of remembered locations from voxel-wise patterns of activity measured with fMRI, and found that selectivity of spatial representations decreased when memory load was increased from one item to two items. Likewise, unit recording in non-human primates (Buschman et al., 2011) have shown that individual units carry less information about each remembered stimulus with greater WM load. Thus, because persistent neural activity tracks mnemonic content and mirrors the declines in behavior observed with increasing mnemonic load, this activity has been thought to play a central role in WM storage (Sreenivasan, Curtis, & D'Esposito, 2014).

Despite robust parallels to behavior findings, prior work does not firmly establish that multiple items are concurrently stored in an *active state* (i.e. represented by persistent neural activity). The dominant view has been that reduced selectivity of neural representations with greater memory load occurs because competition between concurrently stored representations degrades the fidelity of those representations (Bays, 2014; Franconeri et al., 2013). However, recent "activity-silent" models of working memory (Stokes, 2015) have challenged the view that all items maintained in WM are supported by a persistent patterns of neural activity (e.g. continued firing of neurons). Instead these models propose that activity-silent memory mechanisms (e.g. rapid changes to synaptic weights) can support the short-term retention of information (Lewis-Peacock, Drysdale, Oberauer, & Postle, 2012; Rose et al., 2016; Stokes, 2015; Wolff et al., 2017). In this view, when multiple items are stored in WM, they need not be concurrently represented by persistent activity. Instead, each item may transition between active and silent states, with only a single item in an active state at any given time. Consistent with such a model, recent work has found that when two locations must be attended, these locations are sampled sequentially (Busch & Van Rullen, 2010; Fiebelkorn, Saalmann, & Kastner, 2013; Landau & Fries, 2012). Given past work positing functional overlap between spatial WM and spatial attention (Awh & Jonides, 2001; Awh et al., 2006; Chun & Turk-Browne, 2007), sequential representation may also underpin the maintenance of multiple locations in spatial memory. Such a *switching model* (in which only 1 item is actively represented) also predicts an apparent decline in the fidelity of neural representations when mnemonic load is increased, because typical analyses aggregate data across multiple trials. Thus, if an item is represented for a smaller portion of the delay period when load increases, this could mimic the effects of competition between concurrently stored items. Thus, the goal of the present work was to

provide a more decisive test of whether persistent neural activity enables concurrent storage of multiple items in WM.

To test whether multiple memoranda are concurrently represented by persistent activity, we used an inverted encoding model in conjunction with EEG measurements of alpha-band activity to track spatial representations stored in WM with high temporal precision. Oscillatory alpha-band (8–12 Hz) activity is a component of the EEG signal that tracks spatial positions maintained in WM (Foster, Bsales, Jaffe, & Awh, 2017; Foster, Sutterer, Serences, Vogel, & Awh, 2016). To compare predictions of a *switching model* with a *multi-item* model of persistent activity, we compared the selectivity of alpha-band representations for one-item versus two-item arrays. To preview our results, we found that the spatial selectivity of these alpha-band representations was lower in the two-item condition than in the single-item condition, consistent with previous work. However, novel simulations revealed that differences in selectivity between the one- and two-item conditions could not be explained by a switching model in which only a single-item is represented in the ongoing EEG signal. Instead, spatially-specific alpha-band activity provides clear evidence for the simultaneous storage of at least two distinct representations in WM.

Materials and Methods

Participants

Forty-one volunteers participated in the experiment for monetary compensation ($15/hr). Participants were between 18 and 35 years old, reported normal color vision and normal or corrected-to-normal visual acuity, and provided informed consent according to procedures approved by the University of Chicago Institutional Review Board. Participants were excluded from analyses if fewer than 450 trials remained in either the one- or two-item condition after

discarding trials with artifacts. Artifact number exclusion criteria were set during data collection, but before the data were analyzed.

Eight subjects were excluded because of excessive artifacts. Data collection was terminated early for four subjects because of excessive artifacts, and for one subject because of a fire alarm. The final sample included 28 subjects with an average of 581 (*SD* = 75) trials for one-item trials and 596 (*SD* = 69) trials for two-item trials.

Apparatus and Stimuli

We tested participants in a dimly lit, electrically shielded chamber. Stimuli were generated using Matlab (MathWorks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented on a 24" LCD monitor (refresh rate: 120 Hz, resolution: 1024 × 1920 pixels) at a viewing distance of ~100 cm. Stimuli were rendered against a gray background. The spatial WM task (Figure 2-1) required participants to remember the angular location of one or two sample stimuli. Each stimulus was a blue or green circle (equated for luminance; 0.2° in diameter) centered 4° of visual angle from the central fixation point (0.2° in diameter). For one-item trials, blue or green was randomly chosen; for two-item trials, both blue and green were presented. The color of the probed item was pseudo-randomized across trials and conditions such that each color was probed on 50% of trials for each condition. The angular position of each stimulus around the fixation point was drawn from eight position bins, each spanning a 45° wedge of angular positions (bins were centered at 0°, 45°, 90°, and so forth), with jitter added to cover all 360° of possible locations to prevent categorical coding of stimulus-location. On two-item trials, the position bins that each stimulus 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 each position independently. When both stimuli

occupied the same position bin, their exact position within the bin was constrained so that the two items were one stimulus width apart.

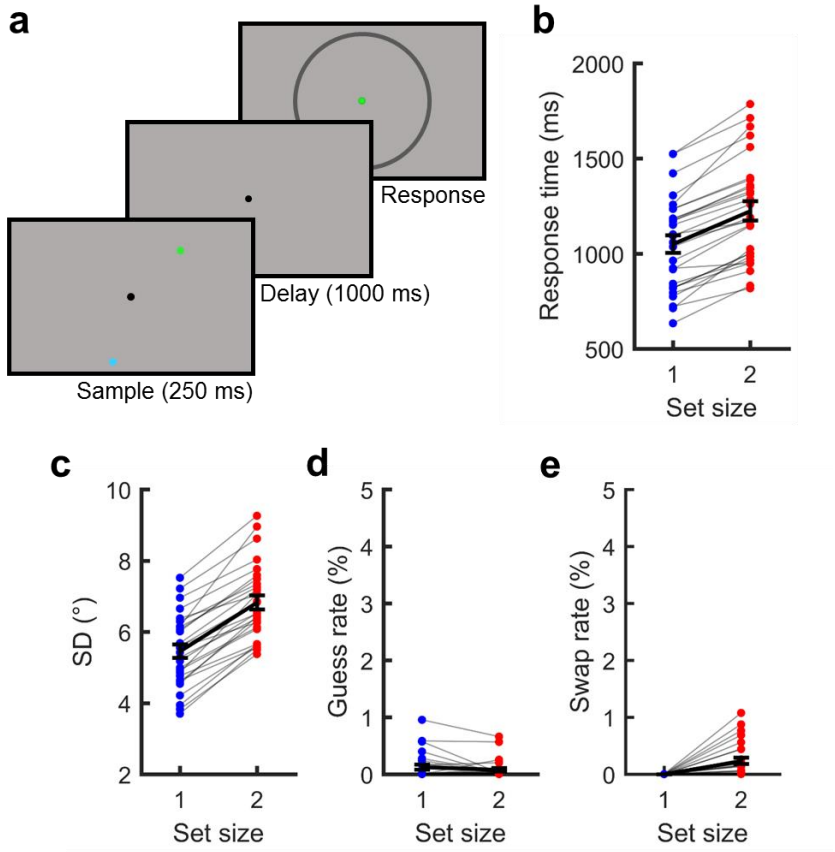

Figure 2-1. Experimental task and behavior.
A. Observers saw a brief sample display (250 ms) that contained one or two colored dots. After a delay period (1000 ms), the fixation point changed color (blue or green) to indicate which item should be reported. Observers reported the angular position of the cued item as precisely as possible by mouse-click on the perimeter of a rim. **B.** Median response time (ms) as a function of memory set size (one vs. two items). Blight grey lines represent individual subjects. Black lines represent the mean. Error bars represent ±1 SEM. **C-E**. Parameter estimates obtained by fitting a three-component mixture model (Bays et al., 2009) to response errors as a function of memory set size. *SD* (c) reflects precision of responses (with higher values indicating worse precision), *pGuess* (d) estimates the probability that the observer produced a random response (i.e., a guess), and *pSwap* (e) estimates the probability that uncued item was misreported instead of the cued-item. Note that *pSwap* is necessarily zero for the single-item condition.

Task procedure

Participants performed a spatial delayed-estimation task (see Fig 1) in which they reported the angular position of a stimulus around the fixation point. Participants initiated each trial with a spacebar press. The trial began with a fixation display lasting between 500 and 800 ms, followed by a memory array that comprised one or two colored circles presented for 250 ms. Subjects were instructed to remember the location of each colored circle as precisely as possible. A 1000 ms delay period, during which only the fixation point remained visible, followed the memory array. After the delay period, a cursor appeared and the fixation dot turned either blue or green to indicate which stimulus should be reported. Participants were instructed to report the remembered location of the probed item. To respond, participants used a mouse to click on the perimeter of a probe ring (8° in diameter, 0.2° thick). Before starting the task, participants completed a brief set of practice trials to ensure that they understood the task.

Electrophysiology

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, F8, F3, F4, Fz, FC5, FC6, FC1, FC2, C3, C4, Cz, CP5, CP6, CP1, CP2, P7, P8, P3, P4, Pz, T7, T8, FT9, FT10, O1, O2, Oz. 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. Data were referenced online to the right mastoid and re-referenced offline to the algebraic average of the left and right mastoids. Eye movements and blinks were also monitored using electrooculogram (EOG). We collected EOG 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 Brain Vision Recorder (Brain Products, Munich, German) running on a PC. Impedance values were kept below 10 kΩ.

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 500 Hz, and data were obtained in remote mode (without a chin rest). We obtained usable eye-tracking data for 19 out of 28 participants.

Artifact rejection

We visually inspected the segmented EEG data for artifacts (excessive muscle noise and skin potentials), and inspected EOG for ocular artifacts (blinks and eye movements). We discarded trials contaminated by artifacts. Data from one or two electrodes were discarded for four participants because of excessive noise (i.e. a broken electrode). The discarded electrodes for each participant are respectively were: T7; F3; CP6 and C4; and P8. For subjects with usable eye tracking data, we also inspected gaze data for ocular artifacts. For the analysis of gaze position, we further excluded trials in which the eye tracker was unable to detect the pupil, operationalized as any trial in which the horizontal gaze position was more than 15° from fixation or the vertical gaze position was more than 8.5° from fixation.

Removal of trials with ocular artifacts was effective. Maximum variation in grand-averaged HEOG waveforms by remembered location bin at any time point was < 4 μV for both one and two-item trials. Thus eye movements corresponded to variations in eye position of < .25° of visual angle (Lins). Analysis of the subset of participants (19) for whom we were able to obtain reliable gaze position data corroborates the HEOG data obtained from all participants.

Variation in grand-average horizontal gaze position as a function of remembered location at any time point was < .12° of visual angle for both one and two-item trials. Variation in grand-average vertical gaze position by remembered location at any time point was < .11° of visual angle for both one and two-item trials. For comparison, HEOG for these participants showed a < 3.2 μV maximum variation at any time point which corresponds to < .21° of visual angle.

Time-frequency analysis

To calculate frequency specific activity at each electrode we first band-pass filtered the baselined raw EEG data using EEGLAB ('eegfilt.m', see Delorme and Makeig, 2004). For alpha band analyses, the data were band-pass filtered between 8 to 12 Hz, consistent with our prior work (Foster et al., 2016). For our exploratory analysis of a broad range of frequencies, we band-pass filtered the data at 1Hz intervals (4–50 Hz), filter order:

$$3 \times \frac{sampling\ rate}{low-pass\ cutoff}$$

We then applied a Hilbert transform (MATLAB Signal Processing Toolbox) and squared the complex magnitude of the complex analytic signal for each trial to calculate instantaneous power before averaging across trials.

Inverted Encoding Model

Following our prior work (Foster et al., 2016), we reconstructed spatially selective channel-tuning functions (CTFs) from the multivariate topographic distribution of oscillatory power across electrodes. We assumed that the power at each electrode reflects the weighted sum of eight spatially selective channels (which we assume reflect the responses of neuronal populations), each tuned for a different angular location (Brouwer & Heeger, 2009; Foster et al., 2016; Sprague, Saproo, & Serences, 2015; Sprague & Serences, 2013). 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 location bins. These 8 location bins each spanned 45° and were centered on 0°, 45°, 90°, and so on.

An IEM routine was applied to each time point in the alpha-band analyses and to each time-frequency point in the time-frequency analyses. We partitioned our data into independent sets of training data and test data (for details see the section, "Assigning trials to training and test sets"). 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 measurements) 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 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 × 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 (i.e., 0° on the "Channel Offset" axis of Figure 2-2a) by aligning the estimated channel responses to the channel tuned for the cued/target location to yield the CTF averaged across the eight remembered locations.

Finally, because the exact contributions of each spatial channel to each electrode (i.e., the channel weights, *W*) varies 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 (Foster et al., 2016; Foster, Sutterer, Serences, Vogel, et al., 2017; Sprague et al., 2015).

Assignment of trials to training and test sets.

Artifact-free trials were partitioned equally into three independent sets to be used as training and test data for the IEM procedure (see Inverted Encoding Model). We down-sampled the data so that each set contained an equal number of trials, and that each location bin within a set also contained the same number of trials. For each of these sets we averaged power across trials for each location bin. We used a cross validation routine such that two sets of estimated power served as the training data and the remaining set served as the test data. We applied the

IEM routine using each of the three matrices as test data, and the remaining two matrices as training data. The resulting CTFs were averaged across each test set.

For analyses in which we examined how location selectivity varied as a function of memory load (i.e. comparison of one item and two item selectivity), we first down-sampled to equate the number of trials assigned from each location across conditions. After completing this additional step, we equated trials across sets and bins in the same manner described above. Finally, we employed the same training procedure described above (2/3 of the total data), but split the final test set into our comparisons of interest. Thus, we used the same training data for both conditions and only the test data varied for each comparison. Using a training set that is an equal combination of both conditions has two advantages. First, common training allowed us to utilize more data to train the IEM improving our signal to noise ratio. Second, common training ensures that the training set matched both conditions equally well, allowing for the straightforward conclusion that any observed difference in selectivity between conditions is the result of true differences in spatial selectivity and not the choice of training set. For the simulation analysis in which we compared two-item trials to one-item trials with 50% noise, we first down-sampled and partitioned trials into sets as described above. Next, we randomized the labels of 50% of the trials in each set of the one-item data. Shuffling the labels in this manner ensured that we preserved a balanced number of labels corresponding to each position for each set. Finally, we conducted the rest of analysis the same way we conducted the analysis comparing the selectivity of one-item trials to two-item trials (see above).

For the analysis in which we compared the frequencies that encode spatial representations in the one and two item conditions we trained and tested the model for each condition separately (train 2/3; test 1/3). We ran the model for each condition separately because we wanted to

maximize the sensitivity of the analysis for observing differences in frequency, and we were not interested in comparing CTF selectivity across conditions.

For analyses in which we assessed whether the similarity of multivariate patterns representing one-item and two-item trials, we used the same procedure described above except that we trained the IEM on two sets of one item data with the number of trials per bin equated as described above, and tested on a single set of the two item data.

Resampling random assignment

To avoid spurious results due to the random assignment of trials, we repeated each analysis multiple times with a different random assignment of trials. For all analyses where we focused solely on alpha-band power, we conducted 50 iterations per time point. For the 4–50 Hz time-frequency analysis (which is a time consuming procedure), we conducted 10 iterations per time frequency point, given the computational time needed for this analysis. In order to decrease computation time further for the 4–50 Hz time-frequency analysis, we down sampled the data matrix of power values to one sample every 20 ms. We down sampled after calculating power so that down sampling did not affect our calculation of power. The data matrix was not down-sampled for analyses restricted to the alpha band.

Statistics

Modeling of response errors. Response error was measured as the number of degrees between the presented angular location and the reported angular location. Errors ranged from 0^0 (a perfect response) to $\pm 180^0$ (a maximally imprecise response). To quantify performance we fit a mixture model to the distribution of response errors for each participant using MemToolbox (Suchow, Brady, Fougnie, & Alvarez, 2013). For one-item trials, 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 0^0), corresponding to trials in which the sample location was remembered, and a uniform distribution, corresponding to guesses in which the reported location was random with respect to the sample location. 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 two-item trials we fit a mixture model that also included an additional von Mises component centered on the location of the un-probed item, corresponding to trials in which participants mistakenly reported the location of the un-probed item (i.e. swaps, Bays, Catalao, & Husain, 2009). We obtained maximum likelihood estimates for the same parameters as in one-item trials, with one additional parameter (P_s), which reflects the probability of swaps.

Calculating CTF selectivity. 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. To generate our null distribution, we randomly shuffled the remembered location labels in each training/test set so that the labels were random with respect to the observed responses at each electrode. We then repeated 1000 iterations of this randomization procedure to obtain a null distribution of *t* statistics at each time point.

Next, we employed a nonparametric cluster approach that corrects for multiple comparisons by taking into account auto-correlation in time and frequency (Cohen, 2014; Maris & Oostenveld, 2007). Specifically, we applied a *t*-value threshold corresponding to $p < .05$ ($t(27) = 1.703$) to identify clusters of pixels (time and frequency analysis) or adjacent time points (alpha only analysis). At the same time, we applied the same threshold to each permutation and

calculated the largest summed *t*-statistic for any cluster in the permutation, resulting in a distribution of maximal summed *t*-statistics for our permuted null distribution. Finally, the sizes of the significant clusters of the non-permuted data were thresholded such that only clusters larger than the 95th percentile of the permuted distribution were considered reliable (Type 1 error less than .05). Therefore, our cluster test was a one-tailed test, corrected for multiple comparisons.

Resampling test.

When comparing CTF slope between conditions, we used a non-parametric resampling procedure across participants (Efron & Tibshirani, 1993). We resampled each participant with replacement 100,000 times. Then, we calculated the number of these resampling iterations in which the differences were different from zero (e.g. difference > 0 and difference < 0) and doubled the smaller p value. Thus this was a two-tailed test. In cases where no iterations were different from zero in the significant direction, we report the p values as $p < .00001$. We deemed results to be reliably above chance if $p < .05$.

Results

Behavior

Observers performed a spatial working memory task (Figure 2-1a). On each trial, observers remembered the spatial position of one or two colored dots. We found that median response times (Figure 2-1b) were slower for two-item trials (M = 1050 ms, SD = 245) than one-item trials (M = 1224, SD = 271), $t(27) = 12.33$, $p < 0.001$. In line with past work (Luck and Vogel, 1997; Wilken and Ma, 2004; Zhang and Luck, 2008; Bays et al., 2009), memory performance declined as memory load increased from one to two items. We analyzed the recall data using a three-component mixture model (Bays et al., 2009) to estimate mnemonic precision

(*SD*, higher values indicate lower precision), the probability that a stimulus was forgotten (*pGuess*), and the probability of reporting a non-target item (*pSwap*). We found that mnemonic precision was worse (i.e., *SD* was higher, Figure 2-1c) when participants maintained two items (M = 6.82°, SD = 1.05) than when they maintained one item (M = 5.45, SD = 1.01), $t(27)$ = 15.66, $p < 0.001$. We saw no reliable difference in the rate of guessing (Figure 2-1d) between one-item (M = 0.12%, SD = 0.24) and two-item (M = 0.08%, SD = 0.16) trials, $t(27)$ = -1.59, p = 0.12. Finally, observers' rates of misreporting the location of the non-target item (Figure 2-1e) on two-item trials (M = 0.23%, SD = 0.30) was reliably greater than zero, $t(27)$ = 4.10, $p <$ 0.001. Note that the combined rate of guessing and swapping were very low (less than 1% for almost all subjects). Thus, the primary change in behavior with memory load was the slowing of response times and reduction in precision of responses.

Alpha-band representations of space degrade with increased memory load

To test how online representations change with increased memory load, we examined oscillatory alpha-band (8–12 Hz) activity, which encodes spatial representations that are maintained in WM (Foster et al., 2016, 2017). We used an inverted encoding model (Brouwer & Heeger, 2009; Sprague et al., 2015) to reconstruct spatial representations encoded by alpha-band activity (Foster et al., 2016). Our encoding model assumed that alpha-band power measured at each scalp electrode reflects the activity of a number of spatially tuned channels (or neuronal populations), each tuned for a different position in the visual field. In a training phase, we estimated the relative contributions of the spatial channels to each electrode on the scalp (called the "channel weights") using a subset of trials during the spatial WM task. Then, in a test phase, using an independent subset of trials, we used these channel weights to estimate the responses of the spatial channels given the pattern of alpha-band power across the scalp. The resulting profile

of responses across the spatial channels (called channel-tuning functions or CTFs) reflects the spatial selectivity of alpha-band activity measured by EEG. We performed this analysis at each time point throughout the trial, which allowed us to test whether active spatial representations were maintained throughout the delay period.

To examine how WM load affects alpha-band representations of the remembered positions, we reconstructed CTFs for the both the one- and two-item conditions (see Materials and Methods). We observed a clear spatially selective CTF, with a peak response in the channel tuned for the remembered location (a channel offset of 0°), which persisted throughout the delay period in both the one-item and two-item conditions (Figure 2-2a and 2-2b). Figure 2-2c shows the spatial selectivity of the CTFs seen in each condition across time (measured as CTF slope, see Materials and Methods). Cluster-based permutation tests revealed that spatial selectivity of alpha-band CTFs was reliably above zero throughout the delay period for both the one- and two-item conditions ($p_{,} < .05$, corrected for multiple comparisons; see markers at the top of Fig 2c). Next, we compared CTF selectivity for the one-item and two-item conditions. A resampling test confirmed that delay-period CTF selectivity (averaged from 250 to 1250 ms after stimulus onset) was reliably lower ($p < .001$) for two-item trials ($M = .059$, $SD = .033$) than for one-item trials ($M = .087$, $SD = .046$). Importantly, we also found that this difference was reliable ($p = .004$) during a late window (800–1250 ms) when alpha-band representations are unlikely to be affected by stimulus-driven activity. Thus, as memory load increases, there is a decline in spatially-selective alpha-band activity that tracks the stored locations.

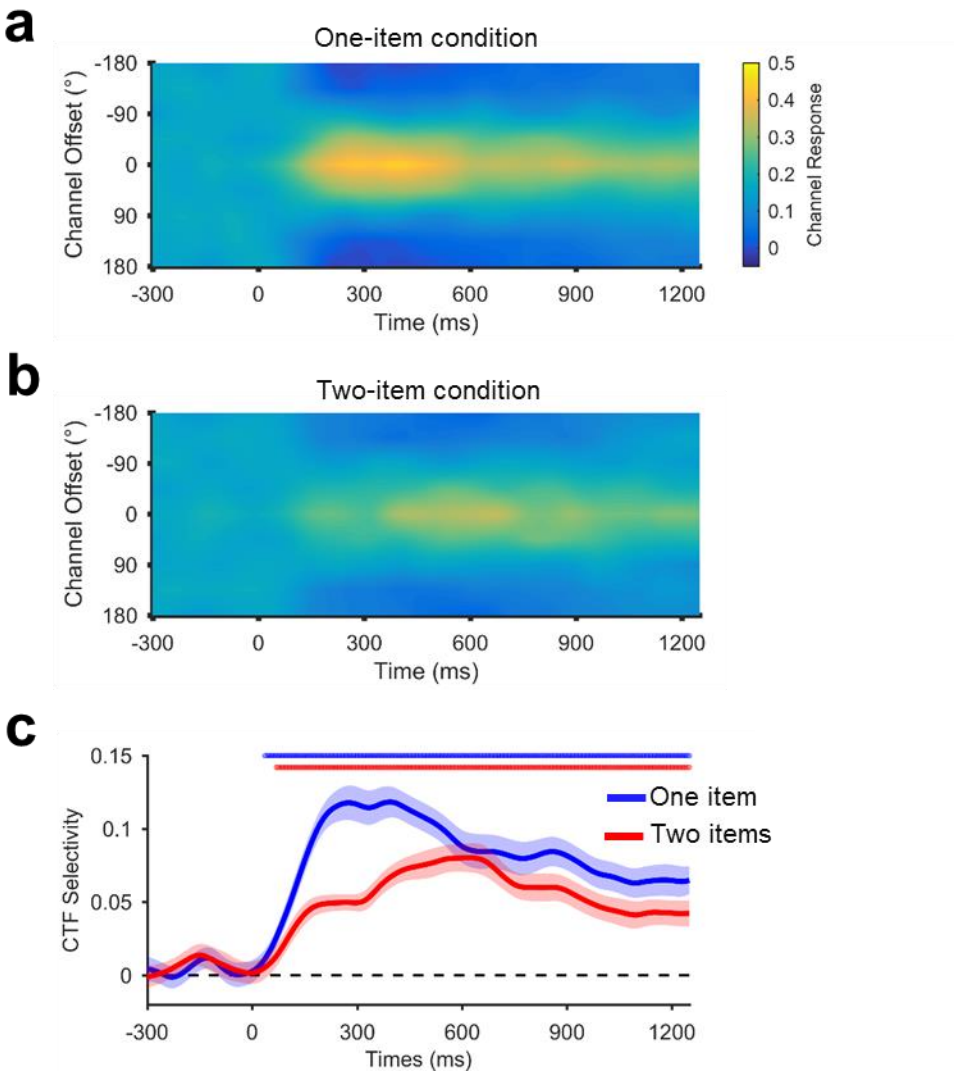

Figure 2-2. Spatial alpha-band CTFs as a function of memory load.
Average alpha-band CTF in the one- and two-item conditions (**A** and **B** respectively). **C.** The spatial selectivity of alpha-band CTFs across time (measured as CTF slope, see Materials and Methods) as a function of memory load. The blue (one-item) and red (two-item) markers at the top of the panel indicate the period of above-chance selectivity obtained using a cluster-based test. The shaded error bars reflect ±1 bootstrapped SEM across subjects.

Alpha-band activity concurrently encodes two spatial representations

Consistent with past work (Buschman et al., 2011; Emrich et al., 2013; Sprague et al., 2014, 2016), we observed that spatially specific alpha-band activity deteriorates as memory load

increases. However, declines in stimulus-specific activity with increasing load can be explained in two ways. On the one hand, the observed decrease in selectivity might reflect the loss of memory fidelity due to competition between representations when multiple stimuli are maintained in an active state (Bays, 2014; Franconeri et al., 2013). On the other hand, it is possible that at any given moment only a single item is maintained in an active state, while the other item stored in an activity-silent state (Figure 3-3a). Critically, this switching account asserts that with a memory load of two items, each item can only be represented 50% of the time on average. To test whether this switching account can explain the CTFs seen during the two-item condition, we simulated the CTF selectivity expected under a switching account. To this end, we generated CTFs from the single-item condition but randomized the position labels for 50% of the trials (see Materials and Methods). We then compared CTF selectivity seen during the two-item condition with the CTF selectivity expected under the switching account. We reasoned that if the switching account was correct, we should see no difference between CTF selectivity for two-item trials and for the simulated switching conditions. However, if we observed a higher CTF selectivity for the observed two-item data, we could conclude that alpha activity reflects the simultaneous maintenance of multiple locations in WM during a given trial.

Figure 3-3b shows CTF selectivity across time for the two-item condition and for simulated switching based on the one-item data. We found that CTF selectivity was higher throughout the delay period (averaged from 100 to 1250 ms after stimulus onset) for the two-item (M = 0.062, SD = 0.036) than expected based on the switching account (M = 0.043, SD = 0.030). A resampling test revealed that this difference was reliable ($p < .0001$). We also observed a reliable difference when we restricted our analysis to a window late in the delay period (800-1250 ms) to minimize the contribution of stimulus-driven activity ($p = .001$). This analysis

provides definitive evidence that multiple locations are simultaneously represented by alpha-band activity, and that observed decreases in CTF selectivity as memory load increases reflect a decline in the quality of active representations rather than a consequence of rapid switching between active representations.

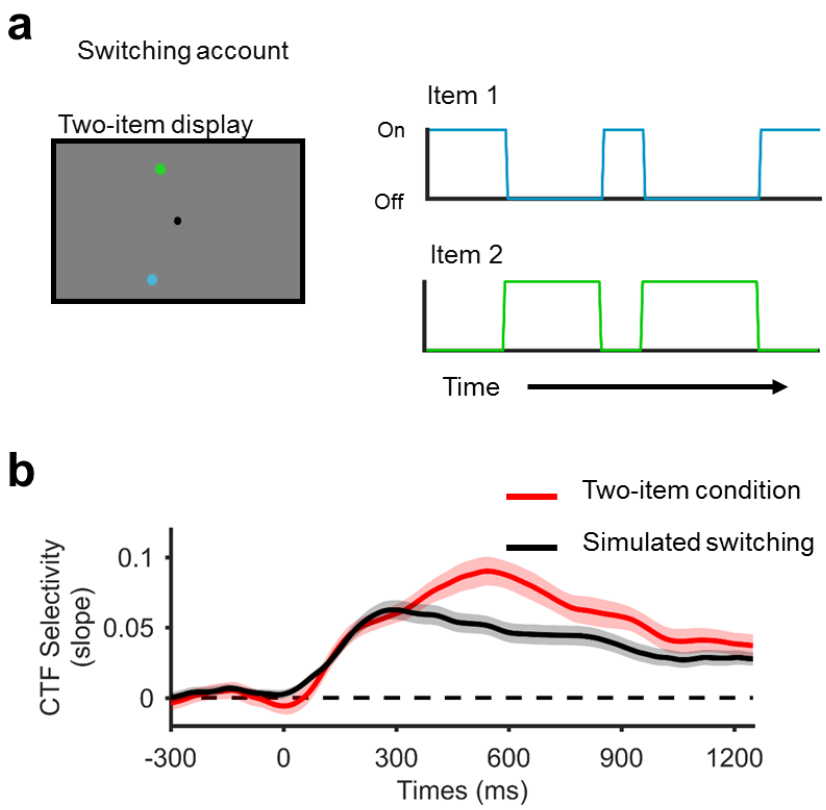

Figure 2-3. Alpha-band activity concurrently represents two spatial positions.
A. If only one item can be represented by alpha-band activity at a time, then in the two item conditions the items might alternate between an active and a activity-silent state such that only one item is actively represented at once. We used data from the single-item condition to simulate the expected CTF based on this switching account. This account holds that each item is represented 50% of the time (on average). Thus, we simulated switching between items by randomizing the position labels for 50% of trials. **B.** Spatial selectivity of alpha-band CTFs across time (measured as CTF slope) for the two-item condition (red) and for simulated switching (black). The shaded error bars reflect ±1 bootstrapped SEM across subjects.

The frequency of oscillations that encode spatial representations does not change with memory load

The decline in the spatial selectivity of alpha-band activity with increasing memory load suggests that the fidelity of the spatial representations decreased as memory load increased. However, another possibility is that the remembered locations were represented by a different frequency band when memory load increased. To test this possibility, we performed the IEM analysis separately for one-item and two-item conditions (i.e., both training and testing within each condition; see Materials and Methods) across a range of frequencies (4–50 Hz). We conducted a cluster-corrected permutation test to identify reliable clusters of above-chance CTF selectivity. Consistent with our past work (Foster et al., 2016), we observed a burst of spatially specific activity across a range of frequencies (4–25 Hz). However, only alpha-band activity (8–12 Hz) tracked the remembered position(s) throughout the delay period (Figure. 4a and 4b). An overlay plot of spatially specific frequencies in both one item and two item trials revealed a strikingly similar frequency profile later during the delay period, when stimulus-driven activity has subsided (Figure. 4c). These findings show that the frequency of oscillatory activity that encodes spatial representations does not change with memory load. Thus, the decrease observed in the spatial selectivity of alpha-band CTFs with increasing memory load reflects a decline in spatially selective activity rather than a shift in the frequency of spatially selective oscillations.

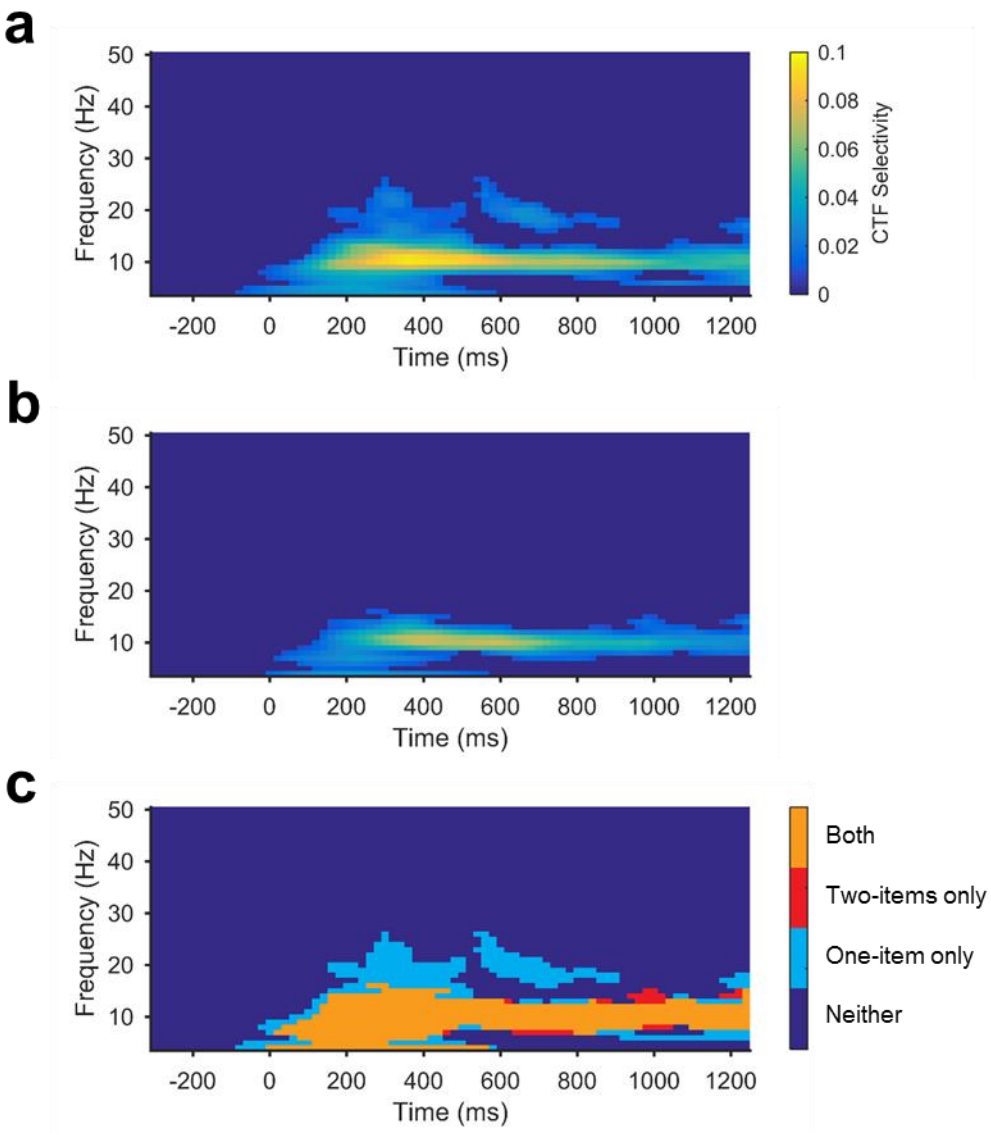

Figure 2-4. Identifying frequencies that track one- and two-item spatial memories
Selectivity of spatial CTFs (measured as CTF slope) reconstructed from the scalp distribution of oscillatory power across a broad range of frequencies for the one-item **A.** and two-item **B.** conditions. Points with no reliable CTF selectivity as determined by the cluster-corrected permutation test are set to dark blue. **C.** Overlay plot marking the clusters of reliable selectivity in the one-item condition (light blue), two-item condition (red), and both (orange).

Discussion

A longstanding view is that the maintenance of information in WM is realized via persistent neural activity that encodes the content of WM. In support of this view, past work has

shown that patterns of delay-period activity track visual features maintained in WM (Foster et al., 2016; Funahashi et al., 1989, 1993; Harrison & Tong, 2009; Larocque et al., 2017; Serences et al., 2009; Sprague et al., 2014; Sreenivasan et al., 2014). However, recent work has challenged content-specific activity as the sole mechanism that supports maintenance in WM, instead proposing that other activity-silent mechanisms can also support the retention of information (Lewis-peacock et al.,, 2012; Stokes, 2015; Rose et al., 2016; Wolff et al., 2017). This activity-silent account warrants a re-evaluation of the evidence that all items maintained in WM are represented in content-specific patterns of activity. The prevailing view has been that multiple items are concurrently represented by content-specific delay activity. Indeed, recent fMRI work finds that increased memory load results in a decline in the content-specific patterns of activity during WM maintenance (Emrich et al., 2013; Sprague et al., 2014, 2016), mirroring behavioral declines in performance, and similar results have been reported using unit recordings from non-human primates (Buschman et al., 2011). However, the activity-silent account raises the possibility that only a single item is actively represented at a time and the *apparent* representation of multiple items might reflect the fact that typical analyses aggregate data across trials, averaging together periods when an item is represented and when it is not represented.

Here, we tested this switching account. We used an inverted encoding model (IEM) combined with EEG measurements of oscillatory alpha-band (8–12 Hz) power to track remembered spatial locations during a WM task in which human observers remembered one or two spatial locations. Consistent with previous fMRI and non-human primate unit recordings that found that content-specific activity declines with WM load (Buschman et al., 2011; Emrich et al., 2013; Sprague et al., 2014),we found that the spatial selectivity of alpha-band activity declined as memory load increased. We tested the switching account by simulating selectivity expected

under a switching model in which only one item was actively represented at a time. We found that the spatial selectivity of alpha-band CTFs during two item trials was greater than would be expected if only one item was represented at once. This test definitively rules out the possibility that only one item is represented in an active state at any given point in time. Thus, our findings provide clear evidence that two items can be represented concurrently in an active state.

At first glance, this finding seems inconsistent with spatial attention studies that have provided evidence for rhythmic sampling when multiple locations are maintained in the focus of attention (e.g., Busch and Van Rullen, 2010). However, the switching account that we tested is a stringent version of rhythmic sampling in which only a single location is represented at any given time. Our simulations do not rule out all classes of rhythmic sampling models, but they do constrain these models. Namely, our data demonstrate that rhythmic sampling accounts of working memory storage must allow for the concurrent representation of two items. One possibility is that there are rhythmic fluctuations in the extent to which each location is actively represented, but that these representations are not fluctuating between the presence of a representation and the complete absence of a representation.

Evidence that two active representations can be maintained concurrently is also relevant to a long-standing debate about active representations in embedded process models of WM. Embedded process models define memory as a common storage space with different levels of activation corresponding to short-term memory and long-term memory (Larocque, Lewis-peacock, & Postle, 2014). One key distinction between competing embedded process models is the number of items that can be represented in an active state in working memory. The proposed number of items ranges from a strict limit of one item (McElree, 2006) to four items (Cowan, 1995; Oberauer & Hein, 2012), and competing models are each able to account for many aspects

of behavioral performance. Evidence for concurrent maintenance challenges models proposing a strict one-item limit for the number of items that can be actively represented. However, it is worth noting that single item activation accounts (McElree, 1998) have suggested that simultaneous displays, like the display used in the present study, may reflect a special case in which multiple items can be actively maintained (McElree & Dosher, 2001). Future work can test this idea by measuring the number of items that can be concurrently represented when items are presented in series.

Here we found that alpha band (8–12 Hz) activity tracks concurrent maintenance of multiple remembered locations. Whether this observation generalizes to previous fMRI and non-human primate unit recording observations that content-specific activity declines with memory load remains an open question. Particularly intriguing is the possibility that the number of items maintained in an active state may vary by brain region. For instance, a recent retro-cuing study found that only cued memory representations were represented in visual cortex, while both cued and un-cued representations were represented in parietal and frontal cortex (Christophel, Iamshchinina, Yan, Allefeld, & Haynes, 2018). This study suggests that there may be interesting differences in the number of items represented across cortex. For instance, it is possible that only a single item can be represented in V1-4 whereas multiple items can be represented in higher order regions. Thus, reanalysis of existing fMRI data (e.g., Sprague et al., 2014, 2016) with explicit tests of switching models might reveal interesting differences in the number of items concurrently maintained across brain regions.

Conclusion

Working memory (WM) is a mental workspace where we temporarily hold information "online" in pursuit of our current goals. However, recent activity-silent models of WM have

challenged the view that all items are held in an "online" state, instead proposing that only a subset of representations in WM – perhaps just one item – are represented by persistent activity. Here we directly tested a single-item model of persistent activity using a multivariate spatial encoding model to read out the strength of two representations from alpha-band power in the human EEG signal. Critically, we provide the first direct evidence that both locations were maintained *concurrently*, ruling out the possibility that declines in content-specific activity are due to storing one of two items in an activity-silent state.

CHAPTER 3. RETRIEVAL PRACTICE ENHANCES THE ACCESSIBILITY BUT NOT THE QUALITY OF MEMORY

Introduction

Numerous studies have demonstrated that retrieval from long term memory (LTM) can enhance subsequent memory performance, a phenomenon labeled the retrieval practice effect (Carrier & Pashler, 1992). The benefits of retrieval practice have been observed with a wide variety of memoranda (Roediger & Karpicke, 2006) including word pairs (Pyc & Rawson, 2009), pictures (Wheeler & Roediger, 1992) and spatial positions (Carpenter & Kelly, 2012; Carpenter & Pashler, 2007; Rohrer, Taylor, & Sholar, 2010).

Varying explanations have been offered for how retrieval practice enhances memory performance. Some have focused on increased elaborative retrieval during testing (Carpenter, 2009) while others have emphasized the narrowing of the retrieval search space via helpful contextual associations (Lehman, Smith, & Karpicke, 2014). One common assumption of these accounts, however, is that retrieval practice enhances the probability of access to a memory rather than the quality of the memory. This focus on accessibility over fidelity may be attributable in part to the fact that past studies have typically used discrete word or picture stimuli (and all-or-none measures of accuracy) that do not allow clear measurements of memory fidelity. That said, some past findings may be consistent with a putative effect of retrieval practice on memory quality. For example, Chan and McDermott (2007) found that retrieval practice improved participants' ability to avoid semantically similar lures during a recognition test, and improved source memory. Likewise, Szpunar, McDermott, and Roediger (2008) found that testing improves list discrimination. However, while each of these findings could reflect a more precise memory (e.g., of specific semantic content, or of the temporal context associated

with an item), the binary nature of the responses in these studies also allows for an interpretation based on retrieval probability.

An approach that may provide more traction for understanding the effect of retrieval practice on the quality of item specific memory is to allow participants to report remembered information along a continuous response space. For example, Carpenter and Kelly (2012) used a continuous response space in a task where subjects recalled the precise positions of different objects. Retrieval practice resulted in a decrease in the average response error for retrieved locations relative to restudied locations. However, although a change in memory quality provides an intuitive explanation of these findings, a reduced guessing rate in the retrieval practice condition would also yield lower average response errors. Thus, the goal of the present work was to examine the retrieval practice effect using an analytic approach that can estimate both the probability of retrieval and the quality of the retrieved representations.

We measured performance in a shape/color recall task in which the possible colors were drawn from a continuous 360 degree space, and we used a mixture-modeling approach (Zhang & Luck, 2008) that provided separate measures of the probability of recall and the quality of the retrieved memories. This analytic approach has been widely applied to the field of working memory (see Luck & Vogel, 2013 for review), and has recently been applied to the study of LTM (Brady et al., 2013). To anticipate our conclusions, retrieval practice elicited robust improvements in the probability of memory access, but absolutely no improvement in the fidelity of the retrieved memories.

Experiment 1

Method

Participants

Twenty-two undergraduates at the University of Oregon completed the experiment for course credit. All participants gave informed consent according to procedures approved by the University of Oregon institutional review board.

Apparatus

Stimuli were generated in MATLAB using Psychophysics Toolbox extension (Brainard, 1997; Pelli, 1997) and were presented on a 17-in. flat CRT computer screen (60-HZ refresh rate). The viewing distance was ~80 cm. Stimuli were 9.2 x 9.2° of visual angle.

Stimuli

Four hundred nameable pictures (e.g., animals, plants, shapes, countries, US states, and symbols) were obtained via a web search for royalty free clip art. One of 360 continuous colors was assigned to each image, with different color/shape sets for each subject.

Task and Procedure

The 400 stimuli were presented in two successive runs, each containing 200 distinct shape/color associations. Each run was comprised of two parts, a learning period and a delayed-retrieval period. During the learning period, images were presented serially in blocks of 10 items, followed either by retrieval practice, during which all ten colors were recalled or by the start of the next block of 10 items (Figure 4-1); thus, subjects did not know during encoding whether or not they would be immediately tested. Images were tested in a random order without feedback.

After viewing all 200 images with retrieval practice for half of the items in the run (~20-30 minutes) subjects were asked to recall the color of each image by clicking on a color wheel

that represented all of the presented colors. Images were tested in a random order relative to their initial presentation. Participants received feedback consisting of the presentation of the shape filled with the correct color and a number denoting the magnitude of the error.

During recall, a white shape cue was displayed for 1 second before the cursor and color wheel appeared (Figure 4-1b). During response selection, the color of the shape cue shifted continuously to match the hue that was indicated by the mouse cursor on the color wheel. Participants indicated their color choice by clicking the mouse. Responses were unspeeded and accuracy was given highest priority; subjects were instructed to choose a response even if they felt they were guessing. When they thought they were guessing, they were instructed to click with the right mouse button rather than the left. The color wheel was randomly rotated across trials (so that position information was irrelevant to the color response). Following completion of the first run of 200 images, the remaining 200 images were presented and tested using the same procedure (i.e. a learning period and delayed-retrieval period) with 200 new images. One image was presented twice during the learning period of run one and was dropped from the delayed

analyses.

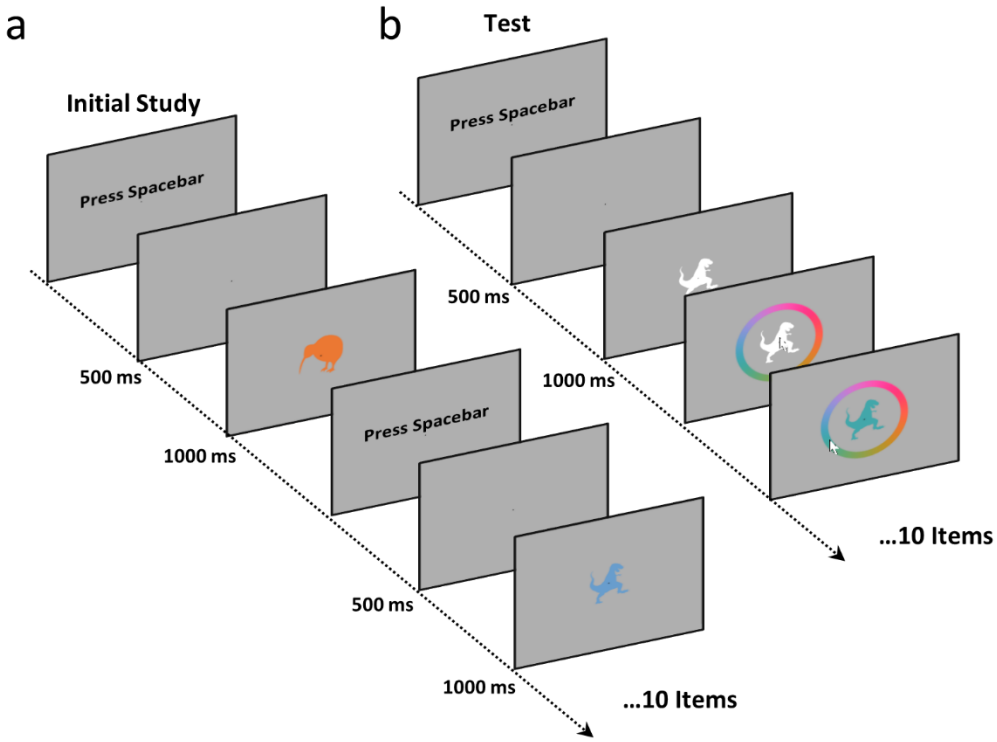

Figure 3-1 Task diagram for Experiment 1
A. Timing of stimuli for initial study opportunity. Participants studied 10 stimuli and then were either tested on those ten images or not tested. Subjects studied 400 images total. **B.** Timing of each item at initial test and at final test.

Data Analysis

Response error was measured as the number of degrees between the presented color and the reported color. Errors ranged from 0° (perfect response) to ±180° (a maximally imprecise response). Responses were centered on 0° but spanned the entire range of responses (for example see, Figure 4-2a). These error histograms are well described as a mixture of two distributions which reflect guesses and correct responses (Zhang & Luck, 2008). On some trials

subjects do not remember the color associated with the shape cue and guess randomly with respect to the target color. This results in a uniform distribution of responses with respect to the target color. On other trials, participants remember the color of the shape cue and provide responses centered on the correct color value but with some degree of error. This distribution is well described by a von Mises distribution (the circular analogue of a Gaussian distribution because the tested color space was circular) centered on the correct response. To obtain an estimate of these two distributions, response errors were fit using Markov Chain Monte Carlo (MCMC) as employed by the "memfit" function of Memtoolbox (Suchow et al., 2013). MCMC repeatedly samples parameter values in proportion to how well they describe the data and the prior (in this case an uninformative Jeffreys prior) to obtain a Maximum a Posteriori (MAP) estimate of three parameters: P_{mem} is the probability that subjects could retrieve nonzero target information, operationalized as the inverse of the height of the uniform distribution (i.e. 1 – proportion of guesses). *SD* is the standard deviation of the von Mises distribution (with larger values reflecting reduced precision). *Mu* (μ), the mean of the von Mises distribution reflects, systematic bias in the error distribution (preferred clockwise or anti-clockwise responses on the color wheel).

These parameters are calculated using the distribution of all responses, which is a *mixture* of responses not guided by memory (guesses) and responses guided by memory. Thus, we can determine the proportion of remembered items and the precision of responses guided by memory, but it is not possible to determine if any individual response was guided by memory.

Results

Aggregate data

All participants' responses were combined into an aggregate error histogram (see Figure 4-2a.) and fit using the "memfit" function of Memtoolbox (Suchow et al. 2013) to obtain parameter estimates and 95% credibility intervals (CrI); there is a 95% chance that the true value of the parameter for the sample lies between the credibility intervals. We will refer to parameters with overlapping credibility intervals as "not significantly different" and parameters with non-overlapping credibility intervals as "significantly different". Unlike confidence intervals, Bayesian credibility intervals are not necessarily symmetrical.

The mixture modeling analysis revealed that 70.7% (CrI:-1.7%,+2.0%) of the items were recalled during the initial test. SD – our operational definition of mnemonic precision – was 21.4° (CrI:-.8°,+.1.1°). At delayed test, subjects recalled significantly more items that they had previously retrieved (53.8%, CrI:-1.9%,+2.3%) than items that that were previously untested (37.9%, CrI:-2.2%,+2.8%, Figure 4-2). Mnemonic precision was not significantly different

between tested (22.9°, CrI:-1.0°,+1.5°) and untested (24.2°, CrI:-1.6°,+2.6°) items.

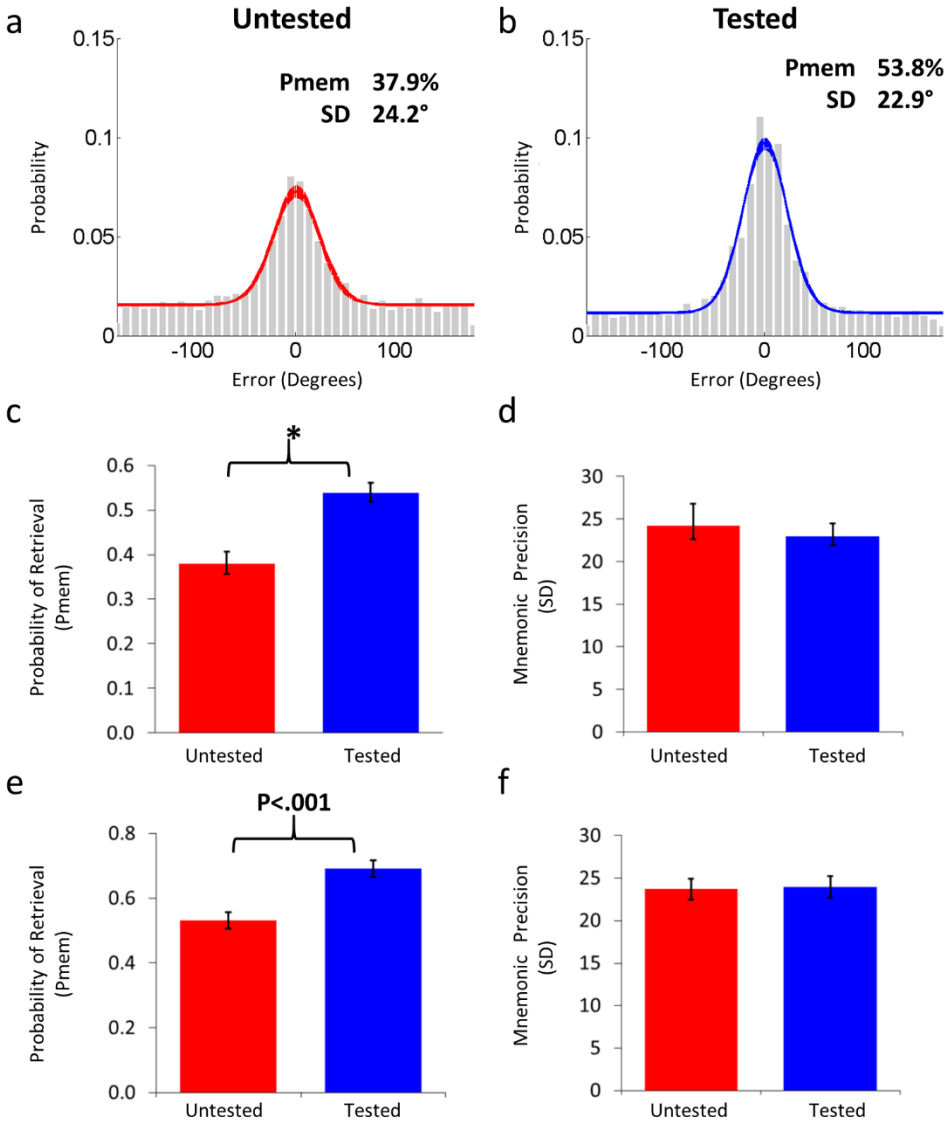

Figure 3-2 Delayed test results from Experiment 1
A. Aggregate fit of delayed response errors from all subjects for initially untested items. **B.** Aggregate fit of delayed response errors from all subjects for initially retrieved items. **C**+**D.** Aggregate parameter estimates of probability of retrieval (P_{mem}) and mnemonic precision (SD) at delayed test*. Error bars represent Bayesian credibility intervals of the fits. (E+F) Average of individual parameter estimates of P_{mem} and SD for subjects who successfully retrieved > 40% of the items at the delayed test. Error Bars represent 95% within-subjects confidence intervals that were calculated by normalizing to remove within subject variance (Loftus & Masson, 1994).
* Denotes non overlapping credibility intervals.

Simulations

We were interested in examining the data at the individual subject level, but simulations showed that there would be consistent biases in the precision estimates if the probability of retrieval was too low. We determined this by generating artificial data that presumed varying P_{mem} values and SD values equal to those observed in our aggregate data (20°). Parameter estimates were obtained from these artificial datasets by sampling 100 times from each dataset and then fitting each sample with a mixture model. These simulations revealed that SD is systematically overestimated when the proportion of successfully retrieved items was less than 40% (Figure 4-3a.). By contrast, the P_{mem} parameter is relatively accurate even when probability of retrieval is low. Thus, to avoid misleading estimates of SD, we compared individual parameter estimates of precision only for subjects who successfully retrieved at least 40% of the items in both the tested and untested conditions. Further simulations confirmed that estimates of the SD parameter would not be affected by high guess rates in the aggregate data, because of the large number of trials run across all subjects (>4000 trials per condition). Thus, in the aggregate analysis, accurate P_{mem} and SD estimates could be obtained even when probability of retrieval was low (Figure 4-3b).

Individual Parameter Comparisons (Delayed Test)

Analysis of the subset of subjects who successfully retrieved 40% or more items in both conditions (n = 12) also showed higher P_{mem} for retrieved items (M = 69.2%, SD = 13.4%) compared to untested items (M = 53.2%, SD = 9.9%, $t(11) = -6.03$ $p < .001$). Also in line with the aggregate data, subjects did not exhibit superior mnemonic precision for items that they had previously retrieved (M = 24.0, SD = 5.3) compared to items that were not retrieved (M = 23.7,

SD = 5.7, *t*(11) = -.22 p = .83).

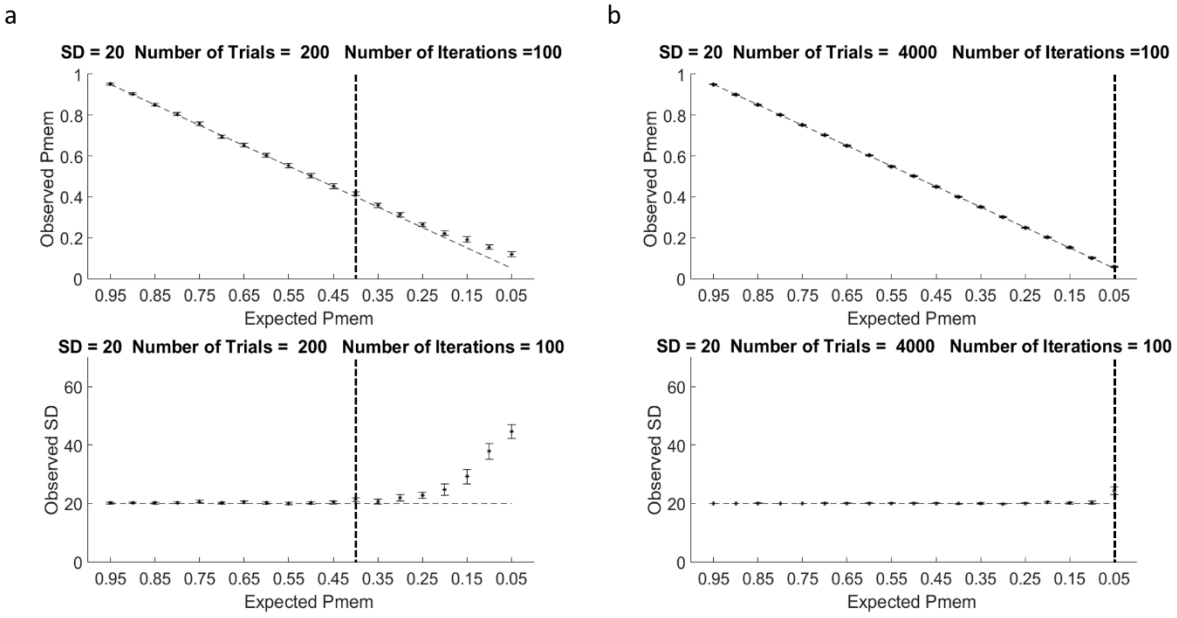

Figure 3-3 Simulations of parameter estimates
A. Parameter estimates obtained by taking 100 iterations of 200 trials each from distributions of known probability of retrieval (P_{mem}) and precision (SD). Error bars represent 95% confidence intervals, and dotted lines represent expected values. Top graph is observed probability of retrieval bottom is observed precision **B.** Parameter estimates obtained by taking 100 iterations of 4000 trials each from distributions of known probability of retrieval (P_{mem}) and precision (SD). Error bars represent 95% confidence interval, and dotted lines represent expected values. Top graph is observed probability of retrieval bottom is observed precision.

Discussion

Experiment 1 suggests that retrieval practice increases the probability that an item can be retrieved in the future but does not improve the precision of that memory. In Experiment 2, we equated the number of times that participants saw and responded to each item by comparing the retrieval practice condition with a restudy condition (e.g., Carrier and Pashler, 1992).

Experiment 2

Method

Participants

Twenty-eight students from the University of Oregon participated in Experiment 2 for course credit or monetary compensation. Six participants were excluded: two did not complete the session, one was excluded during the session for not following instructions, and three participants who completed the session were excluded for responding randomly on restudy trials. Twenty-two participants were included in the analysis of Experiment 2. Six subjects who did not complete all trials in the time allotted were included in the experiment because they had completed the session and followed instructions.

Task

The task in Experiment 2 was the same as experiment one except that during the learning period subjects either completed a memory test or were given a chance to restudy the items after every 10 images (Figure 4-4a). During restudy the shape cue (with the correct hue) was presented simultaneously with a color wheel and subjects were instructed to select the color of the presented item by clicking on the color wheel. The learning period was followed by a recall

test of all of the items (with item-by-item feedback on the degree of response error).

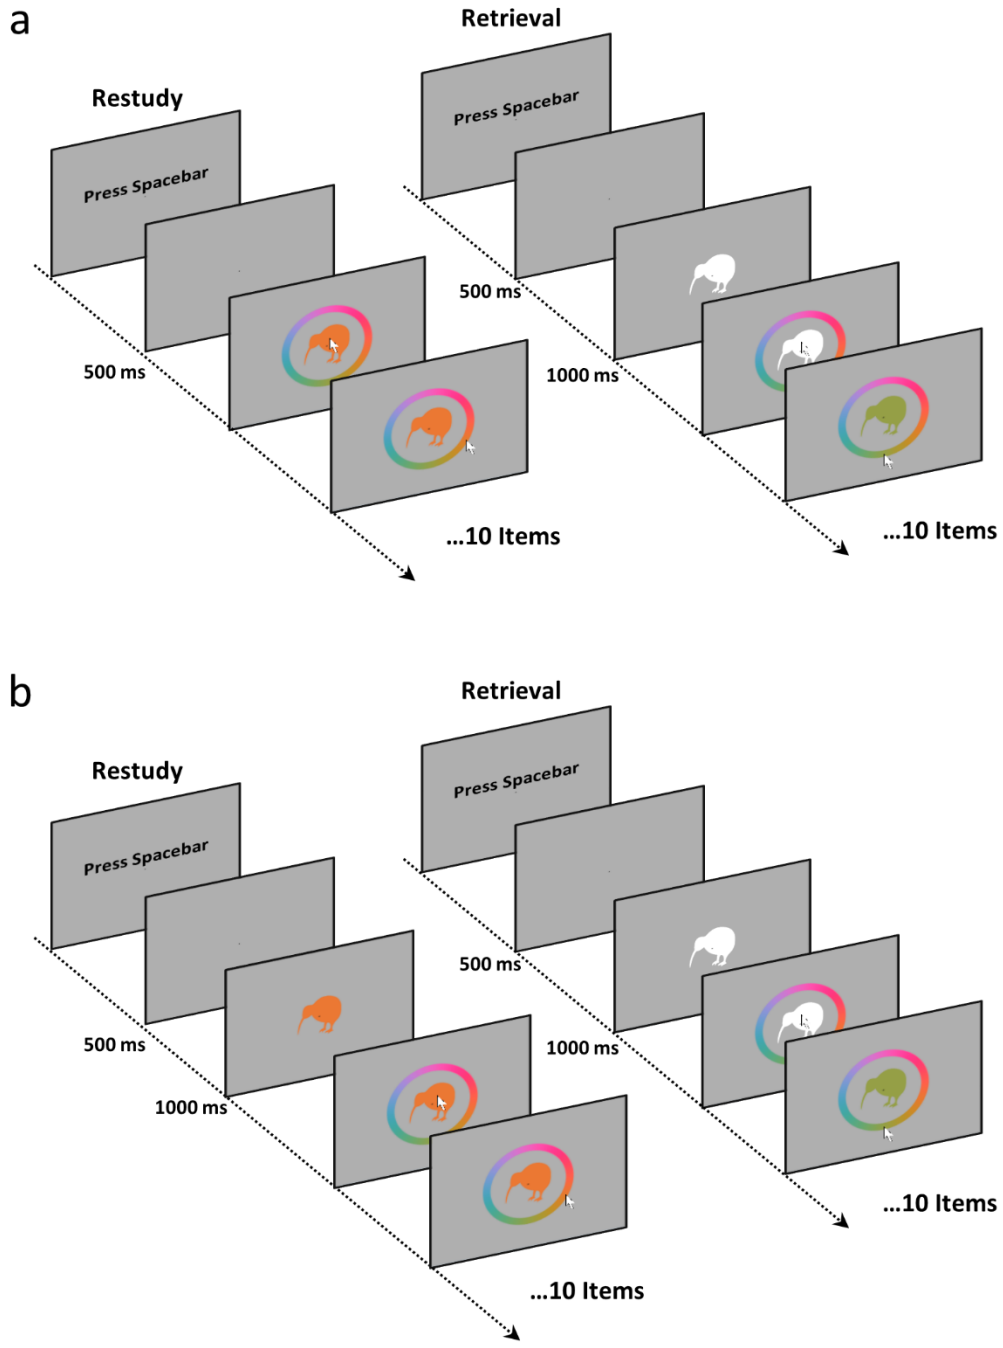

Figure 3-4 Task diagram for Experiment 2 and 3
A. Stimuli timing for restudy and retrieval for Experiment 2. **B.** Stimuli timing for restudy and retrieval for Experiment 3. Final test timing was the same as initial test timing.

Analyses

Similar to experiment one, we relied upon an aggregate fit to assess the mnemonic precision for all subjects, and then looked at individual fits for subjects who retrieved at least 40% of the items (P_{mem} > 40%). Additional simulations with fewer trials revealed that this was also an appropriate cut off for subjects who did not complete all trials.

Results

Aggregate

Seventy four percent (CrI:-2.0%,+1.5%) of the items were recalled during the initial test, and as expected, participants correctly selected responses for over 99% (CrI:-.4%,+.3%) of the items during the restudy task when the stimuli were physically present to guide responses. Not surprisingly, precision was substantially higher for the restudy (SD = 7.2°, CrI:- .2°,+.2 °) than for the memory task (18.6°, CrI:- .8°,+.6°).

At delayed test (Figure 4-5), subjects recalled a significantly higher proportion of items that were initially retrieved (58.2%, CrI:-2.0%,+2.5%) than items that were restudied (50.6%, CrI:-.1%,+2.1%). Unlike the results of Experiment 1, mnemonic precision was significantly better for restudied items (18.8°, CrI:-1.2°,+1.2 °) relative to retrieved items (21.3°, CrI:-1.0°,+1.2°, Figure 4-5). As we report below, however, this relative disadvantage in mnemonic precision in the

testing condition did not replicate in Experiment 3.

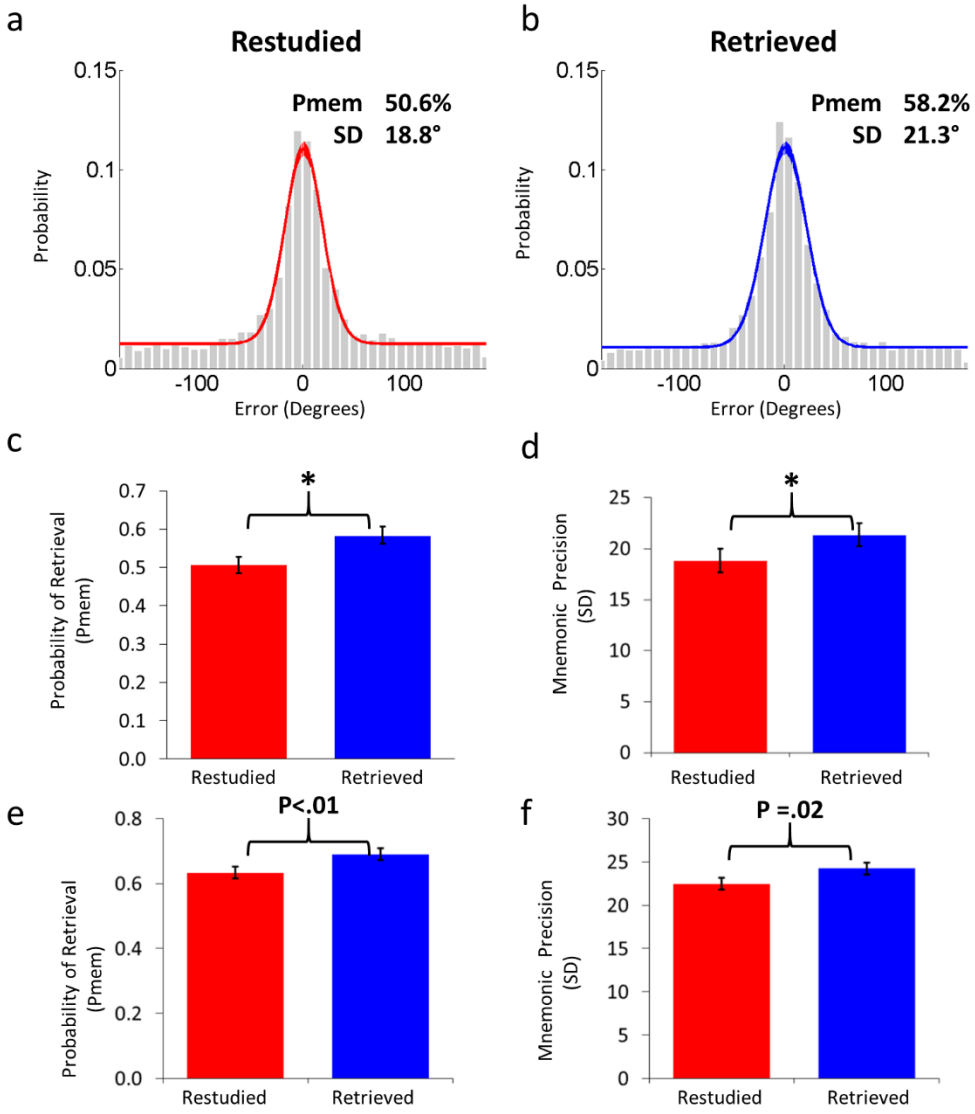

Figure 3-5 Delayed test results from Experiment 2
A. Aggregate fit of delayed response errors from all subjects for initially restudied items. **B.** Aggregate fit of delayed response errors from all subjects for initially retrieved items. **C + D.** Aggregate parameter estimates of probability of retrieval (P_{mem}) and mnemonic precision (SD) at delayed test*. Error bars represent Bayesian credibility intervals of the fits. (E+F) Average of individual parameter estimates of P_{mem} and SD for subjects who successfully retrieved > 40% of the items at the delayed test. Error Bars represent 95% within-subjects confidence intervals that were calculated by normalizing to remove within subject variance (Loftus and Mason, 1994). * Denotes non overlapping credibility intervals.

Individual Parameter Comparisons (Delayed Test)

Analysis of the subset of subjects who remembered at least 40% of items (n = 17) in both conditions revealed that subjects recalled a significantly higher proportion of items they had previously retrieved (M = 69.1%, SD = 14.8%) relative to items which were previously restudied (M = 63.4%, SD = 14.4%, t(16) = -3.06 , p = .008). As in the aggregate data, subjects exhibited superior mnemonic precision for items that they had previously restudied (M = 22.5°, SD = 8.2°) relative to retrieved items (M = 24.3°, SD = 9.2°, t(16) = -2.60, p = .02).

Discussion

As in Experiment 1, retrieval practice improved the probability of successful delayed recall but not mnemonic precision. Thus, the benefits of retrieval practice on probability of retrieval were robust when the control condition allowed extra time to restudy the memoranda. In Experiment 3, we tested whether a similar empirical pattern would emerge when we equated the amount of exposure time between retrieval and restudy, and whether the same pattern would emerge following a >24 hour retention interval.

Experiment 3

Method

Participants

Twenty three students from the University of Oregon participated in Experiment 3 for course credit. Two participants who did not complete all trials in the time allotted were included in the experiment. All participants gave informed consent according to procedures approved by the University of Oregon institutional review board.

Task

The task in Experiment 3 was the same as Experiment 2 except for two differences. First, to equate total presentation time with that in the testing condition, restudied items were displayed for 1s before subjects could respond (Figure 4-4b). Second, to determine if the same pattern of results would emerge over a longer delay, the two runs of the task were completed on separate days, 1-4 days apart. This allowed for a >24hr delayed retrieval of the items learned from the first run before subjects completed the second run of the experiment on day 2. Twenty subjects completed the surprise second retrieval period (three subjects arrived late for the session and skipped the >24hr retrieval to ensure a prompt finish).

Analyses

Analyses were identical to Experiments 1 and 2. Only the aggregate analysis was applied to the >24hr delayed test because subjects were only tested on 100 items in each condition and probability of retrieval was low.

Results

Aggregate

Sixty-three percent of the items (CrI:-2,+2) were recalled during the initial test, and as expected, participants correctly selected responses for 99.7% (CrI:-.2%,+.1%) of the items during the restudy task. Also as expected, precision was significantly higher for restudy (SD = 7.7°, CrI:-.2°,+.2°) than for retrieval (21.3°, CrI:-.9°,+1.1°).

At delayed test (Figure 4-6), subjects recalled a significantly higher proportion of items that were previously retrieved (47.3%, CrI:-2.6,+2.1)) than items that that were previously restudied (40.2%, CrI:-2.2,+2.8), Figure 4-6). Unlike the results from Experiment 2, estimates of mnemonic precision were not significantly different for retrieved (24.9°, CrI:-1.6°,+1.7°) and

restudied (22.5°, CrI:-1.6°,+2.5°) items. Thus, although Experiment 2 revealed relatively better precision in the restudy condition, this does not appear to be a robust empirical pattern.

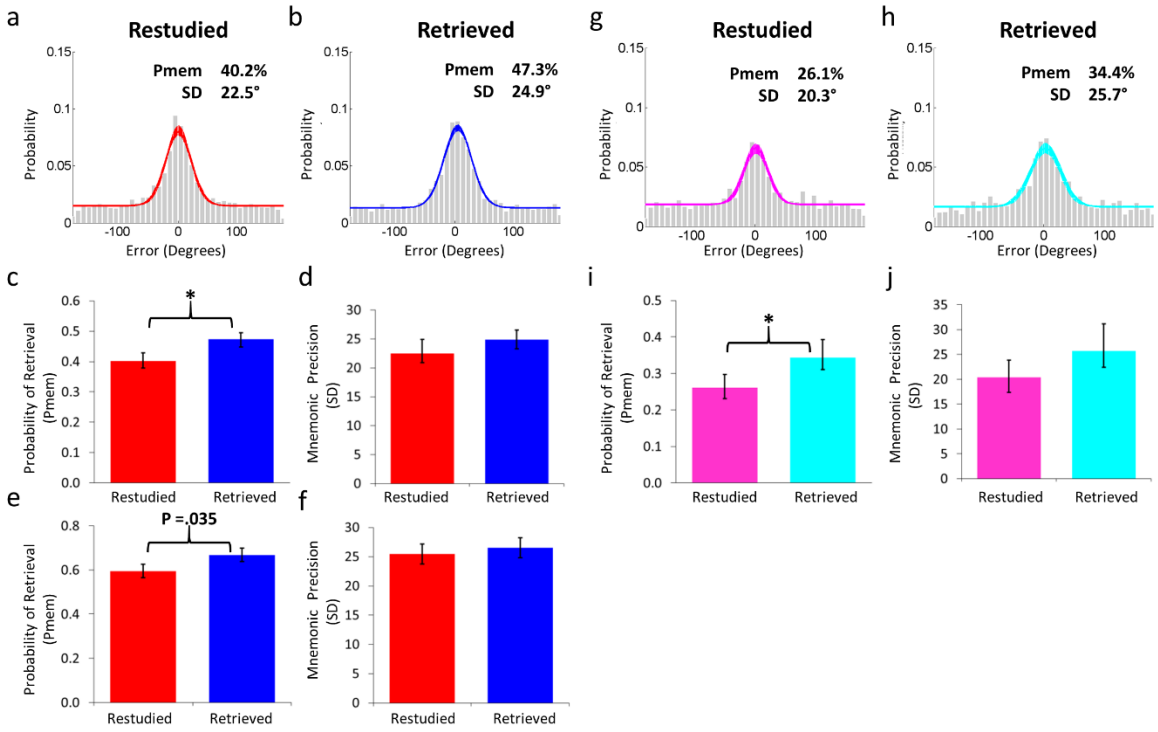

Figure 3-6 Delayed test results from Experiment 3
A. Aggregate fit of 30 min delayed response errors from all subjects for initially restudied items **B.** Aggregate fit of 30 min delayed response errors from all subjects for initially retrieved items. **C+D.** Aggregate parameter estimates of probability of retrieval (P_{mem}) and mnemonic precision (SD) at 30 minute delayed response errors. Error bars represent Bayesian credibility intervals of the fits*. **E+F.** Average of individual parameter estimates of P_{mem} and SD for subjects who successfully retrieved > 40% of the items at the 30 min delayed test. Error Bars represent 95% within-subjects confidence intervals that were calculated by normalizing to remove within subject variance (Loftus and Masson, 1994). **G.** Aggregate fit of >24hr delayed response errors from all subjects for delayed test of initially restudied items **H.** Aggregate fit of >24hr delayed response errors from all subjects for delayed test of initially retrieved items. **I+J.** Aggregate parameter estimate for P_{mem} and SD during >24 hr delay*. Error bars represent Bayesian credibility intervals of the fits.* Denotes non overlapping credibility intervals.

The pattern of results observed during the test after more than 24hrs was similar to the pattern of results for the first delayed test. Subjects recalled a significantly higher proportion of items that they had previously retrieved (34.4%, CrI:-3.2%,+5.0%) than items that they had restudied (26.1%, CrI:-3.1%,+3.5%, Figure 4-6). Estimates of mnemonic precision were not significantly different for retrieved (25.7°, CrI:-3.2°,+5.5°) and restudied (20.4°, CrI:-3.0°,+3.5°) items.

Individual Parameter Comparisons (Delayed Test)

Analysis of the subset of subjects who successfully retrieved at least 40% items (n = 12) revealed that subjects recalled a significantly higher proportion of previously retrieved items (*M* = 66.8%, *SD* = 11.4%) relative to previously restudied items (*M* = 59.5%, *SD* = 10.4%, *t*(11) = -2.35 , *p* = .039). In contrast to the findings from experiment 2, subjects exhibited similar precision for items that they had previously restudied (*M* = 25.5°, *SD* =7.3°) relative to previously retrieved items (*M* = 26.6°, *SD* = 5.0°, *t*(11) = 0.62, *p* = .54).

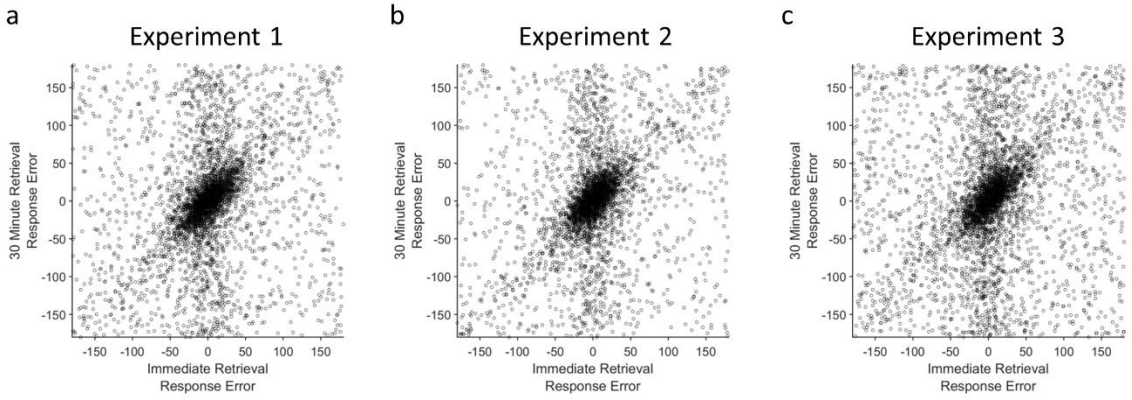

Figure 3-7 Relationship between immediate and delayed responses
A-C. Scatter plot of immediate and 30 minute delayed responses for retrieved items for all subjects in experiments 1, 2, and 3 respectively. Responses from all subjects were pooled for the visualization.

Discussion

The findings from Experiment 3 are in line with the findings from Experiment 1 and Experiment 2. Retrieval practice improves the probability of successful delayed recall, but does not improve mnemonic precision. Thus, the benefits of retrieval practice on recall probability were robust when the control condition allowed extra time to restudy the memoranda, and when delayed recall did not take place for over 24 hours.

General Discussion

In three experiments, we demonstrated that retrieval practice improves probability of retrieval but not mnemonic precision. Furthermore, in Experiments 2 and 3 subjects provided a response to restudied items by selecting the color they were viewing on the color wheel. Thus, we were able to replicate a critical finding of Carpenter and Kelly (2012) that testing effects are still observed when subjects are required to make a response to restudied items. This line of results supports the idea that the benefits of retrieval practice are due to the act of retrieving information from long-term memory and not simply to subjects making a response for tested material but not for restudied material.

Ruling out a verbal code

One alternative explanation for the absence of an improvement in memory quality is that subjects relied on coarse verbal labels for color rather than a fine-grained visual memory. For example, subjects could have remembered the name of a color and randomly responded within a section of the color wheel they associated with that label. Fortunately, the use of a continuous report measure allowed us to observe the relationship between an early retrieval attempt and a later retrieval attempt of the same item. This allowed us to calculate the correlation between response errors during initial test and response errors at delayed test for each subject. Indeed, the

direction of error in the observers' responses in the immediate recall task predicted the direction of error in the delayed recall task. This bias was reliably different from zero for retrieved items in all experiments [Experiment 1 (Z = 3.8, p <.001), Experiment 2 (Z = 4.1, p <.001), and Experiment 3 (Z = 4.2, p <.001), see Figure 4-7 for a visualization of aggregate responses]. These observations provide evidence against the use of a purely verbal code. If subjects were relying solely on a coarse verbal label and then guessing randomly within a section of the color wheel associated with that label when responding, we would expect early and late retrieval errors to be uncorrelated. However, we observed that delayed responses were biased by immediate responses. This bias suggests that subjects are able to maintain more than a coarse verbal representation of color. Furthermore, this is additional evidence for the idea that retrieval from memory is not a passive process and that subjects are reminded of, and re-encode past representations of an item when given a test (Hintzman, 2011). This response bias finding dovetails with the episodic context theory of retrieval practice (Lehman et al., 2014), which maintains that the context associated with a previously studied item is updated during the subsequent retrieval of that item to include features of both contexts. In this case fine-grained prior reports can be added as another layer of context that is encoded when an item is later retrieved.

Conclusion

Extant models of the retrieval practice effect have asserted that testing enhances the accessibility of learned associations rather than the fidelity of the retrieved memories (Carpenter, 2009; Lehman et al., 2014). The evidence for this assertion has been inconclusive, however, because of a heavy reliance on discrete word or picture stimuli that preclude a clear measure of item specific mnemonic precision. Here, we measured performance in a test that required recall

of colors from a continuous 360 degree space, and we used an analytic approach that enables distinct estimates of the probability of successful retrieval and the precision of the retrieved representations. The results were clear at both the aggregate and individual subject levels. Retrieval practice selectively enhances the probability of recall without improving mnemonic precision. Thus, even though both accessibility and fidelity can determine memory performance, the selective effect of retrieval practice on the former highlights the utility of distinguishing these aspects of memory function.

CHAPTER 4. ALPHA-BAND OSCILLATIONS TRACK THE RETRIEVAL OF PRECISE SPATIAL REPRESENTATIONS FROM LONG-TERM MEMORY

Introduction

Episodic memory is defined by the phenomenon of re-experiencing the details of past events, and is thought to be supported by the reactivation of neural activity that was present at encoding. In line with this view, functional magnetic resonance imaging (fMRI) studies have shown that sensory regions involved in the initial processing of information are re-engaged at retrieval (Wagner et al. 2005; Danker & Anderson, 2010), and that the voxel-wise patterns of activity within these regions resembles activity seen during encoding (Bosch, Jehee, Fernandez, & Doeller, 2014; Hindy, Ng, & Turk-Browne, 2016; Ritchey, Wing, LaBar, & Cabeza, 2013). In addition, more time-resolved measures of neural activity such as electroencephalography (EEG) and magnetoencephalography (MEG) have shown that retrieval-related neural activity echoes the broad strokes of encoding-related activity, such as the category of the paired associate (Jafarpour et al., 2014; Morton et al., 2013; Waldhauser et al., 2016; Wimber, Maaß, Staudigl, Richardson-Klavehn, & Hanslmayr, 2012) or the task performed at encoding (J. D. Johnson, Price, & Leiker, 2015). However, an open question is whether EEG activity can provide a temporally resolved means of tracking the retrieval of precise feature values that are associated with specific items.

Here, we addressed this question by measuring EEG activity during the encoding and recall of spatial information from long-term memory (LTM). To reconstruct the spatial representations in a precise manner, we applied an inverted encoding model (IEM) to the

topography of oscillatory activity on the scalp. IEMs have provided a useful approach for reconstructing precise spatial representations from fMRI and EEG activity (Foster et al., 2016; Foster, Sutterer, Serences, Vogel, et al., 2017; Sprague et al., 2014, 2016; Sprague & Serences, 2013). However, this approach has not yet been applied to the study of long-term memory; therefore, it is an open question whether it is possible to track retrieval of spatial long-term memories by applying an IEM to EEG activity. Furthermore, it is unclear which frequency bands would carry this spatially specific information. On the one hand, previous work using an IEM applied to alpha-band EEG activity, has successfully tracked covert spatial attention (Foster, Sutterer, Serences, Vogel, et al., 2017), and spatial representations maintained in working memory(Foster, Bsales, Jaffe, & Awh, 2017; Foster et al., 2016). Recent theories about the role of rhythmic oscillations in memory maintain that same frequencies of oscillations coordinate specific cognitive operations at encoding and retrieval (Siegel, Donner, & Engel, 2012; Watrous & Ekstrom, 2014; Watrous et al., 2015), predicting that alpha-band activity may play a similar role at retrieval. In addition, alpha-band activity has been shown to track hemifield-specific location memory (Stokes, Atherton, Patai, & Nobre, 2012; Waldhauser et al., 2016). On the other hand, other frequency bands, especially theta and beta, are known to play important roles in long-term memory encoding and retrieval (Hsieh & Ranganath, 2014; Morton et al., 2013; Morton & Polyn, 2017; Nyhus & Curran, 2010) and spatial navigation (Bohbot, Copara, Gotman, & Ekstrom, 2017; Watrous, Fried, & Ekstrom, 2011). It is also possible that a combination of these frequencies carry spatially specific information. Thus, we investigated whether precise spatially specific information is reinstated during long-term memory retrieval, and, if so, which frequency bands carry this information.

In two experiments, participants learned to associate objects with specific angular locations. Then, they were asked to precisely report the associated location when presented with an object cue. This continuous recall procedure enabled a fine-grained measurement of mnemonic performance, and recent work has shown that modeling of the response error distribution can provide robust indices of the probability and precision of the stored representations (Harlow & Yonelinas, 2014; Richter et al., 2016; Sutterer & Awh, 2016; Zhang & Luck, 2008). The IEM analysis revealed that spatially specific oscillatory activity tracked the retrieved locations after the presentation of the object cue. Consistent with oscillatory reinstatement accounts, we primarily observed spatially specific patterns of activity in the alpha band, just as observed in past studies of spatial working memory (Foster, Bsales, Jaffe, & Awh, 2017; Foster et al., 2016). Moreover, the alpha-band patterns observed during retrieval matched those observed during the initial encoding of the objects, in line with the hypothesis that encoding-related oscillatory patterns were reinstated during retrieval from LTM. Finally, the selectivity of alpha-band activity tracked memory performance as learning progressed as well as the latency with which participants reported the target locations. Together these findings suggest that LTM retrieval yields a reinstatement of the spatially specific oscillatory activity that is observed during encoding, and that multivariate analysis of alpha-band activity provides a powerful measure of the timing and success of this basic cognitive process.

Materials and Methods

Participants

Sixty-nine adults (33 in Experiment 1, and 36 in Experiment 2; 18–35 years old, 38 female) participated in the study for monetary compensation ($10 per hour in Experiment 1, and $15 per hour in Experiment 2). All participants reported normal or corrected-to-normal vision

and provided informed consent according to procedures approved by the University of Oregon Institutional Review Board (Experiment 1) and the University of Chicago Institutional Review Board (Experiment 2).

Participant exclusions for Experiment 1

For Experiment 1 participants were excluded for poor performance on the task and excessive EEG artifacts. One participant did not return for the second day of the experiment. One participant was excluded for poor performance on the first day (86.1° average response error across all day 1 tests) and data collection was terminated for one participant during the session for excessive artifacts. In addition, participants were excluded from further analysis if they had insufficient artifact-free trials (<550 trials). Artifact number exclusion criteria were set during data collection, but before the data were analyzed. Three participants were excluded due to excessive EEG artifacts. In the final sample, there were 27 participants in Experiment 1 (mean number of artifact-free trials = 799, *SD* = 85).

Participant exclusions for Experiment 2

For Experiment 2, our target final sample size was 24 subjects. Participants were replaced for poor task performance or if too many trials were lost due to recording or ocular artifacts. One participant was excluded for poor performance on LTM trials (87.1° average response error across all retrieval tests), and data collection was terminated for three participants during the session for excessive artifacts. In addition, participants were excluded from further analysis if they had insufficient artifact-free trials (<450 trials for encoding or retrieval). Artifact number exclusion criteria were set during data collection, but before the data were analyzed. We relaxed the exclusion criterion in Experiment 2 because we obtained fewer trials per condition. Eight participants were excluded due to excessive EEG artifacts. In the final sample there were 24

participants in Experiment 2 (mean number of artifact-free encoding trials = 535, *SD* = 46 and recall trials = 545, *SD* = 39).

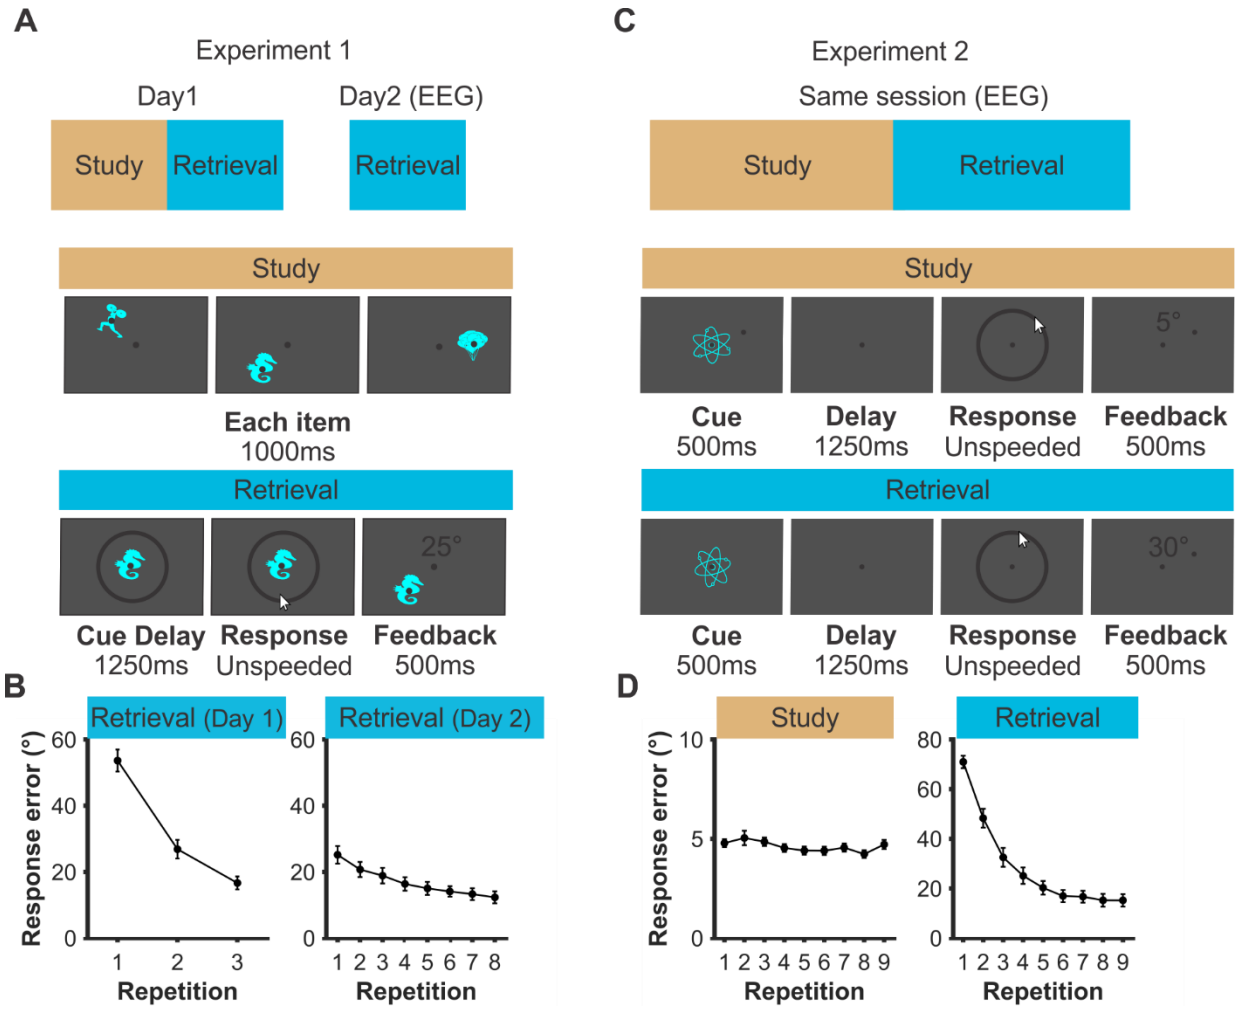

Figure 4-1 Task figure and memory performance for Experiments 1 and 2
A. Schematic of Experiment 1, with example study (Experiment 1, Day 1) and retrieval (Experiment 1, Day 2) trials. Each trial was initiated with a space press. **B.** Average absolute error of retrieval responses demonstrating improvements over retrieval repetitions. **C.** Schematic of interleaved study and retrieval repetitions in Experiment 2 with example study and retrieval trials. Each trial was initiated with a space press. **D.** Average absolute error of study and retrieval responses demonstrating memory accuracy over study and retrieval repetitions. Error bars represent ±1 s.e.m.

Apparatus

Stimuli were presented in MATLAB using Psychtoolbox (Brainard, 1997; Pelli, 1997) and were presented on a 17-in. CRT monitor (60 Hz) for Experiment 1 and on a 24-in. LCD monitor (120 Hz) for Experiment 2.

Experiment 1 task procedure

The experiment was comprised of two sessions run on consecutive days (Figure 4-1a). On Day 1, participants were instructed to learn 120 object-location associations (see Figure 4-1a for example clip art) as accurately as possible for the test the next day. On Day 2, participants were cued with the object and asked to recall and report the associated location while we recorded EEG data.

On Day 1, all 120 object-location pairings were studied over three repetitions with interleaved retrieval practice. Each of these repetitions were randomly divided into 12 "mini-blocks", in which 10 objects were presented followed by a final test on all objects in a random order. Specifically, 10 objects were serially presented in their respective spatial locations (1000 ms per object, each object initiated by pressing spacebar). Next, each of the 10 objects were presented at fixation in a random order (1000 ms per object), and participants clicked that object's location along a ring (unspeeded). Recall performance was assessed by calculating the response error (i.e., difference between the presented and reported location, ranging between –180° and 180°). After each response, participants were presented with the object in its correct location and the response error (500 ms). After completing these mini-blocks, participants again retrieved all 120 objects in a random order. One participant did not complete the final retrieval on the third run, and one participant accidentally aborted the experiment during the presentation of the first 10 objects before completing the rest of the session.

On Day 2, participants repeatedly retrieved the location of all 120 objects while we recorded EEG activity (Figure 4-1a). During each repetition (7–8 in total), the objects were presented in a random order. Each retrieval trial was initiated by a space press. After a variable interval of 1100 to 1500ms, an object was presented at fixation along with the response ring. Participants were instructed to maintain fixation and to avoid blinking or moving the mouse from trial initiation until the cursor appeared. Participants were also instructed to recall the location during the retrieval delay (1250 ms).

Experiment 2 task procedure

Experiment 2 was designed to examine encoding-retrieval similarity within a single session. As such, the experiment was modified to take place within one day by reducing the total number of objects (80 vs. 120). Participants were instructed to learn object-location associations as accurately as possible and that they would alternate between studying and being tested on these associations (Figure 4-1c).

During the study session, participants studied and then recalled each item during each trial. Each study trial was initiated by a space press. After a variable interval of 500 to 800 ms, an object was centrally presented (Figure 4-1c) along with a dot at the paired location (500 ms stimuli) followed by a blank delay (1250 ms). To prevent participants from using a part of the object as a reference to remember the associated location, we randomly varied the orientation of each object (-45 to 45°) for each presentation. As in Experiment 1, participants then reported the to-be-remembered location by clicking on the response ring (unspeeded). Participants were instructed to click with the left mouse button if they were confident in their response, and to click with the right mouse button if they felt that they were guessing. Both confident and guess responses were used for subsequent analyses. After each response, participants were shown the

correct location of the item along with a number denoting the magnitude of the error. After studying all 80 objects, participants underwent another retrieval test for all objects in a random order (Figure 4-1c). The only difference between study and recall trials was the presence of the peripheral dot.

Stimuli

In Experiment 1, 120 clip art objects (e.g., animals, plants, objects) were selected from the Sutterer and Awh (2016) clip art library. All objects were randomly assigned to unique angular locations (0–360°, 3° steps) for each participant. On Day 1, the viewing distance was ~80 cm (1.9° stimuli, 5° response ring, 0.3° fixation dot). On Day 2, the viewing distance was ~100cm (1.5° stimuli, 4° response ring, 0.25° fixation dot). The background of the screen was medium gray, all objects appeared in the color cyan, the response ring was dark grey, and the fixation dot was rendered in black.

In Experiment 2, 80 of the objects from Experiment 1 were randomly paired with a unique location drawn from all 360° of possible locations. In order to assure that the entire space was used, assignment of locations was constrained such that an equal number of locations were drawn without replacement from eight bins each spanning 45 degrees of the possible space. The viewing distance was ~100 cm (1.2° stimuli, 4° response ring, 0.25° fixation dot). The background of the screen was again medium gray, all objects appeared in the color cyan, and the response ring and the fixation dot were dark grey.

Modeling of Response Errors

Response error was measured as the number of degrees between the presented angular location and the reported angular location. Errors ranged from 0^0 (a perfect response) to $\pm 180^0$ (a maximally imprecise response). For each run (Figure 4-1b), we calculated the average

absolute response error for the artifact free trials. Error distributions of this sort have been shown to be well described by a mixture of a uniform distribution for guesses and a Von Mises distribution for correct responses (Brady et al., 2013; Zhang & Luck, 2008). We used MemToolbox (Suchow et al., 2013) to calculate the probability of retrieval (P_{mem}), precision (SD), and the bias (μ) of each participants responses.

EEG acquisition

In Experiment 1, EEG was recorded with 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, OR, PO3, PO4, POz). 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, electrodes were placed ~1 cm from the canthi of each eye to record horizontal electrooculorgram (EOG). To detect blinks and vertical eye movements, a single electrode was placed under the center of the right eye and referenced to the left mastoid to record vertical EOG. The EEG and EOG data were amplified with an SA Instrumentation amplifier, filtered (0.01–80 Hz), and digitized (250 Hz) using LabVIEW 6.1 running on a PC.

In Experiment 2, EEG was recorded from 30 active Ag/AgCl electrodes (Brain Products actiCHamp, Munich, Germany) mounted in an elastic cap positioned according to the International 10-20 system Fp1, Fp2, F7, F3, F4, F8, Fz, FC5, FC6, FC1, FC2, C3, C4, Cz, CP5, CP6, CP1, CP2, P7, P8, P3, P4, Pz, PO7, PO8, PO3, PO4, O1, O2, Oz. A ground electrode was placed in the elastic cap at position FPz. Data were referenced online to the right mastoid and re-referenced offline to the algebraic average of the left and right mastoids. Incoming data were filtered (0.01– 80 Hz) and recorded with a 500 Hz sampling rate using BrainVision Recorder

running on a PC. To detect eye movements and blinks, we used eye tracking to monitor gaze position and electrooculogram (EOG) activity recorded with five electrodes (~1cm from the outer canthi of each eye, above/below the right eye, and a ground electrode placed on the left cheek).

Artifact Rejection

Data from both experiments were visually inspected for EOG and EEG artifacts. Trials containing blinks, eye movements, blocking, and muscle artifacts were excluded from analysis. One electrode for one participant in Experiment 2 was also rejected during recording because it had malfunctioned. We also monitored gaze position during Experiment 2 using a desk-mounted infrared eye tracking system (EyeLink 1000 Plus, SR Research, Ontario, Canada). Gaze position data for Experiment 2 were also visually inspected for ocular artifacts. For the analysis of gaze position, we further excluded trials in which the eye tracker was unable to detect the pupil, operationalized as any trial in which the horizontal gaze position was more than 15° from fixation or the vertical gaze position was more than 8.5° from fixation. We collected useable gaze position data (500 Hz sampling rate) for 18 of 24 participants.

Removal of trials with ocular artifacts was effective: maximum variation in grand-averaged HEOG waveforms by remembered location bin was < 2.5 μV for Experiment 1 and < 2 μV for both the encoding and retrieval in Experiment 2. Thus, eye movements in both experiments corresponded to variations in eye position of < 0.2° of visual angle (Lins, Picton, Berg, & Scherg, 1993), roughly the size of the fixation dot. Analysis of the subset of participants (18) for whom we were able to obtain reliable gaze position data in Experiment 2 corroborates the HEOG data obtained from all participants. Variation in grand-average horizontal gaze position as a function of remembered location was < 0.11° for encoding and < 0.08° of visual

angle for retrieval. Variation in grand-average vertical gaze position data by remembered location was $< 0.14°$ for encoding and $< 0.09°$ of visual angle for retrieval. For comparison, HEOG for these participants showed a < 2.1 μV maximum variation which also corresponds to $< 0.2°$ of visual angle.

Time-frequency analysis

To calculate frequency specific activity at each electrode we first band-pass filtered the raw EEG data using EEGLAB (eegfilt, see Delorme and Makeig, 2004). Alpha band analyses were band-pass filtered between 8 to 12 Hz, which is consistent with our prior work (Foster et al., 2016). For our exploratory analysis of the full range of frequencies, we band-pass filtered the data at 1 Hz intervals (4–50 Hz, down-sampled to 20 Hz, filter order:

$$3 \times \frac{sampling\ rate}{low-pass\ cutoff}$$

We then applied a Hilbert transform (MATLAB Signal Processing Toolbox) and squared the complex magnitude of the complex analytic signal for each trial to calculate instantaneous power before averaging across trials.

Inverted encoding model

Following our prior work (Foster et al., 2016), we reconstructed spatially selective channel-tuning functions (CTFs) from the multivariate topographic distribution of oscillatory power across electrodes. We assumed that the power at each electrode reflects the weighted sum of eight spatially selective channels (which we assume reflect the responses of neuronal populations), each tuned for a different angular location (Brouwer & Heeger, 2009; Foster et al., 2016; Sprague et al., 2015; 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 (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 location bins. These 8 location bins each spanned 45° and were centered on 22.5°, 67.5°, 112.5° etc for Experiment 1 and on 0°, 45°, 90° etc for Experiment 2. Bin centers for each experiment were chosen prior to data collection.

An IEM routine was applied to each time point in the alpha-band analyses and to each time-frequency point in the time-frequency analyses. We partitioned our data into independent sets of training data and test data (for details see the Assigning trials to training and test sets section). 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 measurements) 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 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 × 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 (i.e., 0° on the "Channel Offset" axis of Figure 4-2a) by aligning the estimated channel responses to the channel tuned for the cued/target location to yield the CTF averaged across the eight remembered locations.

Finally, because the exact contributions of each spatial channel to each electrode (i.e., the channel weights, *W*) varies 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 (Foster et al., 2016; Foster, Sutterer, Serences, Vogel, et al., 2017; Sprague et al., 2015).

Assignment of trials to training and test sets

Artifact free trials were partitioned equally into three independent sets to be used as training and test data for the IEM procedure (see Inverted Encoding Model). We down-sampled the data so that each set contained an equal number of trials, and that each location bin within a set also contained the same number of trials. For each of these sets we averaged power across trials for each location bin. We used a cross validation routine such that two sets of estimated power served as the training data and the remaining set served as the test data. We applied the IEM routine using each of the three matrices as test data, and the remaining two matrices as training data. The resulting CTFs were averaged across each test set.

For the analysis in which we ruled out the possibility that the IEM was detecting object-specific information, we assigned all trials with the same object to the same partition. After completing this additional step, we equated trials across sets and bins in the same manner described above**.**

For analyses in which we examined how within participant changes in selectivity related to behavior, we first down-sampled to equate the number of trials assigned from each location across conditions. After completing this additional step, we equated trials across sets and bins in the same manner described above. Finally, we employed the same training procedure described above (2/3 of the total data), but split the final test set into our comparisons of interest. Thus, we used the same training data for both conditions and only the test data varied for each comparison.

For analyses that assessed relationships between CTF selectivity and behavior across participants we down-sampled the number of trials assigned to each location bin for each of the three sets to be equal to the smallest number of trials assigned to each bin in each set for any participant. This down sampling approach precluded individual differences in CTF selectivity driven by the number of the trials included in the analysis for each participant.

In Experiment 2, we sought to compare encoding and retrieval related activity. We closely followed the procedure that examined retrieval-related activity alone, by training on 2/3 of the encoding data and testing on 1/3 of the retrieval data. By maintaining these same ratios of training to test data, we could more directly compare the results from encoding and retrieval.

Resampling random assignment

To avoid spurious results due to the random assignment of trials, we repeated each analysis multiple times with a different random assignment of trials. When comparing between conditions, we conducted 500 iterations per time point. When comparing against a permuted null

distribution (which is a time consuming procedure), we conducted 10 iterations per time point, given the computational time needed for each analysis. In order to decrease computation time further for the 4–50 Hz time-frequency analysis, we down sampled the data matrix of power values to 50 Hz (i.e., one sample every 20 ms). We down sampled after calculating power so that down sampling did not affect our calculation of power. The data matrix was not down-sampled for analyses restricted to the alpha band.

Calculating CTF Selectivity

To quantify selectivity at each time point we calculated the slope of the CTF via linear regression. We collapsed across channels of equidistance (e.g., ±2 bins). As such, higher slope values indicate greater CTF selectivity while lower values indicate less CTF 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. To generate our null distribution, we randomly shuffled the remembered location labels in each training/test set so that the labels were random with respect to the observed responses at each electrode. We then repeated 1000 iterations of this randomization procedure to obtain a null distribution of *t* statistics at each time point.

Finally, to test whether CTF selectivity was reliably above chance we employed a nonparametric cluster approach that corrects for multiple comparisons by taking into account auto-correlation in time and frequency (Cohen, 2014; Maris & Oostenveld, 2007). Specifically, we applied a *t*-value threshold corresponding to $p < .05$ (Experiment 1: $t = 1.706$; Experiment 2: $t = 1.714$) to identify clusters of pixels (time and frequency analysis) or adjacent time points

(alpha only analysis). At the same time, we applied the same threshold to each permutation and calculated the largest summed-*t* statistic for any cluster in the permutation, resulting in a distribution of maximal summed *t*-statistics for our permuted null distribution. Finally, the sizes of the significant clusters of the non-permuted data were thresholded such that only clusters larger than the 95th percentile of the permuted distribution were considered reliable (Type 1 error less than .05). Therefore, our cluster test was a one-tailed test, corrected for multiple comparisons.

Resampling test

When comparing between conditions (i.e., baseline alpha power, CTF slope), we used a non-parametric resampling procedure across participants (Efron & Tibshirani, 1993). We resampled each participant with replacement 100,000 times. Then, we calculated the number of these resampling iterations in which the differences were ≤ 0 (one-tailed). In cases where no iterations were ≤ 0, we report the p values as $p < .001$. We deemed results to be reliably above chance if $p < .05$.

Results

Experiment 1

Experiment 1 was designed to test whether continuous tests of memory accuracy, in conjunction with an IEM applied to EEG data, could be used to track memory retrieval. The design includes two important properties. First, we used a continuous test of memory accuracy by having participants report remembered locations along a ring. This provides a sensitive test of memory accuracy as the deviation from the correct location. Second, we recorded EEG activity during memory retrieval for the purposes of building and evaluating an encoding model. This

inverted encoding model (IEM) can track memory retrieval as a graded function of spatial location.

Behavioral performance.

On Day 1, participants studied 120 object-location associations (Figure 4-1a). On Day 2, participants returned for a retrieval session in which we recorded EEG. Participants received feedback based on their response error (−180° to 180°). During both days, their performance improved (Figure 4-1b). During Day 1, average response error improved significantly from the first (M = 53.6°, SD = 17.7°) to the final test (M = 17.6°, SD = 11.1°) of the session ($t(26)$ = 13.2, p <.001, one-tailed). As a result on continued feedback, memory also improved from the first (M = 25.2°, SD = 14.1°) to the final test (M = 12.7°, SD = 9.3°) during the second session ($t(26)$ = 7.3, p <.001, one-tailed).

Alpha-band (8–12 Hz) topography tracks spatial representations retrieved from LTM.

In Experiment 1, we tested whether oscillatory EEG activity tracks the time-resolved retrieval of precise spatial memories. Because we have previously found that alpha-band activity tracks spatial locations held in working memory (Foster et al., 2016), we were *a priori* interested in whether alpha-band power would also track locations retrieved from long-term memory. Thus, we used an IEM to test whether the multivariate topography of alpha-band power tracked locations retrieved from long-term memory. If the pattern of alpha-band power contains spatially selective information about the remembered location, we would expect to see a graded channel tuning function (CTF) with a peak response in the channel tuned for the remembered location (a channel offset of 0°in Figure 4-2) following the retrieval cue. This graded pattern can be quantified as slope across the position channels as distance from the retrieved location increases.

A slope of zero reflects no spatial selectivity in the CTF, while a positive slope reflects spatial selectivity for the location associated with the cue. To test this hypothesis, we conducted a permutation test (see Materials and Methods) to determine at which time points we observed a CTF slope that was reliably above zero. We detected reliable selectivity for spatial information (i.e., slopes > 0) that was sustained during the retrieval interval (588 – 1250ms; Figure 4-2a).

One possibility is that this graded tuning is an artifact of our selection of a graded basis set (Ester, Sprague, & Serences, 2015; Foster et al., 2016; Saproo & Serences, 2014). To investigate this possibility, we reran this analysis using a delta function basis set that predicts a peak response in the preferred channel and no response in adjacent channels. If the topography of alpha power represents spatial locations in a graded manner, we would still expect a graded pattern of responses. Instead, if the observed results were driven by our selection of a basis set, we would expect a peak response in the correct bin and a little to no response in all other bins. Using a delta function basis set, we observed a graded pattern of responses across remembered locations (Figure 4-2b) that is similar to the pattern of activity we see when we apply the standard basis set (Figure 4-2c). This suggests that our results are not an artifact of our selection of a basis set, but reflect a real graded tuning profile during the retrieval of spatial memories.

Although the aggregate results revealed that channel activity peaked at the remembered location and dropped in a graded fashion as the distance form that location increased, this analysis did not establish that this orderly pattern was present at each location. Indeed, a courser hemi-field or quadrant-based signal could produce such a pattern. If alpha-band activity precisely tracks retrieved spatial locations, we should observe a graded pattern for each remembered location. We examined the average channel response during time points where we previously observed reliable spatial selectivity (588–1250ms) for eight location bins separately. The channel

response for each location revealed graded information throughout the same window, and the channel response for all locations were reliable (All slopes > 0.05, all *p's* < .002). Therefore, alpha-band CTFs track memory retrieval of a precise spatial location.

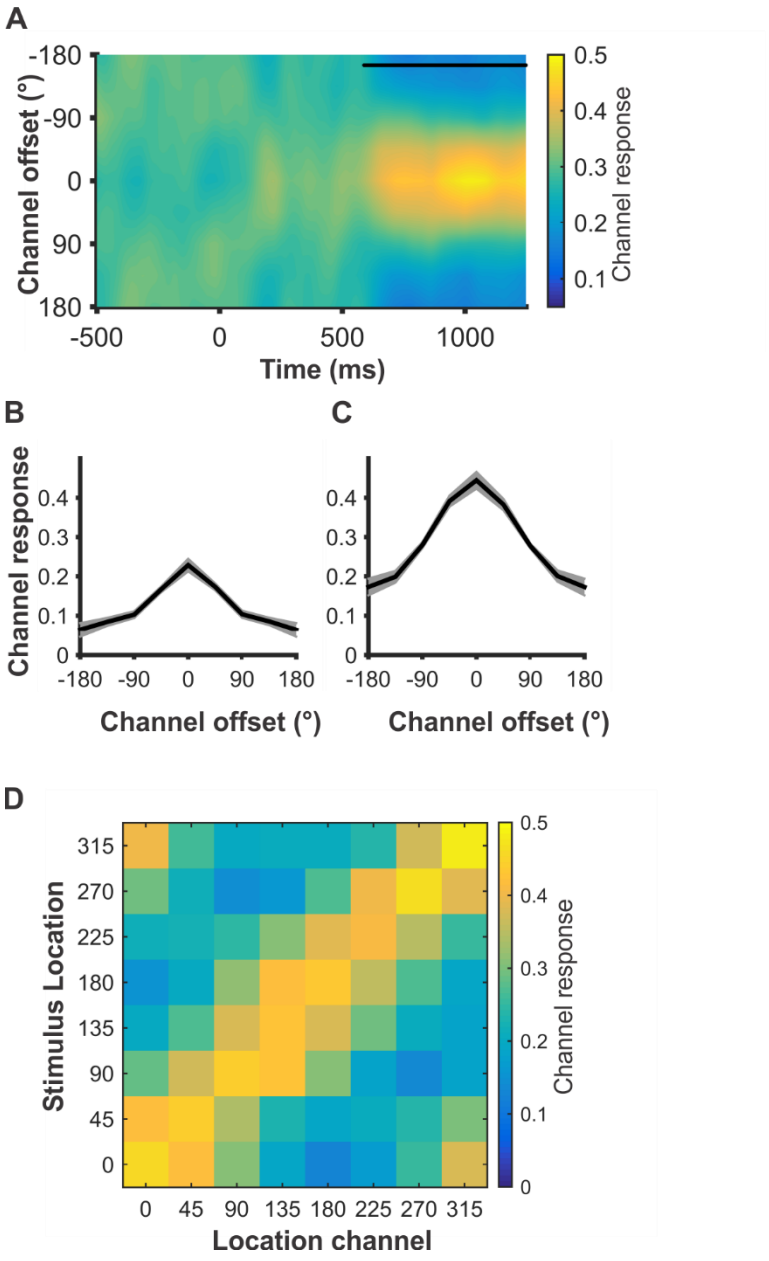

Figure 4-2 Alpha band (8–12 Hz) channel tuning functions (CTF) from Experiment 1
A. Alpha CTF across time. An IEM was used to reconstruct spatially selective CTFs from the topographic distribution of alpha-band power. CTF selectivity was reliable from 588 to 1250ms (quantified as CTF slope; $p < .05$, indicated by the black marker). **B.** Alpha CTF derived with a set of 8 delta functions and averaged across significant time points (588–1250 ms). Delta

function CTFs are graded confirming that the signal carried by the topography of alpha-band power is intrinsically graded. Thus, the use of a graded basis set is appropriate. Shaded area represents ±1 bootstrapped s.e.m. **C.** Alpha CTF derived with a graded basis set and averaged across significant time points (588–1250 ms). Shaded area represents ±1 bootstrapped s.e.m. **D.** Channel responses for each of the eight stimulus location bins averaged across significant time points (588–1250ms). The channel response peaks at the channels preferred location, indicating that alpha activity is selective for the specific remembered location.

Identifying frequencies that track the retrieval of spatial location

A motivating question for the present work was whether spatially selective information was specific to alpha-band activity. On the one hand, prior work has found that alpha-band activity selectively tracks spatial locations that are covertly attended (Foster, Sutterer, Serences, Vogel, et al., 2017) or held in working memory (Foster et al., 2016). On the other hand, theta-band (4–7 Hz) and beta-band (16–25 Hz) activity are known to play an important role in long-term memory (Morton et al., 2013; Nyhus & Curran, 2010). Therefore, we performed the same IEM analysis at each frequency and time point from 4–50 Hz to test whether other frequency bands also carried spatially selective information about the remembered location. We conducted a permutation test at each frequency and time point and used a cluster correction to identify frequencies with CTFs that were reliably above zero (Figure 4-3a). Although we observed brief periods of spatial selectivity in the beta range (16–25 Hz) the most robust and sustained selectivity was in the alpha band (410ms – 1250ms).

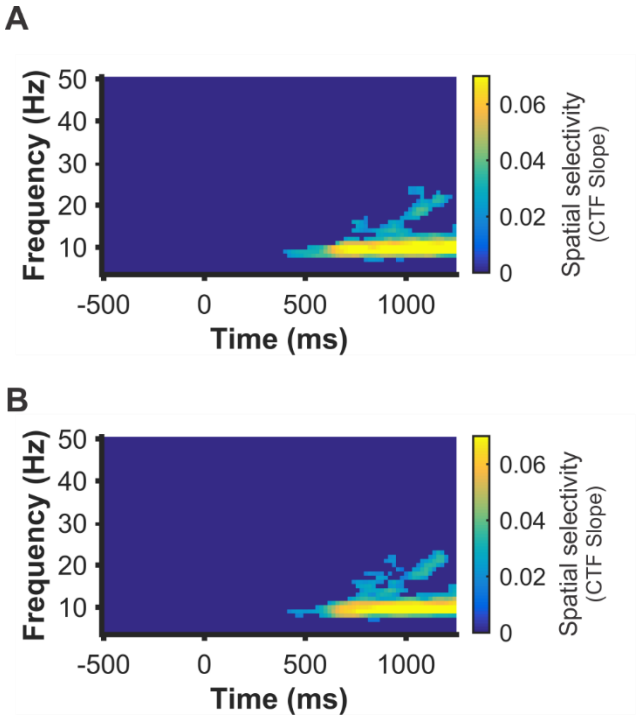

Figure 4-3 Identifying frequencies that track retrieved locations for Experiment 1
A. An IEM was used to reconstruct spatially selective CTFs from the topographic distribution of total power across a range of frequencies (4–50Hz). Alpha and to a lesser extent beta power tracked retrieval of spatial information. **B.** Training and testing across shapes. Alpha power continued to track the retrieval of spatial information when the IEM was trained and tested on separate shape cues, indicating that CTFs reflect remembered locations not the retrieval cue. Points at which CTF slope values were not reliably above zero as determined by a cluster corrected permutation test (p <.05) were set to dark blue.

Spatially selective alpha-band activity generalizes across visual objects associated with the same spatial location

For each participant, each object was associated with a unique location such that object and position were confounded within this analysis. Thus, it is possible that the selectivity we observed across some or all frequencies, reflects patterns of activity elicited by the cue rather than activity related to the retrieval of a spatial position. To investigate this, we re-ran the analysis while ensuring that distinct items were included in the training and test sets (see Materials and Methods). Despite this constraint, we observed similar results (Figure 4-3b),

confirming that the sustained spatial selectivity we observed reflected the position associated with each cue rather than the cue itself.

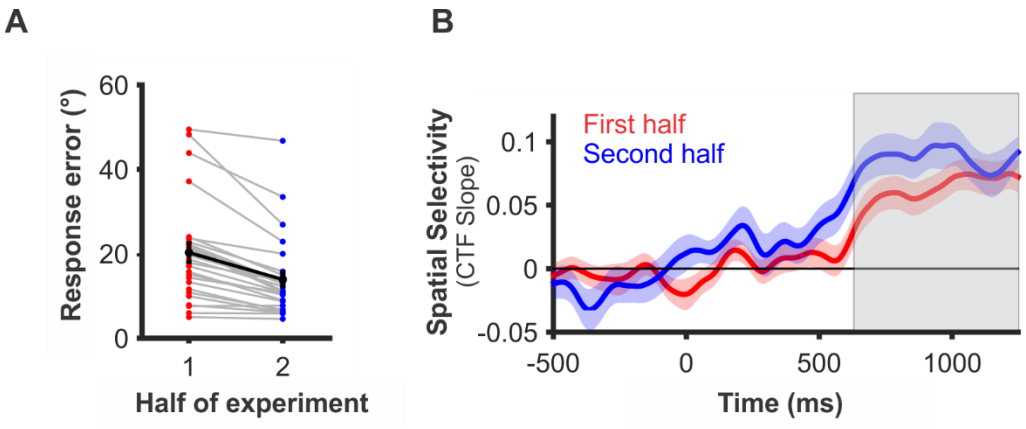

Figure 4-4 Assessing the relationship between alpha selectivity and memory performance for Experiment 1
A. Average response error improved from the first to the second half of the experiment (p <.001). Error bars represent ±1 s.e.m. **B.** Time resolved CTF slopes for trials from the first and second half of the experiment. CTF slope reflects learning across the session and reveals significantly higher spatial selectivity for the second half of the experiment relative to the first half (p <.001). Reliable differences were assessed by averaging across time points where we observed reliable CTFs for all trials (grey box) and comparing CTF slope between the first and second half of the experiment. Shaded error bars represent ±1 bootstrapped s.e.m.

Spatially selective alpha-band activity tracks the accuracy of recall from long-term memory

A consequence of providing feedback during Day 2 is that memory performance improved throughout the session. To examine whether alpha-band CTFs tracked these behavioral improvements, we split the test data into the first half and second half of trials. Behaviorally, we observed that memory performance improved from the first half (M =20.3°, SD = 11.9°) to the second half (M = 13.9°, SD = 9.4°) of the session ($t(26)$ = 7.2, p <.001, one-tailed, Figure 4-4a). Furthermore, a mixture model was used to assess whether these decreases in

average response error were driven by changes in the probability of retrieval and/or mnemonic precision (see Materials and Methods). During Day 2, the probability of retrieval (P_{mem}) increased over time (first half: M = 86.3%, SD = 13.6%; second half: M = 93.6%, SD = 9.5%; $t(26) = -6.3$, $p < .001$, one-tailed) and mnemonic precision improved over time (first half: M = 13.6°, SD = 3.7°; second half: M = 12.3°, SD = 4.1°, $t(26) = 5.52$ $p < .001$, one-tailed). If alpha-band CTFs are sensitive to the accuracy of memory retrieval, we would expect greater spatial selectivity in the second vs. first half of the session. To test this prediction, we used a resampling test (see Materials and Methods) in which we isolated the time points where aggregate data revealed significant alpha CTFs (Fig 2a; 588ms after cue onset until the response), and then compared average CTF slope across the first and second halves of the study. Indeed, spatial selectivity was significantly higher for the second half (CTF slope, M = 0.085, SD = 0.061) relative to the first half (M = 0.06, SD = 0.055) of the experiment (Figure 4-4b) ($p < .001$, one-tailed resampling test).

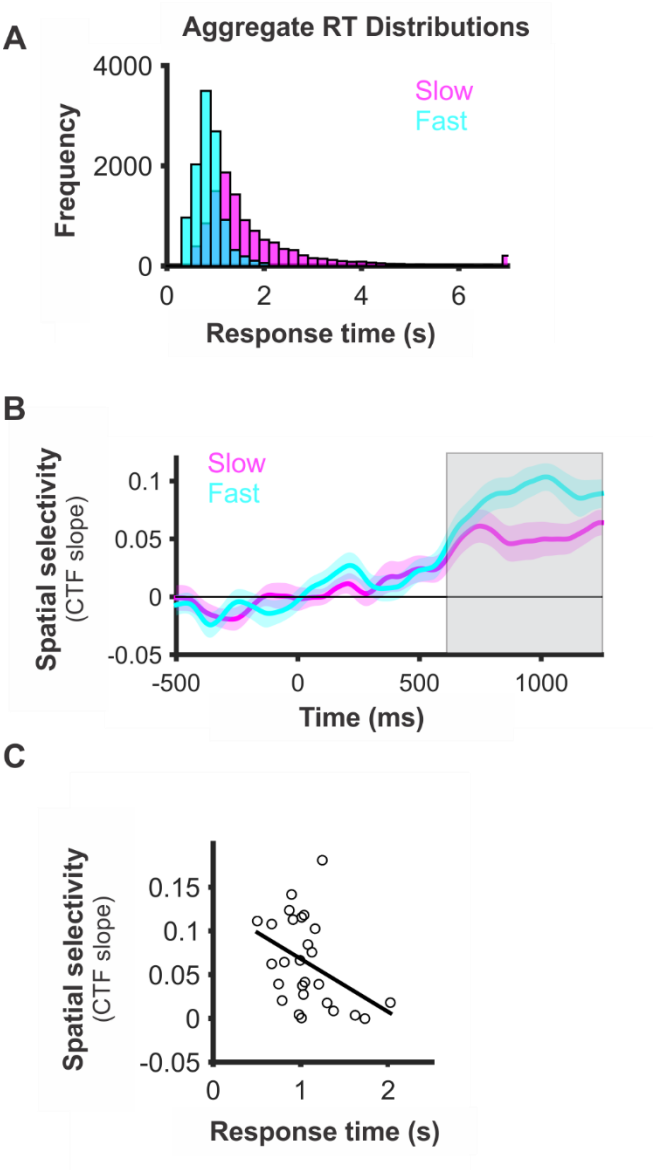

Figure 4-5 Assessing the relationship between alpha selectivity and response times for Experiment 1

A. Aggregate distribution of all participants' fast and slow response times. Response times > 7s are represented in the last bin of the histogram. **B.** Time resolved CTF slope for trials with the fastest and slowest response times. CTF slope reflects response latency and reveals that spatial selectivity was higher for trials when participants responded quickly (p <.001). Reliable differences were assessed by averaging CTF slope across time points where we observed reliable CTFs for all trials (grey box) and comparing splits with a resampling test. Shaded error bars represent ±1 s.e.m. **C.** Alpha selectivity is modestly correlated with response times across subjects although the relationship is not significant ($\rho(26)$= -.36 p = .07).

This reveals that alpha-band activity tracks the improvement in memory performance across learning episodes. Finally, CTF selectivity across the same window did not predict between-subject variations in the accuracy of recall ($\rho(26) = -.11$, $p = .6$). This null result could have numerous explanations but here we offer one hypothesis. While we instructed participants to immediately recall the location that corresponded to the object cue, it may be that some participants waited longer than others to call the correct location to mind while other participants relied on a more prospective strategy in which they immediately recalled the target location. This kind of strategic difference could yield large differences in mean CTF slope that may have been unrelated to whether the critical item could be retrieved. Indeed, the response time analysis in the next section lends further plausibility to this hypothesis.

Spatially selective alpha-band activity tracks within- and between-subject variations in response latency

The latency of memory retrieval varied across trials and participants to a large extent (see Fig 5a). To examine whether alpha-band CTFs tracked these behavioral differences in response time (RT), we split the test data into two halves based on the median RT (average fast RT: M = 854ms, SD = 240ms; average slow RT: M = 1961ms, SD = 1055ms). If alpha-band CTFs track the latency of memory retrieval, we would expect greater location selectivity on trials in which participants responded more quickly. Indeed, location selectivity was significantly greater when participants responded more quickly (M = 0.085, SD = 0.054) than when they responded slowly (M = 0.052, SD = 0.060; $p < .001$, one-tailed resampling test; Figure 4-5b). This pattern supports the hypothesis that participants responded more quickly when they had already retrieved the spatial information prior to the onset of the response cue, yielding a higher level of CTF selectivity during trials with faster responses.

Across participants, we observed substantial variation in median RTs (range = 504 – 2025 ms). To examine whether alpha-band CTFs tracked these individual differences in behavior, we tested whether there was a correlation between median RT and the selectivity of alpha-band CTFs (measured as slope). We predicted that participants who responded more quickly (i.e., faster RTs) would also have greater spatial selectivity (i.e., higher CTF slope). We observed a trending negative relationship between RT and CTF slope, as predicted ($r = -.36$; $p = .07$; Figure 4-5c). In addition to reflecting differences in the immediate accessibility of spatial memories, this relationship could also be driven by individual differences in the extent to which participants engaged in prospective retrieval during the delay interval. This is our working hypothesis, given that the differences in response latency seemed too large to reflect differences in the immediate accessibility of the spatial memories alone.

Experiment 2

In Experiment 2, we replicated and extended Experiment 1 in two important ways. First, to further examine the relationship between alpha-band selectivity and memory performance, we recorded EEG data throughout the learning process, including the first retrieval attempts. Second, we recorded EEG during both encoding and retrieval, which allowed us to test the extent that retrieval-related oscillatory activity resembled encoding-related oscillatory activity.

Behavioral performance

During a single session, participants learned 80 object-location associations (Figure 4-1c) with interleaved study and retrieval. During study trials, participants actively maintained the associated spatial location over a 1250 ms delay interval. During retrieval trials, participants had to retrieve the associated spatial location from long-term memory. During study trials, memory performance was very accurate and improved modestly but reliably from the first half ($M = 4.7°$,

SD = 1°) to the second half (*M* = 4.4°, *SD* = .9°) of the session (*t*(23) = 2.42, *p* = .012, one-tailed; Figure 4-1d). Mixture modelling revealed that this change was due to an improvement in mnemonic precision (first half: *M* = 5.8°, *SD* = 1.2°; second half: *M* = 5.4°, *SD* = 1.1°; *t*(23) = 2.28, *p* = .016, one-tailed) while no change was observed for probability of retrieval (first half: *M* = 99.9%, *SD* = .29%; second half: *M* = 99.9%, *SD* = .16%; *t*(23) = -.81, *p* = .21, one-tailed), which was at ceiling. For the LTM retrieval trials, we observed a substantial improvement in memory performance across the session as learning progressed. Memory error decreased from the first half (*M* = 40.8°, *SD* = 14.0°) to the second half (*M* = 16.2°, *SD* =11.9°; *t*(23) = 17.0, *p* < .001, one-tailed; Figure 4-6a). We replicated our finding in Experiment 1 that the reduction in memory error was driven by both an increase in the probability of retrieval (first half: *M* = 61.1%, *SD* = 16.0%; second half: *M* = 89.9%, *SD* = 13.7%; *t*(23) = -15.4, *p* < .001, one-tailed) and an improvement in mnemonic precision (first half: *M* =13.8°, *SD* = 4.35°; second half: *M* = 11.3°, *SD* =2.76°; *p* <.001, one-tailed). Thus, long-term memory improved throughout the session as participants learned the object-location associations.

One goal for Experiment 2 was to create a larger range of performance throughout the session in which EEG data was recorded. In line with this goal, we observed a much larger range in mean response error in Experiment 2 (71.0° Run 1 – 15.2° Run 9; Figure 4-1d) than in Day 2 of Experiment 1 (25.2° Run 1 – 12.3° Run 8), giving us the opportunity to apply the IEMs approach across the full trajectory of learning.

Spatially selective alpha-band activity tracks the accuracy of recall from long-term memory

In Experiment 1, alpha-band CTFs tracked retrieval of spatial locations from long-term memory. Furthermore, spatial selectivity of alpha-band CTFs increased as memory accuracy

improved (Figure 4-4). Experiment 2 was designed to replicate and extend those results over a larger range of behavior. We predicted that alpha-band CTFs would demonstrate higher selectivity when memories were more accurate. In line with this prediction, the average selectivity (i.e., CTF slopes) was larger in the second half of the session (M = 0.048, SD = 0.032) than in the first half (M = 0.012, SD = 0.022; $p < .001$; one-tailed; Figure 4-6b). Note, for this and all subsequent average CTF analyses, we averaged from 588 ms (the starting time point used in Experiment 1) until the onset of the response cue. As in Experiment 1, CTF slope did not track memory performance between participants ($\rho(23) = -.14$, $p = .52$). Thus, the spatial selectivity of alpha activity tracked broad improvements in recall accuracy across the session.

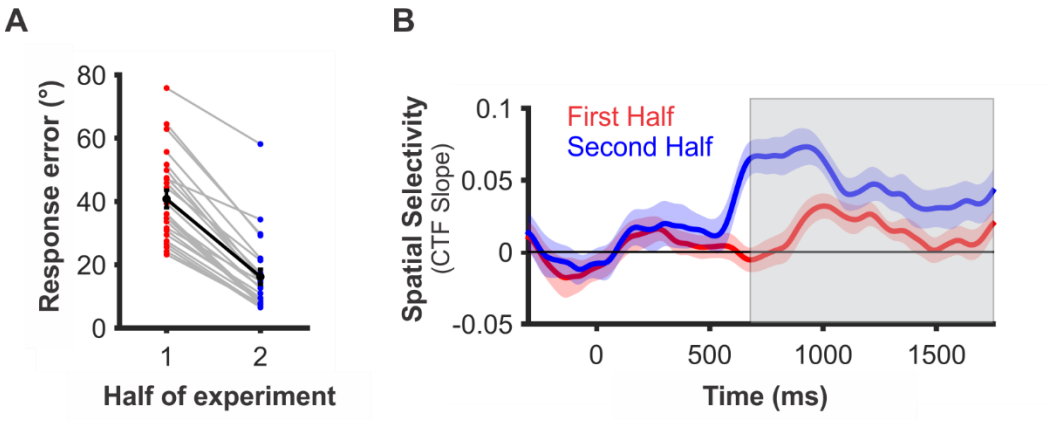

Figure 4-6 Assessing the relationship between alpha selectivity and memory performance for Experiment 2
A. Average response error improved from the first to the second half of the experiment ($p <$ .001). Error bars represent ±1 s.e.m. **B.** Time resolved CTF slopes for trials from the first and second half of the experiment. CTF slope reflects learning across the session and reveals significantly higher spatial selectivity for the second half of the experiment relative to the first half ($p <$.001). Reliable differences were assessed by averaging across time points where we observed reliable CTFs for all trials (grey box) and comparing CTF slope between the first and second half of the experiment. Shaded error bars represent ±1 bootstrapped s.e.m.

Spatially selective alpha-band activity tracks response latency

As in Experiment 1, we found that the selectivity of alpha-band CTFs tracked within- and between-subject variations in RT (Figure 4-7). A median split on RT revealed greater spatial selectivity for trials with fast RTs (CTF slope: $M = 0.035$, $SD = 0.03$) than trials with slow RTs ($M = 0.021$, $SD = 0.018$), $p = .005$, one-tailed; Figure 4-7b). We also replicated our finding that participants with faster RTs showed greater spatial selectivity of alpha-band CTFs ($r = -.49$; $p = .02$; Figure 4-7c). This link between CTF slope and RTs may reflect strategic differences between participants who prospectively recalled the associated location quickly following cue onset and those that waited until closer to the response window to bring that information to mind.

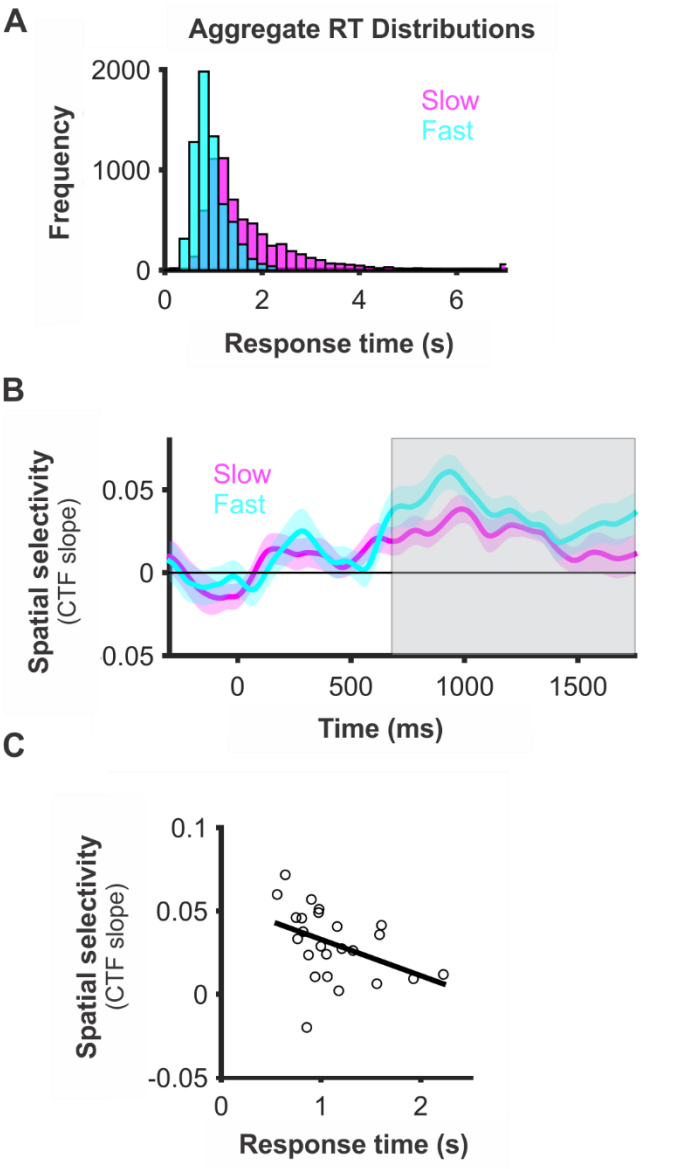

Figure 4-7 Assessing the relationship between alpha selectivity and response times for Experiment 2

A. Aggregate distribution of all participants' fast and slow response times. Response times > 7s are represented in the last bin of the histogram. **B.** Time resolved CTF slope for trials with the fastest and slowest response times. CTF slope reflects response latency and reveals that spatial selectivity was higher for trials when participants responded quickly (p=.005). Reliable differences were assessed by averaging CTF slope across time points where we observed reliable CTFs (grey box) for all trials and comparing splits with a resampling test. Shaded error bars represent ±1 bootstrapped s.e.m. C. Alpha selectivity is modestly correlated with response times across subjects ($\rho(23)$= -.49 p = .02).

Comparing frequency specificity at encoding and retrieval

In Experiment 1, we found that oscillatory activity in the alpha band (8–12 Hz) tracked retrieved locations following a memory cue. In Experiment 2, we replicated this finding, with cluster corrected permutation tests showing that primarily oscillations between 8 and 12 Hz, and to a lesser extent oscillations between 12 and 18 Hz, tracked the retrieved location (~500–1250ms; Figure 4-8a). Note, that in order to obtain the most robust measurement of spatially sensitive frequencies at retrieval, we only tested our IEM on trials from the second half of the experiment when memory performance and spatial selectivity were highest (Figure 4-6a). For consistency we applied the same approach to study trials (Figure 4-6b). Applying the IEM to study trials revealed that spatially selective information was represented across a wider range of low frequencies (Figure 4-8b; 4–8 Hz; 0–500ms; 25–30 Hz, 500–600ms; 8–20 Hz, ~200–1750ms). Although we observed the most sustained spatial selectivity in the alpha band (8–12 Hz). These results replicate past work that has shown that alpha-band activity tracks locations held in working memory (Foster et al., 2016). Finally, an overlay plot of frequency bands carrying spatially specific information in both the encoding and retrieval tasks (Figure 4-8c), revealed considerable overlap in the 8–12 Hz band across conditions. Together these findings suggest that the range of frequencies carrying spatially specific information during encoding is

similar to the range of spatially specific frequencies in a WM task as well as during LTM retrieval.

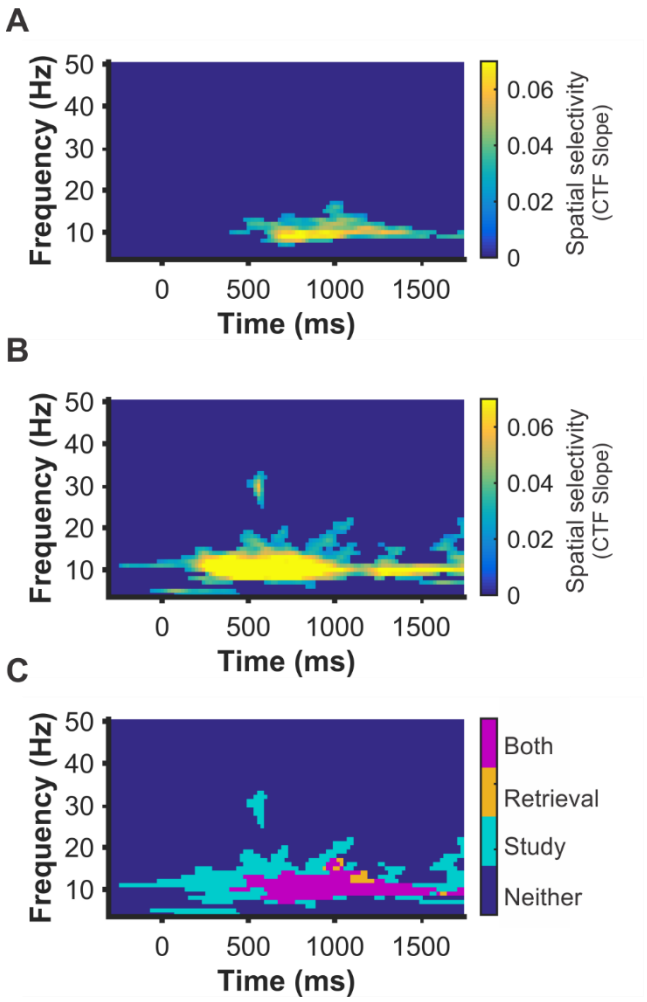

Figure 4-8 Identifying frequencies that track encoded and retrieved locations for Experiment 2

An IEM was used to reconstruct spatially selective CTFs from the topographic distribution of total power across a range of frequencies (4–50Hz) at retrieval and study. To ensure robust retrieval memory performance, only trials from the second half of the session were used in the test set for this analysis. **A.** Primarily alpha power tracked spatial information during retrieval trials. **B.** Initially, a broad range of frequencies tracked the encoded location (4–35Hz) while primarily alpha power tracked the remembered location through the entire delay. **C.** Overlay plot of significant activity during retrieval and study. Teal points reflect reliable spatial selectivity during study, orange points reflect reliable selectivity at retrieval, and magenta points reflect overlap selectivity at study and retrieval. Points at which CTF slope values were not reliably above zero as determined by a cluster corrected permutation test (p <.05) were set to dark blue.

Patterns of alpha-band activity generalize across encoding and retrieval

While the same frequency band carried spatially spatial information during both study and retrieval, this does not necessarily mean that the multivariate patterns of activity corresponding to each location are also similar during encoding and retrieval. To provide a comprehensive test of encoding-retrieval similarity, we trained the IEM using study trials and tested the model on retrieval trials. We observed robust spatial selectivity throughout the retrieval interval (520–1750ms; $p < .05$; Figure 4-9). This provides evidence that the multivariate pattern of alpha activity during retrieval is well-described as a re-instantiation of the pattern of alpha-band activity seen during encoding.

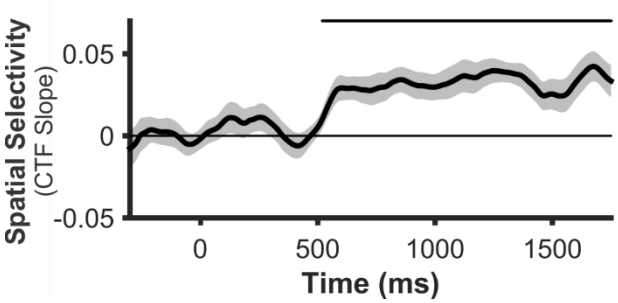

Figure 4-9 Testing whether the multivariate patterns of alpha power at study are reinstated at retrieval
Alpha power tracks the retrieval of spatial information when the IEM was trained on study data and tested on retrieval data, indicating that the pattern of alpha-band activity observed during study is reactivated at retrieval. Shaded error bars represent ±1 bootstrapped s.e.m. Markers on the top of the plot mark the periods of reliable spatial selectivity ($p < .05$).

Discussion

The present work represents a new approach for tracking and understanding the neural mechanisms underlying retrieval of precise feature memories. Over two experiments, we employed a combination of a continuous report task, in which participants learned to associate individual objects with specific spatial locations, with the application of an inverted encoding model to ongoing EEG activity. We demonstrated that IEMs and rhythmic brain activity can be

used track and reconstruct the reinstatement of spatially selective information from long-term memory.

We modeled long-term memory retrieval performance using a mixture modeling approach. This approach is commonplace in the field of visual working memory (Wilken & Ma, 2004; Zhang & Luck, 2008) but has only recently begun gaining traction in the field of long-term memory (Brady et al. 2013). This approach is able to disentangle improvements in mnemonic precision and the probability that memories are retrieved (Fan & Turk-Browne, 2013; Harlow & Yonelinas, 2014; Sutterer & Awh, 2016). Initial studies have found that these parameters are reflected by distinct neural signals (Murray et al., 2015; Richter et al., 2016), providing further evidence that separately modeling mnemonic precision and probability of retrieval is a meaningful distinction. Our results demonstrate that both the probability of retrieving long-term memories and the precision with which those memories are retrieved continue to improve with feedback over many repetitions. We propose that this more sensitive approach of assessing memory accuracy will continue be a successful direction for the field of long-term memory.

A recent model put forth by Watrous and colleagues (2014), the spectro-contextual encoding and retrieval theory (SCERT), asserts that both the frequencies supporting cognitive operations at encoding and retrieval and the specific patterns of activity within those frequencies should overlap. In line with this prediction, we observed considerable overlap in the frequencies carrying spatial memory representations between encoding and retrieval. Furthermore, we found that the multivariate patterns of alpha-band activity reinstated during retrieval are strikingly similar to those patterns observed during the initial encoding of locations. These observations provide new evidence that encoding-retrieval oscillatory similarity extends to the representation of precise feature representations at the population level, supporting the idea that oscillatory

brain activity plays a critical role in memory formation and reinstatement. However, it is worth noting that a broader range of frequencies tracked to-be-remembered locations at encoding than at retrieval, suggesting that not all spatially sensitive frequencies engaged during stimulus presentation are later reinstated.

Indeed, a novel aspect of our work was the ability to search for frequencies that code for precise spatial memories. We demonstrated robust and sustained spatial selectivity during long-term memory retrieval, primarily in the alpha band (8–12 Hz). These results are similar to what has been observed in the field of visual working memory. However this stands in contrast with other findings that suggest the role of theta and beta activity in long-term memory. In particular, past studies have found that theta activity (4–7 Hz) plays a key role in episodic memory (Hsieh & Ranganath, 2014; Nyhus & Curran, 2010) and in the hippocampus during spatial navigation (Bohbot et al., 2017; Watrous et al., 2011). However, it is consistent with some work that suggests a role for alpha in memory and memory guided attention (Stokes et al., 2012; Waldhauser et al., 2016; Waldhauser, Johansson, & Hanslmayr, 2012). It is possible that our scalp EEG signal was relatively insensitive to theta signals prominent in the hippocampus (Hsieh & Ranganath, 2014). Future work from modalities that more directly index hippocampal activity (i.e., MEG/ECOG) might provide some insight into the role of theta in precise spatial memory reinstatement.

Another promising application of the approach employed here is the ability to compare the time course with which fine-grained and courser memory representations emerge. For example, spatially selective alpha activity emerged considerably later than some prior observations of hemifield-selective activity. While hemifield selective activity has been observed within 200 ms of the onset of a retrieval cue (Waldhauser et al., 2016), our time by frequency

analysis revealed no evidence of activity related to the specific retrieved location, in any frequency band, until at least 410 ms after the retrieval cue. One possible explanation for this latency difference is that hemisphere reactivation and retrieval of fine-grained spatial representations rely on different processes. For instance, Gratton et al. (1997), suggest that the hemisphere bias they observe may be more structural, resulting from the formation of a stronger trace in the hemisphere contralateral to the hemifield in which the stimulus was presented; while in the present study, the relatively slower onset of alpha CTFs implies a more effortful retrieval of precise spatial information. Another potential explanation is that context reinstatement during an object recognition task could occur more rapidly than object-cued retrieval of spatially selective information. Further work is needed to explore this difference in latency between hemifield effects and the reactivation of the fine-grained alpha topography that tracks specific locations.

Prominent models have argued that spatial-temporal context is the backbone of episodic memory (Ekstrom & Ranganath, 2017; O'Keefe & Nadel, 1978) serving as an index for the retrieval of specific past experiences. Thus, a method that allows temporally-resolved tracking of spatial retrieval from LTM may provide a powerful tool for understanding human memory. Here, we present such a method, and show that it tracks both the accuracy and latency of memory-guided behavior. Moreover, we provide new evidence confirming a clear prediction of reinstatement models of LTM retrieval. The format of oscillatory activity during encoding into LTM is recapitulated during the subsequent retrieval of those memories.

CHAPTER 5. GENERAL CONCLUSIONS

A hallmark of episodic memory is the phenomenon of mentally re-experiencing the details of past events. Contemporary models of memory posit that maintenance of recently encoded information (WM) and retrieval of information encountered in the distant past (LTM), both rely on active representations of memoranda via persistent patterns of neural activity. Indeed, embedded process models (Cowan, 1995; McElree, 2006; Oberauer, 2002) characterize WM as an activated portion of LTM, suggesting that WM and LTM are part of the same memory system and that active memory representations across delays should rely on the same patterns of activity. Despite growing support for this view, numerous details about the degree of overlap between these two states of memory remain untested. In particular, an open question is whether the most basic features of visual memory representations are supported by the same patterns of rhythmic neural activity across delays. Here, I leveraged precise tests of feature-specific memory by combining a continuous space with the application of a spatial inverted encoding model, to examine three questions about how the brain supports feature-specific, online memory representations. In Chapter 2, I tested whether more than one active feature-specific representation can be held in working memory at once. To this end, I examined spatial representations encoded by delay-period alpha-band (8–12 Hz) activity in a WM task where observers remembered either one or two spatial locations. In line with fMRI (Emrich et al., 2013; Larocque et al., 2017; Sprague et al., 2014, 2016) and unit recordings in non-human primates (Buschman et al., 2011), we observed a decline in the spatial selectivity of alpha-band activity as memory load increased. Critically, this reduction in selectivity could not be explained by models in which only a single item was actively maintained. Simulations revealed that the selectivity of spatial representations in the two-item condition was greater than expected if only

one item was represented at once. Thus, our findings provide clear evidence that multiple items can be represented concurrently by persistent delay-period activity. In Chapter 3, instead of tracking the neural representations directly, I focused on the downstream consequences of activating a memory representation via memory retrieval. Specifically, we assessed whether retrieval practice improved the probability of successful memory retrieval or the quality of the retrieved representation. To answer this question, we assessed the effects of retrieval practice on memory for specific colors. Performance in a subsequent delayed recall test revealed a robust retrieval practice effect. Observers recalled a significantly higher proportion of items that they had previously retrieved relative to items that were untested or that they had restudied. Interestingly, retrieval practice did not elicit any improvement in the precision of the retrieved memories. The same empirical pattern was also observed on a test taken more than a day later. Thus, retrieval practice increases the probability of successful memory retrieval but does not improve memory quality. Finally, in Chapter 4 we investigated whether multivariate patterns of alpha-band power can track memory retrieval of specific spatial locations from LTM. Across two experiments, we established that an inverted encoding model applied to multivariate alpha topography can track retrieval of precise spatial memories. In line with reinstatement accounts of memory, we demonstrate that the pattern of multivariate alpha activity at study is similar to the pattern observed during retrieval. Finally, we observed that these encoding models predict memory retrieval behavior, including the accuracy and latency of recall. Taken together these findings highlight the broad potential for using encoding models to characterize online memory representations for both WM and LTM. Before concluding, I will leave off with a few open questions that are interesting avenues for future research.

Alpha-band activity as a cognitive domain general mechanism of selective spatial attention

Over the course of several studies, we observed that alpha-band oscillations track the precise locus and time-course of the deployment of spatial attention (Foster, Sutterer, Serences, Vogel, et al., 2017), representation in spatial working memory (Foster, Bsales, Jaffe, & Awh, 2017; Foster et al., 2016), and retrieval of long-term memories. Taken together, it is tempting to conclude that alpha-band activity is, or reflects, a cognitive domain general mechanism for selective spatial attention. This interpretation is consistent with a core assumption of embedded process models that maintenance of active memory representations is coordinated via the interaction of attention and long-term memory (Cowan, 1995). This idea is also in line with empirical work suggesting that covert spatial attention facilitates the rehearsal of spatial memories (Awh & Jonides, 2001; Awh, Jonides, & Reuter-lorenz, 1998; Awh et al., 2006). A critical finding in this line of work is that interrupting the delay period of a spatial memory task with distracting spatial information, but not a spatially neutral stimulus, impairs spatial WM performance (Awh et al., 1998), suggesting that spatial attention plays a functional role in spatial memory maintenance. Recent work has shown that similar to memory performance, spatially selective alpha-band activity is disrupted when participants must complete an interleaved spatial attention task during the delay period (van Moorselaar et al., 2017). Furthermore, these disruptions were also accompanied by a decrease in spatial WM performance. It was also possible to train the IEM on attention data and test the model on memory locations. Taken together these results suggest that the patterns of alpha-band power supporting memory and deployment of covert attention are at least somewhat overlapping. Despite initial evidence that these signals overlap, the exact relationship between spatial attention and spatial memory is complex and is an active area of debate. Indeed there is a growing literature dedicated to

assessing the behavioral and neuronal overlap between externally guided and memory guided attention (Hutchinson & Turk-Browne, 2012; Myers, Chekroud, Stokes, & Nobre, 2017; Myers, Walther, Wallis, Stokes, & Nobre, 2015; Rosen, Stern, Michalka, Devaney, & Somers, 2015). Of particular interest are recent fMRI observations that stimulus and memory guided attention rely on a combination of overlapping and separate brain regions (Hutchinson et al., 2014). Thus, further work is needed to determine the extent to which alpha-band oscillations that support spatial attention and spatial memory are driven by the same patterns of activity and brain regions.

Establishing a functional role for alpha oscillations

Another open question is whether alpha-band oscillations actively facilitate ongoing memory maintenance or if these patterns of alpha topography are epiphenomenal, reflecting but not playing a causal role in the maintenance of spatial representations. For instance, the first published report of human EEG (Berger, 1929) described elevated 10 Hz oscillations over posterior electrode sites that increased in magnitude when the eyes were closed and decreased when they were opened. Repeated observations of this effect lead to the hypothesis that alpha-band activity may reflect a cortical idling system that serves to keep populations of neurons ready to function, but plays no crucial role in representing information (Pfurtscheller, Stancák, & Neuper, 1996). Going against this hypothesis are findings that the phase of ongoing alpha-band oscillations before visual stimulation predict perceptual outcomes (Mathewson, Gratton, Fabiani, Beck, & Ro, 2009) and can be guided by top-down control to achieve the optimal phase before stimulus onset (Samaha, Bauer, Cimaroli, & Postle, 2015; Wutz, Melcher, & Samaha, 2018). Moreover, transcranial magnetic stimulation (TMS) of parietal cortex can bias the power of ongoing alpha-band oscillations, and these changes in alpha-band power are related to modulation in working memory and attention (Emrich, Johnson, Sutterer, & Postle, 2016;

Massihullah, Slagter, Tononi, & Postle, 2009) performance. However, given that alpha-band oscillations are likely capable of supporting multiple cognitive functions (J. S. Johnson, Sutterer, Acheson, Lewis-Peacock, & Postle, 2011; Palva & Palva, 2007), further work with brain stimulation and unit recordings will be necessary to elucidate the precise role alpha band oscillations play in the maintenance of online spatial representations.

Conclusion

Completing almost any task in everyday life requires us to rapidly maintain, store, and retrieve memory representations. Central questions for the field of memory research are what neuronal representations contribute to active memory representations, when long-term memories are retrieved into an active state, and how these representations evolve over long delays. To address these questions, I developed and applied an approach in which we use an inverted encoding model (Brouwer & Heeger, 2009; Sprague et al., 2015) and multivariate patterns of electroencephalography activity to track the maintenance of online spatial memory representations. Our team first developed an approach showing that alpha-band (8–12 Hz) oscillations track a single position maintained in working memory (Foster et al., 2016) and that these oscillations also enable time resolved tracking of the deployment of the covert locus of attention (Foster, Sutterer, Serences, Vogel, et al., 2017). I then leveraged a combination of this approach and careful simulations to assess whether multiple representations can be held in mind concurrently. Consistent with both embedded process models of memory that allow for multiple active representations (Cowan, 1995; Oberauer, 2002) and neuroscience accounts of capacity limits in visual WM (Bays, 2014; Franconeri et al., 2013), we found that observers can maintain simultaneous active representations of two locations in WM. We also found that this same approach, and indeed the same patterns of activity observed during memory encoding, can be

used to track the moment-by-moment retrieval of remembered locations from long-term memory. Moreover, we found that the spatial selectivity of alpha-band activity tracked memory performance as learning progressed as well as the latency of behavioral recall. Taken together, this pattern of results suggests that multivariate analysis of alpha-band activity provides a powerful platform for tracking both the initial encoding and subsequent retrieval of online spatial representations.

References

Atkinson, R. C., & Shiffrin, R. M. (1968). Human Memory: A Proposed System and its Control Processes. *Psychology of Learning and Motivation*, *2*, 89–195.

Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*(3), 119–126. http://doi.org/10.1016/S1364-6613(00)01593-X

Awh, E., Jonides, J., & Reuter-lorenz, P. A. (1998). Rehearsal in Spatial Working Memory, *24*(3), 780–790.

Awh, E., Vogel, E. K., & Oh, S. H. (2006). Interactions between attention and working memory. *Neuroscience*, *139*(1), 201–208. http://doi.org/10.1016/j.neuroscience.2005.08.023

Baddeley, A. D., & Warrington, E. K. (1970). Amnesia and the distinction between long- and short-term memory. *Journal of Verbal Learning and Verbal Behavior*, *9*(2), 176–189. http://doi.org/10.1016/S0022-5371(70)80048-2

Bays, P. M. (2014). Noise in Neural Populations Accounts for Errors in Working Memory. *Journal of Neuroscience*, *34*(10), 3632–3645. http://doi.org/10.1523/JNEUROSCI.3204-13.2014

Bays, P. M. (2018). Reassessing the Evidence for Capacity Limits in Neural Signals Related to Working Memory. *Cerebral Cortex*, (April), 1432–1438. http://doi.org/10.1093/cercor/bhx351

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), 7–7. http://doi.org/10.1167/9.10.7

Berger, H. (1929). Uber das Elektrenkephalogramm des Menschen. Zweite Mitteilung. *Arch. Psychiatr. Nervenkr.*, *87*, 528–570.

Bohbot, V. D., Copara, M. S., Gotman, J., & Ekstrom, A. D. (2017). Low-frequency theta oscillations in the human hippocampus during real-world and virtual navigation. *Nature Communications*, *8*, 14415. http://doi.org/10.1038/ncomms14415

Bosch, S. E., Jehee, J. F. M., Fernandez, G., & Doeller, C. F. (2014). Reinstatement of Associative Memories in Early Visual Cortex Is Signaled by the Hippocampus. *Journal of Neuroscience*, *34*(22), 7493–7500. http://doi.org/10.1523/JNEUROSCI.0805-14.2014

Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2013). Visual Long-Term Memory Has the Same Limit on Fidelity as Visual Working Memory. *Psychological Science*, *24*(6), 981–990. http://doi.org/10.1177/0956797612465439

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

http://doi.org/10.1163/156856897X00357

Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *The Journal of Neuroscience*, *29*(44), 13992–4003. http://doi.org/10.1523/JNEUROSCI.3577-09.2009

Busch, N. A., & Van Rullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(37), 16048–53. http://doi.org/10.1073/pnas.1004801107

Buschman, T. J., Siegel, M., Roy, J. E., & Miller, E. K. (2011). Ethics 2011000789. *Proceedings of the National Academy of Sciences*, *108*(27), 1–4. http://doi.org/10.1073/pnas.1104666108/-/DCSupplemental.www.pnas.org/cgi/doi/10.1073/pnas.1104666108

Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–15. http://doi.org/10.1016/j.tics.2010.09.001

Carpenter, S. K. (2009). Cue Strength as a Moderator of the Testing Effect: The Benefits of Elaborative Retrieval. *Journal of Experimental Psychology: Learning Memory and Cognition*, *35*(6), 1563–1569. http://doi.org/10.1037/a0017021

Carpenter, S. K., & Kelly, J. W. (2012). Tests enhance retention and transfer of spatial learning. *Psychonomic Bulletin & Review*, *19*(3), 443–8. http://doi.org/10.3758/s13423-012-0221-2

Carpenter, S. K., & Pashler, H. (2007). Testing beyond words: using tests to enhance visuospatial map learning. *Psychonomic Bulletin & Review*, *14*(3), 474–8.

Carrier, M., & Pashler, H. (1992). The influence of retrieval on retention. *Memory & Cognition*, *20*(6), 633–42.

Christophel, T. B., Iamshchinina, P., Yan, C., Allefeld, C., & Haynes, J.-D. (2018). Cortical specialization for attended versus unattended working memory. *Nature Neuroscience*. http://doi.org/10.1038/s41593-018-0094-4

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409. http://doi.org/10.1016/j.neuropsychologia.2011.01.029

Chun, M. M., & Turk-Browne, N. B. (2007). Interactions between attention and memory. *Current Opinion in Neurobiology*, *17*(2), 177–184. http://doi.org/10.1016/j.conb.2007.03.005

Cohen, M. X. (2014). *Analyzing Neural Time Series Data: Theory and Practice*. *MIT Press*. http://doi.org/10.1017/CBO9781107415324.004

Cowan, N. (1995). *Attention and memory: An integrated framework. Oxford Psychology Series*

(Vol. 26).

Cowan, N. (1999). An embedded-processes model of working memory. *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control.* http://doi.org/10.1017/S0140525X01003922

Craik Fergus, I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, *11*, 671–684.

Danker, J. F., & Anderson, J. R. (2010). The ghosts of brain states past: remembering reactivates the brain regions engaged during encoding. *Psychological Bulletin*, *136*(1), 87–102. http://doi.org/10.1037/a0017937

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

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

Ekstrom, A. D., & Ranganath, C. (2017). Space, time, and episodic memory: The hippocampus is all over the cognitive map. *Hippocampus*, (June), 1–8. http://doi.org/10.1002/hipo.22750

Emrich, S. M., Johnson, J. S., Sutterer, D. W., & Postle, B. R. (2016). Comparing the effects of 10-Hz repetetive TMS on tasks of visual STM and attention. *Journal of Cognitive Neuroscience*, *29*(2), 286–297. http://doi.org/10.1162/jocn

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. *The Journal of Neuroscience*, *33*(15), 6516–23. http://doi.org/10.1523/JNEUROSCI.5732-12.2013

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*(4), 893–905. http://doi.org/10.1016/j.neuron.2015.07.013

Fan, J. E., & Turk-Browne, N. B. (2013). Internal attention to features in visual short-term memory guides object learning. *Cognition*, *129*(2), 292–308. http://doi.org/10.1016/j.cognition.2013.06.007

Fell, J., & Axmacher, N. (2011). The role of phase synchronization in memory processes. *Nature Reviews. Neuroscience*, *12*(2), 105–18. http://doi.org/10.1038/nrn2979

Fiebelkorn, I. C., Saalmann, Y. B., & Kastner, S. (2013). Rhythmic sampling within and between objects despite sustained attention at a cued location. *Current Biology*, *23*(24), 2553–2558. http://doi.org/10.1016/j.cub.2013.10.063

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. http://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*, 168–177. http://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*(7), 929–941. http://doi.org/10.1177/0956797617699167

Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: Competitive content maps for attention and memory. *Trends in Cognitive Sciences*, *17*(3), 134–141. http://doi.org/10.1016/j.tics.2013.01.010

Fukuda, K., Awh, E., & Vogel, E. K. (2010). Discrete capacity limits in visual working memory. *Current Opinion in Neurobiology*, *20*(2), 177–182. http://doi.org/10.1016/j.conb.2010.03.005

Fukuda, K., Mance, I., & Vogel, E. K. (2015). Power Modulation and Event-Related Slow Wave Provide Dissociable Correlates of Visual Working Memory. *Journal of Neuroscience*, *35*(41), 14009–14016. http://doi.org/10.1523/JNEUROSCI.5003-14.2015

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*(2), 331–349. http://doi.org/10.1016/j.neuron.2012.12.039

Funahashi, S., Chafee, M. V., & Goldman-Rakic, P. S. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, *365*(6448), 753–756. http://doi.org/10.1038/365753a0

Fuster, J. M., & Alexander, G. E. (1971). Neuron Activity Related to Short-Term Memory. *Science*, *173*(3997), 652–654.

Grimault, S., Robitaille, N., Grova, C., Lina, J.-M., Dubarry, A.-S., Jolicœur, P., & Jolicoeur, P. (2009). Oscillatory activity in parietal and dorsolateral prefrontal cortex during retention in visual short-term memory: Additive effects of spatial attention and memory load. *Human Brain Mapping*, *30*(10), 3378–92. http://doi.org/10.1002/hbm.20759

Harlow, I. M., & Donaldson, D. I. (2013). Source accuracy data reveal the thresholded nature of human episodic memory. *Psychonomic Bulletin & Review*, *20*(2), 318–25. http://doi.org/10.3758/s13423-012-0340-9

Harlow, I. M., & Yonelinas, A. P. (2014). Distinguishing between the success and precision of recollection. *Memory*, *8211*(January 2015), 1–14. http://doi.org/10.1080/09658211.2014.988162

Harrison, S. a, & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–5. http://doi.org/10.1038/nature07832

Hindy, N. C., Ng, F. Y., & Turk-Browne, N. B. (2016). Linking pattern completion in the hippocampus to predictive coding in visual cortex. *Nature Neuroscience*, *19*(5), 665–667. http://doi.org/10.1038/nn.4284

Hintzman, D. L. (2011). Research Strategy in the Study of Memory: Fads, Fallacies, and the Search for the "Coordinates of Truth." *Perspectives on Psychological Science*, *6*(3), 253–271. http://doi.org/10.1177/1745691611406924

Hsieh, L., & Ranganath, C. (2014). Frontal midline theta oscillations during working memory maintenance and episodic encoding and retrieval. *NeuroImage*, *85*, 721–729. http://doi.org/10.1016/j.neuroimage.2013.08.003

Hutchinson, J. B., & Turk-Browne, N. B. (2012). Memory-guided attention : control from multiple memory systems. *Trends in Cognitive Sciences*, *16*(12), 576–579. http://doi.org/10.1016/j.tics.2012.10.003

Hutchinson, J. B., Uncapher, M. R., Weiner, K. S., Bressler, D. W., Silver, M. A., Preston, A. R., & Wagner, A. D. (2014). Functional heterogeneity in posterior parietal cortex across attention and episodic memory retrieval. *Cerebral Cortex*, *24*(1), 49–66. http://doi.org/10.1093/cercor/bhs278

Jafarpour, A., Fuentemilla, L., Horner, A. J., Penny, W., & Duzel, E. (2014). Replay of very early encoding representations during recollection. *The Journal of Neuroscience*, *34*(1), 242–8. http://doi.org/10.1523/JNEUROSCI.1865-13.2014

James, W. (1890). The principles of psychology. *New York Holt*, *1*, 697. http://doi.org/10.1037/10538-000

Johnson, J. D., Price, M. H., & Leiker, E. K. (2015). Episodic retrieval involves early and sustained effects of reactivating information from encoding. *NeuroImage*, *106*, 300–310. http://doi.org/10.1016/j.neuroimage.2014.11.013

Johnson, J. S., Sutterer, D. W., Acheson, D. J., Lewis-Peacock, J. A., & Postle, B. R. (2011). Increased alpha-band power during the retention of shapes and shape-location associations in visual short-term memory. *Frontiers in Psychology*, *2*(128), 1–9. http://doi.org/10.3389/fpsyg.2011.00128

Jonides, J., Lewis, R. L., Nee, D. E., Lustig, C. A., Berman, M. G., & Moore, K. S. (2008). The Mind and Brain of Short-Term Memory. *Annual Review of Psychology*, *59*(1), 193–224. http://doi.org/10.1146/annurev.psych.59.103006.093615

Kamiński, J., Sullivan, S., Chung, J. M., Ross, I. B., Mamelak, A. N., & Rutishauser, U. (2017). Persistently active neurons in human medial frontal and medial temporal lobe support working memory. *Nature Neuroscience*, *20*(4), 590–601. http://doi.org/10.1038/nn.4509

Kornblith, S., Quiroga, R. Q., Koch, C., Fried, I., & Mormann, F. (2017). Persistent single neuron activity during working memory in the human medial temporal lobe. *Current Biology, This issue*, 1026–1032. http://doi.org/10.1016/j.cub.2017.02.013

Kuhl, B. A., Rissman, J., Chun, M. M., & Wagner, A. D. (2011). Fidelity of neural reactivation reveals competition between memories. *Proceedings of the National Academy of Sciences*, *108*(14), 5903–5908. http://doi.org/10.1073/pnas.1016939108

Landau, A. N., & Fries, P. (2012). Attention samples stimuli rhythmically. *Current Biology*, *22*(11), 1000–1004. http://doi.org/10.1016/j.cub.2012.03.054

Larocque, J. J., Lewis-peacock, J. A., & Postle, B. R. (2014). Multiple neural states of representation in short-term memory ? It â€™ s a matter of attention Article type : Received on : Accepted on : Frontiers website link : Citation : Multiple neural states of representation in short-term memory ? It â€™ s a matter of . http://doi.org/10.3389/fnhum.2014.00005

Larocque, J. J., Riggall, A. C., Emrich, S. M., & Postle, B. R. (2017). Within-Category Decoding of Information in Different Attentional States in Short-Term Memory. *Cerebral Cortex*, *27*(10), 4881–4890. http://doi.org/10.1093/cercor/bhw283

Lehman, M., Smith, M. a, & Karpicke, J. D. (2014). Toward an episodic context account of retrieval-based learning: Dissociating retrieval practice and elaboration. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *40*(4), 1–8. http://doi.org/10.1037/xlm0000012

Lewis-peacock, J. A. (2012). Jarrod A. Lewis-Peacock, *28*(2008), 2008–2011.

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural evidence for a distinction between short-term memory and the focus of attention. *Journal of Cognitive Neuroscience*, *24*(1), 61–79. http://doi.org/10.1162/jocn_a_00140

Lins, O. G., Picton, T. W., Berg, P., & Scherg, M. (1993). Ocular artifacts in recording EEGs and Event-Related potentials II: Source dipoles and source components. *Brain Topography*, *6*(1), 65–78. http://doi.org/10.1007/BF01234128

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–81. http://doi.org/10.1038/36846

Luck, S. J., & Vogel, E. K. (2014). Visual Working Memory Capacity: From Psychophysics and Neurobiology to individual Differences. *Trends Cogn Sci.*, *17*(8), 391–400. http://doi.org/10.1016/j.tics.2013.06.006.Visual

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–90. http://doi.org/10.1016/j.jneumeth.2007.03.024

Massihullah, H., Slagter, H. A., Tononi, G., & Postle, B. R. (2009). Repetetive transcranial

magnetic stimulation affects behavior by biasing endogenous cortical oscillations. *Frontiers in Integrative Neuroscience*, *3*(14), 1–12. http://doi.org/10.3389/neuro.07

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(9), 2725–32. http://doi.org/10.1523/JNEUROSCI.3963-08.2009

Matsushima, A., & Tanaka, M. (2014). Different Neuronal Computations of Spatial Working Memory for Multiple Locations within versus across Visual Hemifields. *Journal of Neuroscience*, *34*(16), 5621–5626. http://doi.org/10.1523/JNEUROSCI.0295-14.2014

McElree, B. (1998). Attended and Non-Attended States in Working Memory: Accessing Categorized Structures. *Journal of Memory and Language*, *38*(38), 225–252. http://doi.org/10.1006/jmla.1997.2545

McElree, B. (2006). Accessing Recent Events. *Psychology of Learning and Motivation - Advances in Research and Theory*, *46*(6), 155–200. http://doi.org/10.1016/S0079-7421(06)46005-9

McElree, B., & Dosher, B. A. (2001). The focus of attention across space and across time. *Behavioral and Brain Sciences*, *24*(1), 129–130. http://doi.org/10.1017/S0140525X01373922

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. http://doi.org/10.1093/cercor/bhl145

Morton, N. W., Kahana, M. J., Rosenberg, E. A., Baltuch, G. H., Litt, B., Sharan, A. D., … Polyn, S. M. (2013). Category-specific neural oscillations predict recall organization during memory search. *Cerebral Cortex*, *23*(10), 2407–2422. http://doi.org/10.1093/cercor/bhs229

Morton, N. W., & Polyn, S. M. (2017). Beta-band activity represents the recent past during episodic encoding. *NeuroImage*, *147*(December 2016), 692–702. http://doi.org/10.1016/j.neuroimage.2016.12.049

Murray, J. G., Howie, C. A., & Donaldson, D. I. (2015). The neural mechanism underlying recollection is sensitive to the quality of episodic memory: Event related potentials reveal a some-or-none threshold. *NeuroImage*, *120*, 298–308. http://doi.org/10.1016/j.neuroimage.2015.06.069

Myers, N. E., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2017). Benefits of Flexible Prioritization in Working Memory Can Arise Without Costs. *Journal of Experimental Psychology: Human Perception and Performance*. http://doi.org/10.1037/xhp0000449

Myers, N. E., Walther, L., Wallis, G., Stokes, M. G., & Nobre, A. C. (2015). Temporal dynamics

of attention during encoding versus maintenance of working memory: complementary views from event-related potentials and alpha-band oscillations. *Journal of Cognitive Neuroscience*, *27*(3), 492–508. http://doi.org/10.1162/jocn

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, *10*(9), 424–430. http://doi.org/10.1016/j.tics.2006.07.005

Nyberg, L., Habib, R., McIntosh, A. R., & Tulving, E. (2000). Reactivation of encoding-related brain activity during memory retrieval. *Proceedings of the National Academy of Sciences*, *97*, 11120–11124. http://doi.org/10.1073/pnas.97.20.11120

Nyhus, E., & Curran, T. (2010). Functional role of gamma and theta oscillations in episodic memory. *Neuroscience and Biobehavioral Reviews*, *34*(7), 1023–1035. http://doi.org/10.1016/j.neubiorev.2009.12.014

O'Keefe, & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. *Oxford University Press.* http://doi.org/10.1017/CBO9781107415324.004

Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 411–421. http://doi.org/10.1037//0278-7393.28.3.411

Oberauer, K., & Hein, L. (2012). Attention to Information in Working Memory. *Current Directions in Psychological Science*, *21*(3), 164–169. http://doi.org/10.1177/0963721412444727

Palva, S., & Palva, J. M. (2007). New vistas for alpha-frequency band oscillations. *Trends in Neurosciences*, *30*(4), 150–8. http://doi.org/10.1016/j.tins.2007.02.001

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442. http://doi.org/10.1163/156856897X00366

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

Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, *310*(5756), 1963–1966. http://doi.org/10.1126/science.1117645

Pyc, M. a., & Rawson, K. a. (2009). Testing the retrieval effort hypothesis: Does greater difficulty correctly recalling information lead to higher levels of memory? *Journal of Memory and Language*, *60*(4), 437–447. http://doi.org/10.1016/j.jml.2009.01.004

Rajsic, J., & Wilson, D. E. (2012). Remembering where : Estimated memory for visual objects is better when retrieving location with colour. *Visual Cognition*, *6285*(April 2013), 37–41. http://doi.org/10.1080/13506285.2012.726477

Richter, F. R., Cooper, R. A., Bays, P. M., & Simons, J. S. (2016). Distinct neural mechanisms underlie the success, precision, and vividness of episodic memory. *eLife*, *5*, 1–18. http://doi.org/10.7554/eLife.18260

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

Ritchey, M., Wing, E. A., LaBar, K. S., & Cabeza, R. (2013). Neural Similarity Between Encoding and Retrieval is Related to Memory Via Hippocampal Interactions. *Cerebral Cortex*, *23*(12), 2818–2828. http://doi.org/10.1093/cercor/bhs258

Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: taking memory tests improves long-term retention. *Psychological Science*, *17*(3), 249–55. http://doi.org/10.1111/j.1467-9280.2006.01693.x

Rohrer, D., Taylor, K., & Sholar, B. (2010). Tests enhance the transfer of learning. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *36*(1), 233–9. http://doi.org/10.1037/a0017678

Rose, N. S., Larocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. *Science*, *354*(6316), 1136–1139.

Rosen, M. L., Stern, C. E., Michalka, S. W., Devaney, K. J., & Somers, D. C. (2015). Influences of Long-Term Memory-Guided Attention and Stimulus-Guided Attention on Visuospatial Representations within Human Intraparietal Sulcus. *Journal of Neuroscience*, *35*(32), 11358–11363. http://doi.org/10.1523/JNEUROSCI.1055-15.2015

Samaha, J., Bauer, P., Cimaroli, S., & Postle, B. R. (2015). Top-down control of the phase of alpha-band oscillations as a mechanism for temporal prediction. *Proceedings of the National Academy of Sciences*, (experiment 1), 201503686. http://doi.org/10.1073/pnas.1503686112

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. http://doi.org/10.1162/jocn

Saproo, S., & Serences, J. T. (2014). Attention Improves Transfer of Motion Information between V1 and MT. *Journal of Neuroscience*, *34*(10), 3586–3596. http://doi.org/10.1523/JNEUROSCI.3484-13.2014

Sauseng, P., Klimesch, W., Heise, K. F., Gruber, W. R., Holz, E., Karim, A. A., … Hummel, F.

C. (2009). Brain Oscillatory Substrates of Visual Short-Term Memory Capacity. *Current Biology*, *19*(21), 1846–1852. http://doi.org/10.1016/j.cub.2009.08.062

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. *European Journal of Neuroscience*, *22*(11), 2917–2926. http://doi.org/10.1111/j.1460-9568.2005.04482.x

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions: Memory and memories-looking back and looking forward. *Journal of Neurology, Neurosurgery and Psychiatry*, *20*, 11–21. http://doi.org/10.1136/jnnp-2015-311092

Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci*, *20*(2), 207–214. http://doi.org/10.1111/j.1467-9280.2009.02276.x

Shallice, T., & Warrington, E. K. (1970). Independant functionning of verbal memory stores : a neuropsychological study. *Quarterly Journal of Experimental Psychology*, *22*, 261–273.

Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience*, *13*(2), 121–134. http://doi.org/10.1038/nrn3137

Singer, W. (1999). Neuronal Synchrony : A Versatile Code for the Definition of Relations ? Most of our knowledge about the functional organization, *24*, 49–65.

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), 2174–2180. http://doi.org/10.1016/j.cub.2014.07.066

Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, *91*(3), 694–707. http://doi.org/10.1016/j.neuron.2016.07.006

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*, 1–12. http://doi.org/10.1016/j.tics.2015.02.005

Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*(12), 1879–87. http://doi.org/10.1038/nn.3574

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. http://doi.org/10.1016/j.tics.2013.12.001

Stokes, M. G. (2015). "Activity-silent" working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405.

http://doi.org/10.1016/j.tics.2015.05.004

Stokes, M. G., Atherton, K., Patai, E. Z., & Nobre, A. C. (2012). Long-term memory prepares neural activity for perception. *Proceedings of the National Academy of Sciences*, *2011*(3), 188–197. http://doi.org/10.1073/pnas.1108555108/-/DCSupplemental.www.pnas.org/cgi/doi/10.1073/pnas.1108555108

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

Sutterer, D. W., & Awh, E. (2016). Retrieval practice enhances the accessibility but not the quality of memory. *Psychonomic Bulletin & Review*, *23*(3), 831–841. http://doi.org/10.3758/s13423-015-0937-x

Thut, G. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *Journal of Neuroscience*, *26*(37), 9494–9502. http://doi.org/10.1523/JNEUROSCI.0875-06.2006

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

Treder, M. S., Bahramisharif, A., Schmidt, N. M., van Gerven, M. A. J., & Blankertz, B. (2011). Brain-computer interfacing using modulations of alpha activity induced by covert shifts of attention. *Journal of NeuroEngineering and Rehabilitation*, *8*(1). http://doi.org/10.1186/1743-0003-8-24

van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial Comparison of Working Memory Models, *121*(1), 124–149. http://doi.org/10.1037/a0035234

van den Berg, R., Shin, H., Chou, W.-C. W.-C., George, R., & Ma, W. J. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(22), 8780–8785. http://doi.org/10.1073/pnas.1117465109

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), 900–5. http://doi.org/10.1073/pnas.0908821107

van Moorselaar, D., Foster, J. J., Sutterer, D. W., Theeuwes, J., Olivers, C. N. L., & Awh, E. (2017). Spatially Selective Alpha Oscillations Reveal Moment-by-Moment Trade-offs between Working Memory and Attention. *Journal of Cognitive Neuroscience*, 1–11. http://doi.org/10.1162/jocn

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*(6984), 748–51.

http://doi.org/10.1038/nature02447

Wagner, A. D., Shannon, B. J., Kahn, I., & Buckner, R. L. (2005). Parietal lobe contributions to episodic memory retrieval. *Trends in Cognitive Sciences*, *9*(9), 445–53. http://doi.org/10.1016/j.tics.2005.07.001

Waldhauser, G. T., Braun, V., & Hanslmayr, S. (2016). Episodic Memory Retrieval Functionally Relies on Very Rapid Reactivation of Sensory Information. *The Journal of Neuroscience*, *36*(1), 251–260. http://doi.org/10.1523/JNEUROSCI.2101-15.2016

Waldhauser, G. T., Johansson, M., & Hanslmayr, S. (2012). A/B Oscillations Indicate Inhibition of Interfering Visual Memories. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(6), 1953–61. http://doi.org/10.1523/JNEUROSCI.4201-11.2012

Watrous, A. J., & Ekstrom, A. D. (2014). The Spectro-Contextual Encoding and Retrieval Theory of Episodic Memory. *Frontiers in Human Neuroscience*, *8*. http://doi.org/10.3389/fnhum.2014.00075

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. http://doi.org/10.1016/j.conb.2014.07.024

Watrous, A. J., Fried, I., & Ekstrom, A. D. (2011). Behavioral correlates of human hippocampal delta and theta oscillations during navigation. *Journal of Neurophysiology*, *105*, 1747–1755. http://doi.org/10.1152/jn.00921.2010.

Wheeler, M. A., & Roediger, H. L. (1992). Disparate Effects of Repeated Testing: Reconciling Ballard's (1913) and Bartlett's (1932) Results. *Psychological Science*, *3*(4), 240–245. http://doi.org/10.1111/j.1467-9280.1992.tb00036.x

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), 1120–35. http://doi.org/10:1167/4.12.11

Wimber, M., Maaß, A., Staudigl, T., Richardson-Klavehn, A., & Hanslmayr, S. (2012). Rapid memory reactivation revealed by oscillatory entrainment. *Current Biology*, *22*(16), 1482–1486. http://doi.org/10.1016/j.cub.2012.05.054

Wolff, M. J., Jochim, J., Akyurek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working memory guided behaviour 2. *Nature Neurosci.*, (April), 1–35. http://doi.org/10.1038/nn.4546

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *20*(6), RC63. http://doi.org/Rc63

Wutz, A., Melcher, D., & Samaha, J. (2018). Frequency modulation of neural oscillations according to visual task demands. *Proceedings of the National Academy of Sciences*, 201713318. http://doi.org/10.1073/pnas.1713318115

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–5. http://doi.org/10.1038/nature06860