

THE UNIVERSITY OF CHICAGO

DISTINCT NEURAL SIGNALS REVEAL A DISSOCIATION BETWEEN STORAGE OF
INDIVIDUATED ITEMS AND SPATIAL ATTENTION

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

GISELLA KARELIA DIAZ

CHICAGO, ILLINOIS

AUGUST 2022

TABLE OF CONTENTS

LIST OF FIGURES.....	iii
LIST OF TABLES	iv
ACKNOWLEDGEMENTS	v
ABSTRACT	vi
INTRODUCTION.....	1
PART I. A generalizable human signature of spatial attention	4
CHAPTER 1. ALPHA TOPOGRAPHY TRACKS ATTENDED LOCATION ACROSS INDIVIDUALS.	5
<i>Materials and Methods</i>	6
<i>Results</i>	20
<i>Discussion</i>	28
PART II. Tracking individuated items in mind	31
CHAPTER 2. DISTINCT SIGNALS TRACK ITEMS AND LOCATIONS IN WORKING MEMORY.	32
<i>Materials and Methods</i>	35
<i>Experiment 2-1 Results</i>	44
<i>Experiment 2-2 Results</i>	47
<i>Discussion</i>	51
CHAPTER 3. EVIDENCE FOR AN ITEM-BASED POINTER SYSTEM.	55
<i>Materials and Methods</i>	57
<i>Results</i>	67
<i>Discussion</i>	78
GENERAL DISCUSSION.....	81
REFERENCES.....	85

LIST OF FIGURES

Figure 1-1. Redrawn figures of tasks.	7
Figure 1-2. Redrawn figure of Bae and Luck (2018) task.....	14
Figure 1-3. Inverted encoding model (IEM) approach.....	17
Figure 1-4. Alpha power channel tuning function (CTF) across participants.....	21
Figure 1-5. Channel responses for each stimulus location based on location channel.....	22
Figure 1-6. CTF slope over time.	23
Figure 1-7. CTFs and their slope over time across projects and sites.	24
Figure 1-8. Channel responses across projects and sites.....	25
Figure 1-9. CTFs over time across projects and sites.....	26
Figure 1-10. CTF and CTF slope across research groups.	27
Figure 2-1. Change detection tasks and perceptual grouping displays.	42
Figure 2-2. Negative slow wave and alpha power over time in Experiment 2-1.	46
Figure 2-3. Negative slow wave and alpha power over time in Experiment 2-2.	50
Figure 3-1. Sequential change detection task.	61
Figure 3-2. mvLoad approach.	65
Figure 3-3. Alpha power over time.	69
Figure 3-4. Classification accuracy over time.....	71
Figure 3-5. Classifier evidence over time when novel positions were occupied.	75
Figure 3-6. Classifier evidence over time when the same positions were occupied.	77

LIST OF TABLES

Table 1-1. Experiment design information for the eight tasks.	8
Table 2-1. Repeated measures ANOVAs for neural data in Experiment 2-2.....	49

ACKNOWLEDGEMENTS

I would like to extend my sincere gratitude to my advisors, Ed Awh and Ed Vogel, for their invaluable support and guidance throughout my PhD. I would also like to thank Wilma Bainbridge and Howard Nusbaum for their feedback and insightful suggestions on this dissertation. I am also grateful for the continuous support and friendship of my labmates. Finally, I would like to acknowledge my family, friends, and partner for their unwavering love and encouragement.

ABSTRACT

At any given moment, there exists an overwhelming amount of information in our surroundings. Yet, only a subset of that information is relevant or critical to our current goals. *Attention* allows us to prioritize the selection of relevant information, thereby making that information accessible to higher order cognitive processes, like reasoning and decision-making. Meanwhile, *working memory* allows us to maintain and manipulate the selected information in a “mental workspace.” In everyday life, attention and working memory are inextricably linked. For instance, mentally solving an arithmetic problem requires that we select, maintain, and manipulate its key elements. In this dissertation, I demonstrate that active maintenance of information relies on several subcomponents whose operations are tracked by distinct neural signals. In Chapter 1, I demonstrate that there exists a generalizable human signature of attention that tracks the selection and maintenance of spatial positions in mind. In Chapter 2, I provide evidence that neural activity tracking items based on their spatial positions can be dissociated from activity tracking the number of individuated items. Finally, Chapter 3 proposes that signals tracking the number of items in mind are instead indexing the product of a more fundamental cognitive operation. Namely, I provide evidence that this neural activity represents spatiotemporal “pointers” that enable the continuous tracking of items without necessarily containing information about the items themselves. Together, the current work highlights that neural activity related to attention and working memory reflects parallel but distinct processes that jointly contribute to successful behavior.

INTRODUCTION

Our everyday behaviors are comprised of seemingly simple tasks. In a crowded scene, we can differentiate between words and nonwords and recognize our name as especially meaningful when we hear it. We can reorient ourselves and scan our surroundings for a familiar face. The mundane task of finding a friend in the city center thus appears seamless and almost effortless. Our continuous experience of the world gives an undeniable but deceptive impression of an infallible brain. Introspection allows us to break down a complex behavior, such as searching for a friend, into its constituent parts, but it does little to inform the underlying brain processes. At first glance, we might be tempted to underestimate the complexity of the brain, but it is the complexity of it that allows for our uninterrupted experience of the world, where mistakes happen but are often quickly and efficiently corrected.

Our experience of the world might also falsely suggest that we have instant access to a great deal of information from our environment. At any given moment, however, the human brain is quite limited in its ability to actively hold information in mind. We have access to only a subset of information. *Working memory* is a capacity-limited system that allows us to temporarily maintain information in a readily accessible state so that we can act upon it. It is a “mental workspace” where information can be manipulated. In everyday life, working memory is inextricably linked to *attention*, which allows us to prioritize the selection of currently relevant information. For instance, mentally solving an arithmetic problem requires that we select, maintain, and manipulate its key elements. Attention and working memory thus ensure that critical or relevant information is made accessible to higher order cognitive operations, like reasoning and decision making.

It is therefore consistent that there is a well-documented relationship between working memory and other cognitive abilities. For instance, working memory capacity refers to the amount of information that one can hold in mind at any given moment, and it is a strong predictor of fluid intelligence (Conway et al., 2003). Working memory capacity is also a stable trait that is unlikely to be improved with training (Shipstead et al., 2012). In addition to capacity, individual differences in attentional control also contribute to the relationship between working memory and fluid intelligence (Unsworth et al., 2014). Thus, there is great interest in the field of cognitive neuroscience to understand the mechanisms underlying working memory, as well as its interactions with attentional processes.

The field is greatly informed by neuroscience research on non-human primates. Single-unit recordings show persistent and sustained levels of firing during the delay period of a working memory task across various cortical areas (Fuster & Alexander, 1971; Fuster & Jervey, 1982; Miller et al., 1996). Similar findings of sustained delay activity have been found in humans using functional magnetic resonance imaging (fMRI) (Courtney et al., 1997; Pessoa et al., 2002; Srimal & Curtis, 2008). Advances in neuroimaging techniques have moved the field forward by providing a means to study the information maintained in neural activity. Multivariate analyses have been used to decode simple visual features or spatial positions held in mind from delay period activity (Christophel et al., 2012; Harrison & Tong, 2009; Jerde et al., 2012; Serences et al., 2009). Additionally, inverted encoding models (IEMs) have been used to recover representations of features or spatial positions held in mind also from delay period activity (Brouwer & Heeger, 2009; Ester et al., 2013; Foster et al., 2016; Sprague et al., 2014). The active maintenance of information then has been characterized by sustained increases in neural activity that may also carry information about the content of the representations in mind.

In this dissertation, I focus on this traditional view that working memory relies on the active maintenance of neural representations, although recent views propose that active maintenance is not always necessary for working memory (Lewis-Peacock et al., 2012; Rose et al., 2016; Stokes, 2015).

This dissertation seeks to further characterize the neural correlates of working memory and spatial attention, or the selection of information on the basis of its spatial position. Spatial attention is just one of many ways to bring information into working memory. In the tasks discussed, participants were required to direct their attention to stimuli at various spatial positions, while maintaining eye fixation on the center of the display. Participants remembered the stimuli over brief blank delay intervals, while electroencephalogram (EEG) activity was recorded. Distinct neural signals revealed a dissociation between the storage of individuated items and spatial attention to those items, demonstrating that multiple cognitive processes work together for effective maintenance.

In Chapter 1, I first demonstrate that there exists a generalizable human signature of attention that tracks the selection and maintenance of spatial positions in mind. In Chapter 2, I provide evidence that delay activity tracking items based on their spatial positions can be dissociated from activity tracking the number of individuated items. Finally, Chapter 3 proposes that signals tracking the number of items in mind are instead indexing the product of a more fundamental cognitive operation. Namely, I provide evidence that this neural activity represents spatiotemporal “pointers” that enable the continuous tracking of items without necessarily containing information about the items themselves. Together, the current work highlights that neural activity related to attention and working memory reflects parallel but distinct processes that jointly contribute to successful behavior.

PART I.

A generalizable human signature of spatial attention

CHAPTER 1.

Alpha topography tracks attended location across individuals.

At any given moment, we are presented with far more sensory information than we can process at once. Thus, it is critical for observers to use attention to focus limited processing resources on the information that is most relevant to their current goals. Spatial attention, specifically, plays a key role in enhancing the quality of information at attended locations and increasing the rate at which that information is processed (Carrasco & McElree, 2001; Posner, 1980). Because of this fundamental role in cognition, there is great interest in understanding the neural basis of spatial attention.

Using electroencephalography (EEG), previous work has linked spatial attention with rhythmic brain activity in the alpha frequency band, ~8 – 12 Hz (Jensen & Mazaheri, 2010; Kelly et al., 2006; Sauseng et al., 2005). For example, alpha power at posterior electrodes is reduced contralateral to the attended hemifield (Thut et al., 2006). Further, when participants are instructed to attend one of eight positions equidistant from fixation, the distribution of alpha power across the scalp systematically varies with the attended position, enabling tracking of the specific spatial location held in mind (Foster et al., 2016; Rihs et al., 2007). Visual inspection reveals graded changes in alpha topography as the attended position changes, with nearby attended locations eliciting more similar topographies. Thus, it appears that the relationship between alpha topography and spatial attention is stereotyped across individuals.

Motivated by this observation, our goal in the present work was to examine the degree to which the precise mapping between alpha topography and the locus of covert spatial attention generalizes across distinct observers. Using inverted encoding models (IEMs) and a database containing hundreds of unique subjects, we used cross-training analyses to show that alpha

topography precisely tracks the locus of spatial attention in a manner that robustly generalizes across individuals. Thus, while it has been known for some time that alpha power is reduced in locations contralateral to attended positions, our findings show that there is also a generalizable pattern of alpha activity that tracks specific positions within each hemifield. By making this dataset publicly available, our goal is to enable future studies to use our models to achieve time-resolved tracking of spatial attention in novel subjects. This approach could be especially useful in research settings where data might be difficult or costly to obtain, and may also be adopted in clinical settings where individual differences in spatial attention have been linked to attention deficit hyperactivity disorder (Epstein et al., 1997; McDonald et al., 1999).

Materials and Methods

Participants

We used EEG data from participants at the University of Oregon ($n = 56$) and the University of Chicago ($n = 90$). The data came from three publicly available datasets from published papers and one newly available dataset (Foster et al., 2016, 2017; Sutterer et al., 2019). Given that most of the data come from published papers, readers are encouraged to see the published papers for complete details on data collection procedures and experiment design. Only the final unpublished dataset will be explained in detail below. All data will be publicly available via the Open Science Framework.

In addition to this compiled dataset, we also reanalyzed publicly available data from 16 participants at the University of California, Davis (Bae & Luck, 2018). The data and task will be briefly described below, although readers are encouraged to refer to the published paper for complete details on data collection procedures and experiment design.

Task Details

Although the data come from different tasks, all data were collected while participants attended and remembered positions around fixation. In total, the data come from eight different experiments (or tasks). Generally, the tasks proceeded in three stages: stimulus, delay, and response. For all tasks, participants were first presented with one or two simple stimuli (circles, squares, gratings, or triangles; Figure 1-1) equidistant from fixation ($\sim 4^\circ$ visual angle).

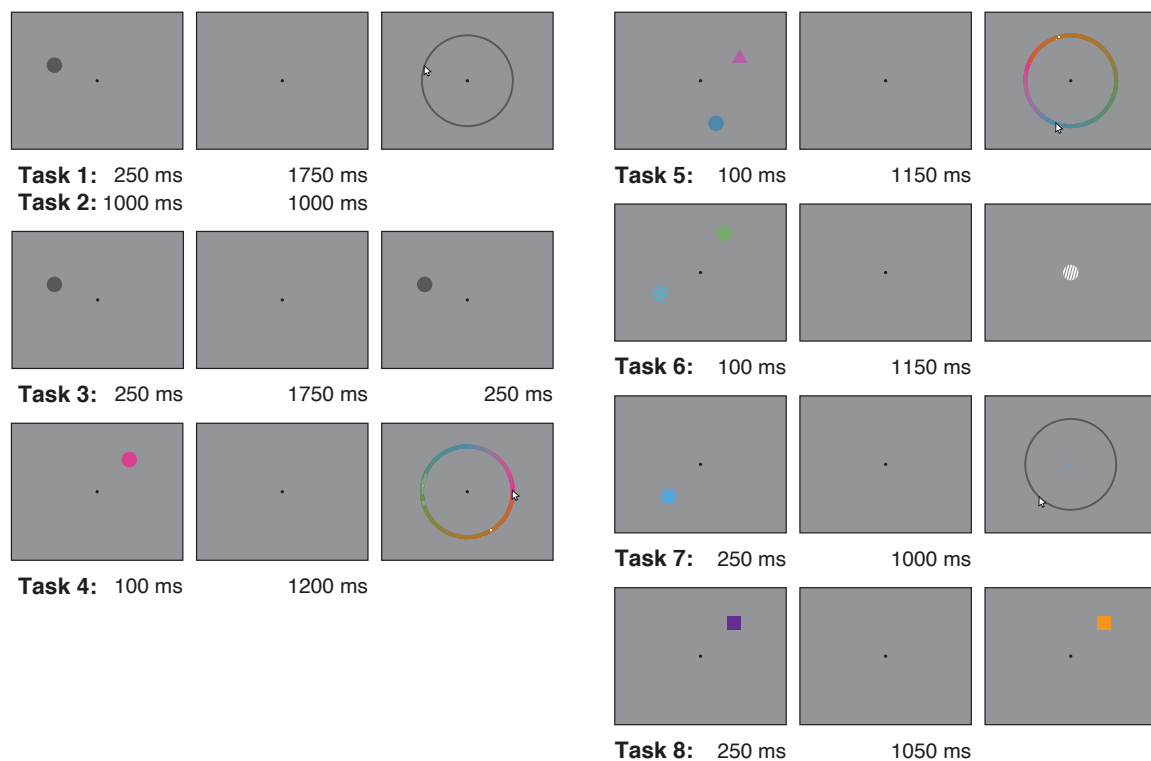


Figure 1-1. Redrawn figures of tasks.

Redrawn figures of example trials for each of the eight tasks. All tasks proceeded in three stages: stimulus array, blank delay, and response.

The stimuli were presented for varying durations ranging from 100 msec to 1000 msec (Table 1-1). Next, participants were presented with a blank delay containing only a fixation point

with durations varying between 1000 msec to 1750 msec. Finally, participants were presented with a response screen that contained either a ring presented around fixation or a single probe stimulus depending on the task. Six of the tasks were recall tasks that required participants to indicate the location, color, or orientation of the target stimulus. The remaining two tasks were change detection tasks that required participants to indicate whether there was a change in the probe stimulus.

Task	Stimulus	Delay	Response	Type	Positions	Site	<i>n</i>	Project
1	250	1750	Until response	Recall	Continuous	Oregon	15	Foster et al., 2016
2	1000	1000	Until response	Recall	Continuous	Oregon	15	Foster et al., 2016
3	250	1750	250	Change detection	Continuous	Oregon	14	Foster et al., 2016
4	100	1200	Until response	Recall	Continuous	Oregon	12	Foster et al., 2017
5	100	1150	Until response	Recall	Continuous	Chicago	18	Foster et al., 2017
6	100	1150	Until response	Recall	Continuous	Chicago	14	Foster et al., 2017
7	250	1000	Until response	Recall	Continuous	Chicago	24	Sutterer et al., 2019
8	250	1050	Until response	Change detection	Discrete	Chicago	34	Unpublished

Table 1-1. Experiment design information for the eight tasks.

Specific Task Details. Importantly, all tasks required that participants attend to locations around fixation either explicitly or implicitly (inadvertently) regardless of the response required. Specifically, in the University of Oregon sample, participants completed one of three tasks in which they had to remember either the spatial position ($n = 44$) or color ($n = 12$) of a circle stimulus presented around fixation throughout a blank delay. In the University of Chicago sample, participants completed one of four tasks. In the first set of tasks, participants were presented with two stimuli, either two color gratings or a circle and a triangle, around fixation.

Importantly, participants were cued in advance to attend and remember only one of the stimuli based on its color or shape, respectively. After a brief delay, participants then reported either the orientation or color of the target stimuli. In the second set of tasks, participants were presented with either a circle or square stimulus presented either continuously around fixation or in one of eight discrete positions equidistant from fixation. Participants then reported either the location of the circle stimulus or whether there was a change in the square stimulus. Notably, all tasks required that participants attend spatial positions around fixation and encouraged participants to maintain those spatial representations through a blank delay.

Modifications to Original Data

Due to differences in experiment design (e.g., stimulus duration) and data collection procedures (e.g., sampling rate of EEG), analyses were carried out after 1) including only the electrodes shared between experiments, 2) down sampling the EEG to the lowest sampling rate available, and 3) segmenting the data to match the shortest experiment trial. Specifically, data were analyzed from the following thirteen electrodes: C3, C4, Cz, F3, F4, Fz, O1, O2, P3, P4, P7, P8, and Pz. Data that were originally collected at 500 Hz (Chicago sample) were down sampled to 250 Hz. Given the differences in stimuli and delay durations across the tasks, individual trial durations could range from 1,250 msec to 2,000 msec. For the purposes of our analyses, data were analyzed from stimulus onset (0 msec) to 1,250 msec. This meant that the blank delay was cut short for most tasks with the new delay duration ranging from 250 msec to 1,150 msec depending on the task. Our time window of interest, then, was comprised of activity from stimulus presentation and delay period (Table 1-1).

In the original analyses presented in the published papers, encoding model analyses revealed that the distribution of alpha power across the scalp mapped onto the location participants were holding in mind. The inverted encoding model is described in detail below. In brief, the encoding model procedure occurs in two stages. In the first stage, encoding models are estimated that describe the selectivity profile of hypothesized neural populations under each electrode. Using training data, this step characterizes the relative contribution of each neural population to alpha power at each electrode. In other words, a set of data is used in the first stage to *train* the encoding model, which “learns” the relationship between alpha power topography and the possible attended locations. In the second stage, a novel set of data is used to *test* the encoding model and reconstruct the responses of the neural populations to each of the possible locations. By aggregating the responses, it is possible to determine which location was being attended. As it is typically applied, the data are split into training and testing subsets for each participant. Here, we instead test encoding models across individuals to investigate whether alpha power topography is generalizable. That is, can we train encoding models on one set of participants and successfully test the models on another set of participants?

Task 8

Participants for Task 8. Participants were recruited from the University of Chicago and the surrounding community. In total, 41 (23 female; mean age = 25 years, S.D. = 4.4) participants were recruited for the experiment. Data from seven participants were excluded due to excessive EEG artifacts.

Experimental procedures were approved by the institutional review board at the University of Chicago. All participants gave informed consent and were compensated for their

participation at a rate of \$15 per hour. Participants reported normal color vision and normal or corrected-to-normal visual acuity.

Apparatus for Task 8. Participants were tested in a dimly lit, electrically shielded chamber. Stimuli were generated using MATLAB (The Mathworks) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were presented on a 24-in. LCD monitor (refresh rate: 120 Hz, resolution: 1080 X 1920 pixels) at a viewing distance of approximately 75 cm and against a dark gray background.

EEG Acquisition for Task 8. We recorded EEG activity using 30 active Ag/AgCl electrodes mounted in an elastic cap (Brain Products actiCHamp). We recorded from International 10-20 sites, Fp1, Fp2, F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, Cz, C4, CP5, CP1, CP2, CP6, P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, and O2. Two additional electrodes were placed on the left and right mastoids, and a ground electrode was placed at position Fpz. All sites were recorded with a right-mastoid reference and were re-referenced off-line to the algebraic average of the left and right mastoids. We recorded EOG using passive electrodes with a ground electrode placed on the left cheek. Horizontal EOG (HEOG) was recorded with a bipolar pair of electrodes placed ~1 cm from the external canthus of each eye, and vertical EOG (VEOG) with a bipolar pair of electrodes placed above and below the right eye. Data were filtered on-line (low cutoff = 0.01 Hz, high cutoff = 80 Hz, slope from low-to-high cutoff = 12dB/octave). Data were digitized at 500 Hz using BrainVision Recorder (Brain Products) running on a PC. During preparation, impedances were set to be below 10k Ω .

Eyetracking for Task 8. We recorded gaze position using a desk-mounted infrared eye-tracking system (EyeLink 1000 Plus, SR Research). Gaze position was sampled at 1000 Hz. Stable head position was maintained during the task using a chin rest. The eye tracker was

recalibrated as needed throughout the session, including whenever participants removed their chin from the chin rest.

Artifact Rejection for Task 8. For artifact rejection, each trial was segmented into -600 msec pretrial and 1300 msec poststimulus array onset epochs. We used an automated procedure to flag trials that were contaminated by ocular (eye movements, blinks) or EEG artifacts, including voltage drift, dropout, blocking, and noise. When eye tracking data were not available, we used HEOG to detect saccades and VEOG to detect blinks. Next, we used this procedure as a guideline during manual visual inspection where it was ultimately determined which trials were to be rejected. Experimenters were blind to condition when inspecting the data for artifacts. Trials contaminated by artifacts were excluded from EEG analyses but not from behavioral analyses. Participants were excluded from the final sample if they had fewer than 60 artifact-free trials per location (average number of minimal trials per location = 86 out of 108).

Experiment Procedure for Task 8. Participants performed a change detection task (Figure 1-1, Task 8). The trial began with a black fixation dot (diameter = 0.25°) presented at the center of a dark gray background for a randomly determined duration between 600 msec and 1500 msec. The fixation dot remained visible throughout the trial. A stimuli array followed consisting of a single colored square (length = 1°) presented for 250 msec. The square was randomly rendered in one color from ten possible colors (red, green, blue, yellow, magenta, cyan, white, black, orange, and violet) and presented in one of eight discrete locations equidistant from fixation (4°). Participants were asked to remember the color of the stimulus over a 1,050 msec blank delay interval where only the fixation dot remained on the screen. After the delay, a single probe stimulus reappeared in the same location that was previously occupied and was either rendered in the same original color or in a different color drawn randomly from the

remaining colors. Participants used a keyboard button press to indicate whether this probe stimulus was the same or not. Participants pressed the “z” or “/” key to indicate whether it was the “same” or “different”, respectively.

Within a block, half of the trials were “same” trials, and the remaining half were “different” trials. The stimulus was equally presented in each of the eight discrete locations, although the order of trials was randomly shuffled within a block so participants could not anticipate the relevant location.

Participants completed nine blocks with each containing 96 trials. The blocks were intermixed with five blocks of another related but not currently relevant task. The task for the first block was counterbalanced across participants. The task for the second block was always the “other” task. With the exception of the first block, participants completed two blocks of Task 8 before switching to one lone block of the other (currently irrelevant) task.

Participants initiated each block by pressing the spacebar key. The experiment session was scheduled to take 3.5 hr, but the actual duration of the session depended on each participant’s pace since each block was self-initiated and breaks between blocks occurred at each participant’s request.

Bae and Luck, 2018

The reanalyzed data come from Experiment 2 of Bae and Luck (2018). The task was similar to the eight tasks described above and proceeded in the same three stages: stimulus (200 msec), delay (1300 msec), and response (Figure 1-2).




Figure 1-2. Redrawn figure of Bae and Luck (2018) task.

Redrawn figure of example trial for Bae and Luck (2018) Experiment 2 task. Participants remembered the teardrop orientation over a brief blank delay and then recreated the orientation by adjusting the orientation of the probe teardrop using the mouse cursor. The teardrop tip could appear in one of sixteen positions equidistant from fixation.

Participants were first presented with a teardrop-shaped stimulus whose tip was oriented towards one of 16 possible angular locations, discretely dividing the 360° space. Additionally, the tip of the teardrop was located at one of 16 discrete locations equidistant to fixation (2.17°). On any given trial, the teardrop orientation and location were chosen randomly out of the 256 possible combinations (16 orientations, 16 locations). Participants were asked to remember the orientation of the stimulus over a blank delay containing only a fixation point. Participants were then presented with a second teardrop stimulus at a randomly chosen location with a randomly chosen orientation, and participants used a mouse cursor to adjust the orientation to match the remembered orientation of the first teardrop. It should be noted that participants were not required or encouraged to remember the stimulus location given its lack of informative value regarding the orientation.

Given differences in experiment design and data collection procedures, some modifications were made to the available Bae and Luck (2018) data before our analyses. First, we included only the thirteen electrodes shared with our other datasets: C3, C4, Cz, F3, F4, Fz, O1, O2, P3, P4, P7, P8, and Pz. Next, trials were analyzed from stimulus onset (0 msec) to 1,250 msec, which meant that the blank delay was cut short from a duration of 1,300 msec to 1,050

msec. Finally, the 16 discrete locations were combined to create eight discrete position bins corresponding to the eight bins used in our analyses.

Alpha-Band Analysis

EEG data were segmented into epochs time locked to stimulus onset (0 msec) and beginning 800 msec before stimulus onset. For the first eight tasks, the epoch ended 1,748 msec after stimulus onset, while the Bae and Luck (2018) epochs ended 1,496 msec after stimulus onset. Alpha-band analyses were performed in MATLAB using the Signal Processing toolbox and the EEGLAB toolbox (Delorme & Makeig, 2004). The EEG signal was bandpass filtered using a two-way least squares finite impulse response filter (“eegfilt.m”, EEGLAB Toolbox). Using the Signal Processing Toolbox, a Hilbert transform was then applied to the bandpass-filtered data, which produces the complex analytic signal of the filtered EEG. The complex analytic signal was extracted for each electrode using the MATLAB syntax: `hilbert(eegfilt(data, F, f1, f2))'`, where `data` is a matrix of raw EEG (number of trials \times number of samples), `F` is the sampling frequency (250 Hz), `f1` is the lower bound of the filtered frequency band (8 Hz), and `f2` is the upper bound of the filtered frequency band (12 Hz). To compute power, the complex magnitude of the complex analytic signal was squared and then averaged across trials.

To prevent bias in our analyses, we equated the number of trials across position bins. We first calculated the minimum number of trials for any given position bin across all participants. Alpha power was then calculated for each position bin, resulting in an $p \times m \times s$ matrix of power values, where p is the number of position bins, m is the number of electrodes, and s is the number of time samples.

Given that we equated the number of trials across position bins, a random subset of trials was not included for each participant. The procedure was, therefore, repeated ten times with a

randomly generated set of trials included in each iteration, each resulting in its own $p \times m \times s$ power matrix. The IEM procedure, described below, was applied to the matrices of power values for each iteration, and their outputs (i.e., channel response profiles) were averaged. This approach was taken to reduce the influence of idiosyncrasies in estimates of alpha power that might have been specific to a certain iteration (i.e., set of chosen trials).

Inverted Encoding Model

We used an inverted encoding model (IEM) approach to reconstruct location-selective channel tuning functions (CTFs) from the distribution of alpha power across electrodes. The approach assumes that power at each electrode reflects the weighted sum of eight spatial channels, or neuronal populations, each tuned for a different angular location (Brouwer & Heeger, 2009; Foster et al., 2016; Sprague & Serences, 2013). We modeled the response of each spatial channel across angular locations as a half sinusoid raised to the seventh power based on previous work and given by: $R = \sin(0.5\theta)^7$, where θ is angular location (ranging from 0° to 359°), and R is the response of the spatial channel in arbitrary units (Figure 1-3A). This response profile was circularly shifted for each channel, so that the peak response of each spatial channel was centered over one of the eight position bins. The predicted channel responses for each position bin were derived from these basis functions using the angular location at the center of each bin.

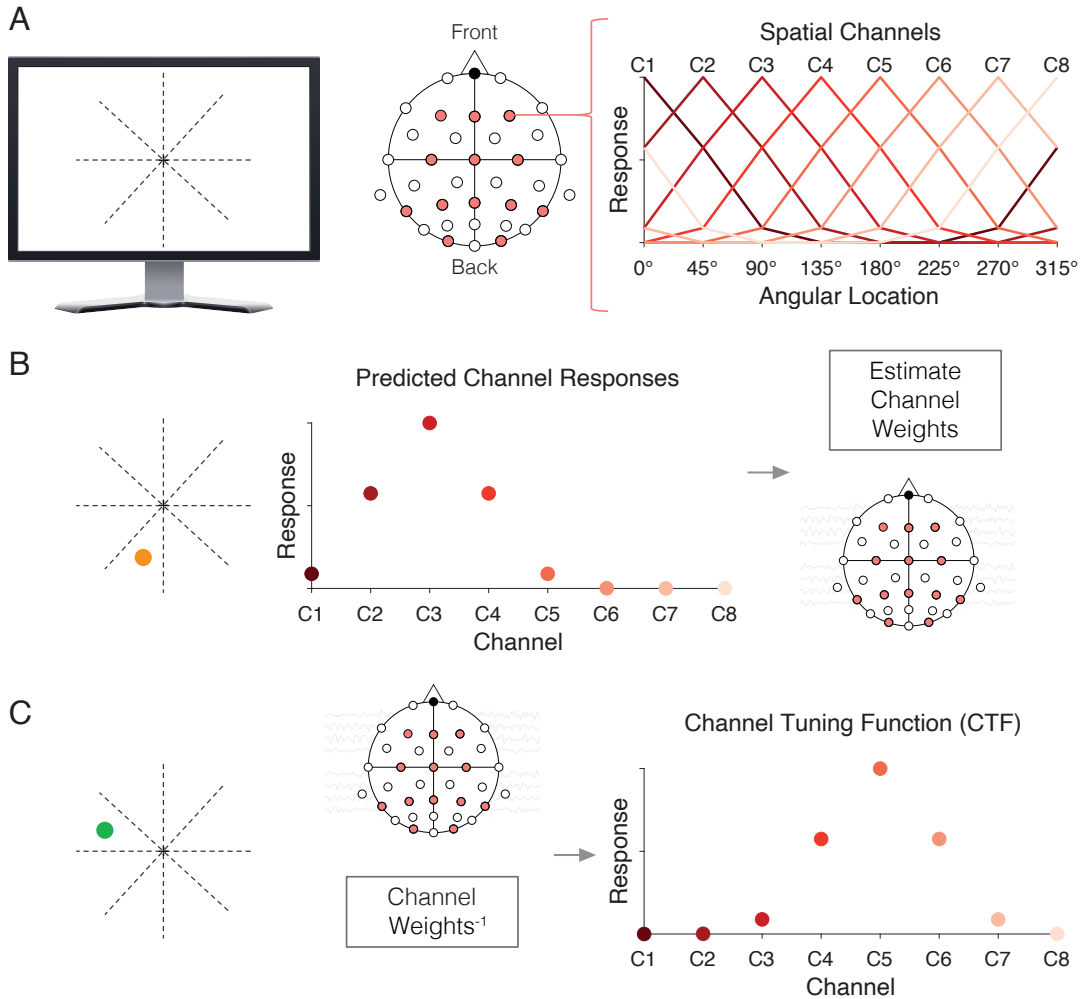


Figure 1-3. Inverted encoding model (IEM) approach.

An overview of the spatial encoding model approach used to reconstruct channel-tuning functions using alpha power topography. Alpha power at each electrode was modeled as the weighted sum of eight spatially tuned channels, or neuronal populations (C1-C8), each of which is tuned to one of the eight positions. Each curve shows the predicted response of one of the channels across the eight possible attended positions (A). In the training phase, for each position, predicted channel responses were used to estimate a set of channel weights that characterized the relative contribution of each of the spatial channels to alpha power measured at each electrode (B). In the test phase, using an independent set of data, we used the channel weights to estimate the channel responses from the observed pattern of alpha power topography for each of the eight positions (C). The channel tuning function (CTF) reflects the spatial selectivity of population-level alpha power with the peak response appearing in the channel tuned for the remembered position and gradually diminishing responses for channels tuned for positions further away.

An IEM routine was performed separately for each time point from 500 msec prior to stimulus onset through 1250 msec post-stimulus onset. This routine proceeded in two stages. In the training stage, training data (B_1) were used to estimate weights that approximated the relative contribution of the eight spatial channels to the alpha power response at each electrode (Figure 1-3B). If B_1 (m electrodes \times n_1 observations) is the power at each electrodes for each measurement in the training set, C_1 (k channels \times n_1 observations) is the predicted response of each spatial channel for each measurement, and W (m electrodes \times k channels) is the weight matrix that characterizes a linear mapping from “channel space” to “electrode space”, then 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: $\hat{W} = B_1 C_1^T (C_1 C_1^T)^{-1}$.

In the testing stage, we used the weights calculated in the training stage above and inverted the encoding model to transform the observed test data B_2 (m electrodes \times n_2 observations) into estimated channel responses, \hat{C}_2 , (k channels \times n_2 observations): $\hat{C}_2 = (\hat{W}^T \hat{W})^{-1} \hat{W}^T B_2$ (Figure 1-3C). Each estimated channel response function was circularly shifted to a common center by aligning the estimated channel responses to the channel tuned for that stimulus bin to obtain CTFs.

For our analyses, we used either a “ k -fold” or “leave-one-out” cross validation procedure to ensure that our training and testing data were independent. In essence, some data of estimated power values served as the training data, B_1 and were used to estimate W , and the remaining data served as our testing data, B_2 and was used to estimate our channel responses, C_2 . The manner in which data were split depended on the analysis and on the relevant question of interest. The procedure was repeated until each split of the data was held out as the test set, and the resulting CTFs were averaged across all iterations.

Unlike previous work, the training data for the IEM procedure consisted of multiple subjects. That is, we took an approach of averaging power values across groups of subjects. On the other hand, the testing data was always comprised of an individual subject as in previous work. Importantly, each training subject contributed the same amount of data given that we equated the number of trials per position bin across all subjects. We also z-scored the raw EEG for each subject to ensure that the training data were not heavily influenced by any outliers. This standardization was employed at each time point to minimize any influence that might arise due to the differences in timing between the different tasks. Although it is likely that the *exact* contribution of each spatial channel to alpha power at each electrode varies by subject, we aimed to isolate a generalizable signal of spatial attention, if it exists, by averaging power values across subjects.

Statistical Analysis

As in previous work, we used linear regression to estimate the slope of our channel response after collapsing across channels that were equidistant from the channel tuned to the location of the stimulus. To control for differences across participants, the channel responses were first scaled to a fixed range from 0 to 1 for each participant. A bounded range was used to suppress the effect of potential outliers, which could occur if a participant's channel responses showed a steep slope but rather small channel responses or large but not graded channel responses that gave rise to steep slopes. The scaling was performed at each time point to further suppresses any effect of timing given the different tasks. We operationalized CTF slope as location selectivity, where steeper slopes indicate greater location selectivity.

To test whether the selectivity of the CTF was reliably above chance, we tested whether CTF slope was greater than zero using a one-sample *t*-test. Given that mean CTF slope may not

be normally distributed under the null hypothesis, we used a Monte Carlo randomization procedure to empirically approximate the null distribution of the t -statistic. Specifically, we implemented the IEM procedure as described above but randomized the position labels within each run, so that the labels were random with respect to the responses recorded at each electrode. This randomization procedure was repeated 1,000 times to obtain a null distribution of t -statistics. We then calculated the probability of obtaining a t -statistic from the surrogate null distribution greater than or equal to the observed t -statistic (i.e., one-tailed) to determine whether location selectivity was reliably above chance (i.e., the probability of a Type 1 error). The location selectivity of the CTF then was considered reliably above chance if the probability of a Type 1 error was less than 0.01. This procedure was applied at each time point to see whether location selectivity differed throughout the trial.

Results

Location selectivity across participants

First, we tested whether training on a large group of subjects showed location selectivity for an independent set of subjects. Following a k -fold cross-validation procedure, participants ($n = 146$) were randomly split into ten independent groups of 14 or 15 participants each. Encoding models were then trained using the data from nine groups and tested on each of the participants in the group that was left out. This was repeated until each group served as the test case.

If the pattern of alpha power across electrodes carries information about the stimulus location, the IEM should reveal a graded CTF profile with a peak in the channel tuned for the remembered location. If there is no location tuning, then the IEM should produce a flat CTF profile, indicating that the pattern of alpha power across electrodes does not contain any

location-specific information. We found a graded CTF profile after averaging CTFs across the trial (0-1250 msec, Figure 1-4).

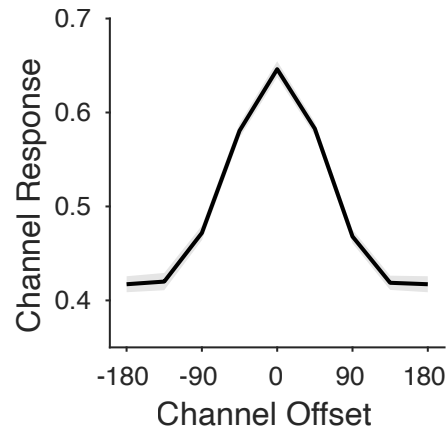


Figure 1-4. Alpha power channel tuning function (CTF) across participants.

Averaged CTF reconstructed from alpha (8–12 Hz) power after stimulus onset (averaged from 0–1,250 msec). Participants ($n = 146$) were divided into ten groups. The alpha topography of nine groups was used to decode the spatial position of the participants in the remaining group, until each group served as the held-out group.

We further examined the reconstructed channel response profile for each of the eight positions to investigate whether the eight positions were equally represented in the averaged CTF. We observed a graded response profile for each of the eight positions, such that the peak channel responses were for the channel tuned for the remembered location, suggesting that alpha power topography *precisely* tracked the angular location throughout the trial (Figure 1-5).




Figure 1-5. Channel responses for each stimulus location based on location channel.

Channel responses reconstructed from alpha power after stimulus onset (averaged from 0–1,250 msec) shown for each of the eight stimulus location bins.

A permutation test was performed at each time point in the trial to identify the points at which CTF slope was reliably above zero. We found that CTF slope was reliably above zero shortly after stimulus presentation and remained reliably above zero throughout the trial (Figure 1-6). This finding suggests that alpha power topography generalizes across participants and can be used to track locations across participants in a *sustained* and *precise* manner.




Figure 1-6. CTF slope over time.

Slopes of averaged CTFs reconstructed from alpha (8–12 Hz) power across time. Black markers indicate the time points at which CTF slope was reliably above chance.

Location selectivity remained after removing task- and site-specific information

Next, we set out to investigate the extent to which the encoding models relied on task- or site-specific information that might have been shared between participants. Participants were split into eight groups based on task. Following a leave-one-out procedure, encoding models were trained using the data from seven of the groups and tested on each of the participants in the group that was left out.

In a separate analysis, participants were split into two groups based on research site. Following a similar leave-one-out procedure, encoding models were trained using one group of subjects and tested on each of the participants in the other group. This latter analysis was important to rule out the possibility of duplicate participants across datasets. Specifically, it is possible that subjects participated in more than one of the tasks included. As much as was possible, duplicate subjects were removed, but identifiers were not available for all subjects. The possibility remains that our large sample of subjects contains data from the same subject.

However, it is highly unlikely that any participant was able to participate at both the University of Oregon and the University of Chicago.

For both of these analyses, we again found a graded CTF profile after averaging CTFs across the trial (0-1250 msec; Figure 1-7A). Permutation tests revealed that CTF slope was reliably above zero for nearly the entire trial (Figure 1-7B).

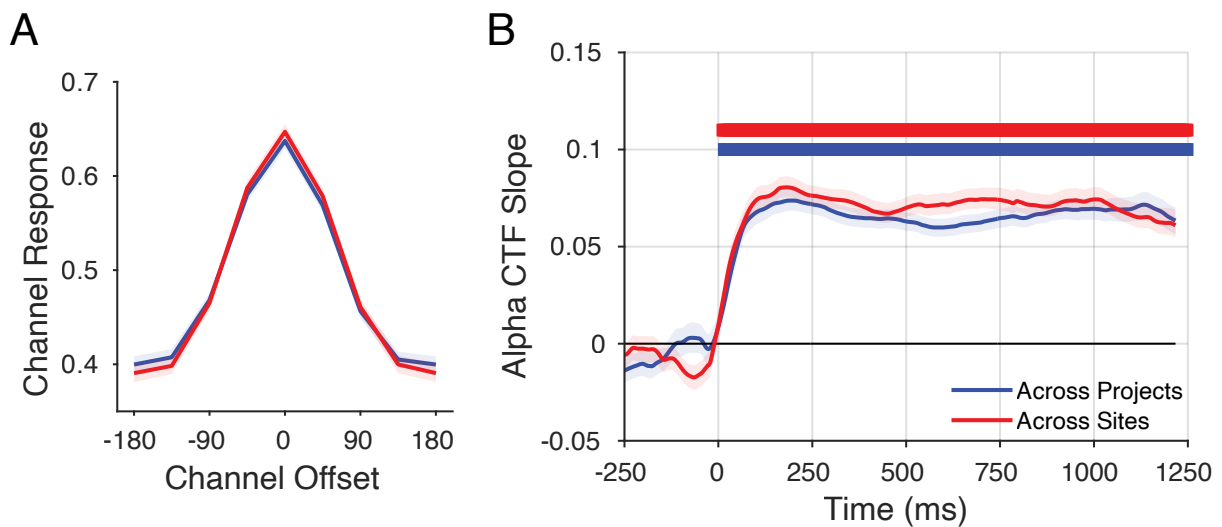


Figure 1-7. CTFs and their slope over time across projects and sites.

Averaged CTFs reconstructed from alpha power after stimulus onset (0 – 1250 msec) across projects (blue; A) and across sites (red; A). Slopes of averaged CTFs across time. Blue and red markers indicate the time points at which CTF slope was reliably above chance (B).

However, the response profiles for each of the eight position bins was orderly but not precise, such that the peak response did not always occur in the channel tuned for the remembered location (Figure 1-8).

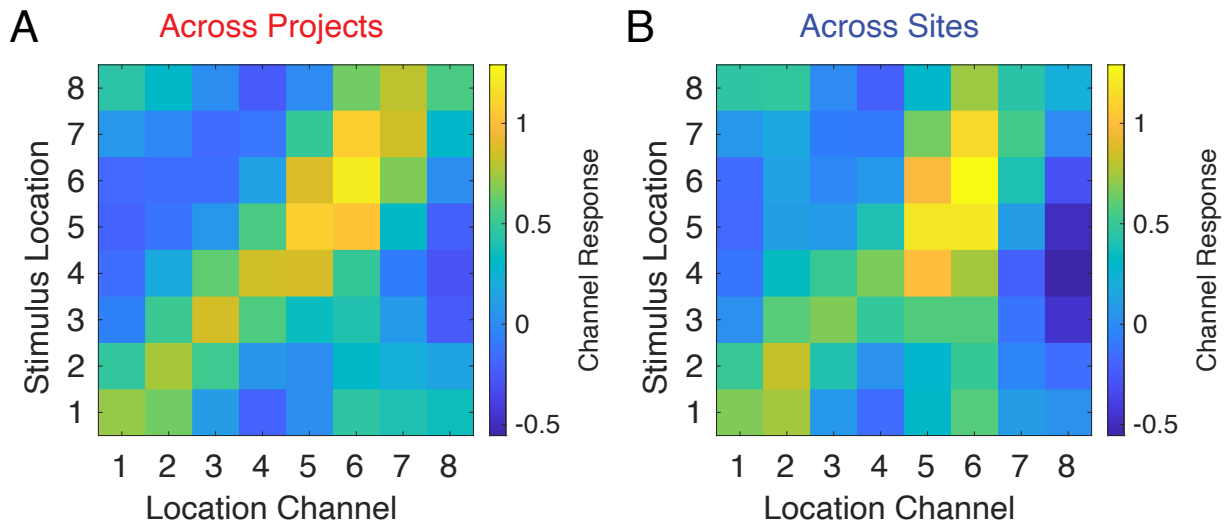


Figure 1-8. Channel responses across projects and sites.

Channel responses reconstructed from alpha power after stimulus onset (averaged from 0–1,250 msec) for each of the eight stimulus location bins across projects (A) and across sites (B).

Nonetheless, averaging channel responses across positions that were equidistant from the remembered location showed that, on average, the peak response occurred in the channel tuned for the remembered location across the trial (Figure 1-9). Thus, location selectivity using alpha power remained after removing any task- and site-specific information that might have been shared between participants. The contribution of shared site-specific information (e.g., EEG system) is a potential avenue for future research.

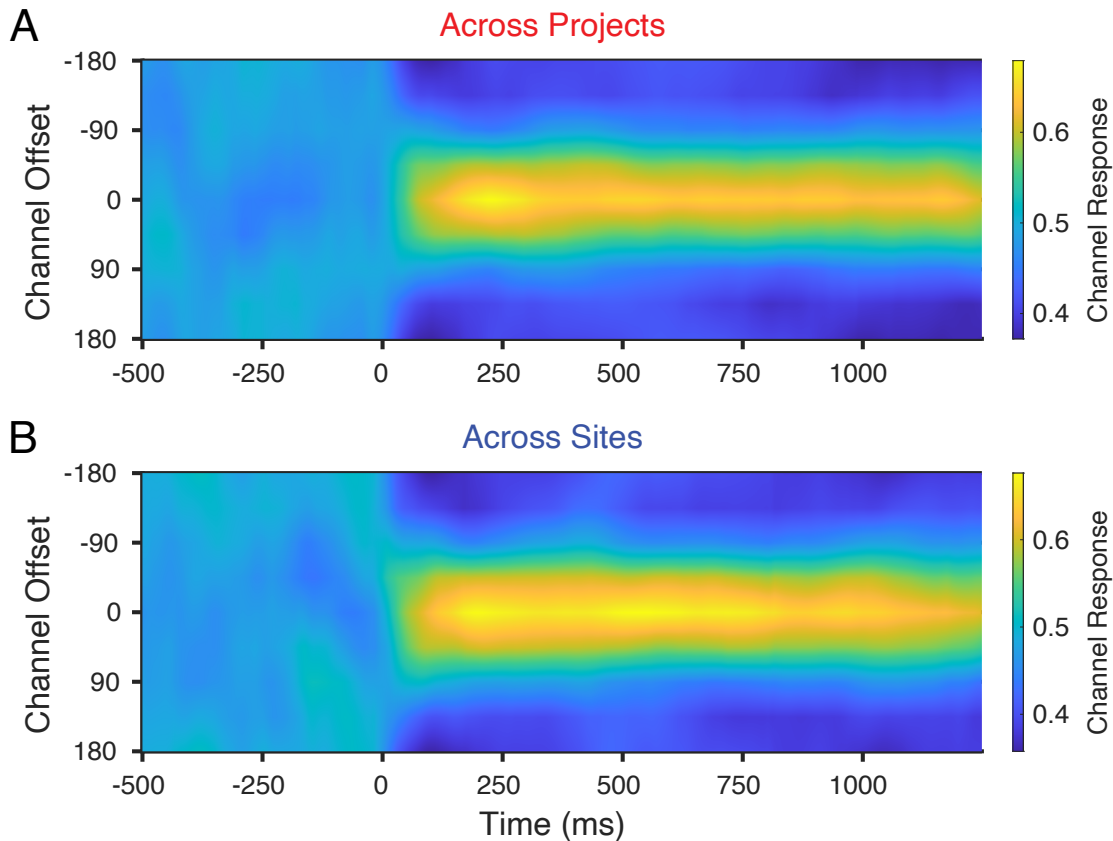


Figure 1-9. CTFs over time across projects and sites.

CTFs reconstructed from alpha activity across the entire trial across projects (A) and across sites (B).

Location selectivity across research groups

Finally, we tested the generalizability of alpha power topography by examining whether training encoding models using our entire group of participants revealed location information about a group of participants from another research group. Encoding models were trained using our set of participants ($n = 146$) and tested on each of the participants in Experiment 2 of Bae and Luck (2018). In this task, participants ($n = 16$) were presented with teardrop stimuli for 200 msec (Figure 1-2). Importantly, the teardrop stimulus could appear in 16 discrete positions around fixation. After a brief blank delay of 1300 msec, participants adjusted a teardrop probe to

match the remembered orientation of the first teardrop. Although participants were not instructed to remember the location of the stimulus, Bae and Luck (2018) found that alpha power could be used to track the location of the teardrop in memory, therefore making this dataset a good candidate for testing whether alpha power topography generalizes across research groups.

We found a graded CTF profile after averaging CTFs across the trial (0-1250 msec; Figure 1-10A). Permutation tests revealed that CTF slope was reliably above zero shortly after stimulus onset and lasting around 1,000 msec (Figure 1-10B). Alpha power, then, can generalize across subjects from different research groups using different tasks and different procedures.

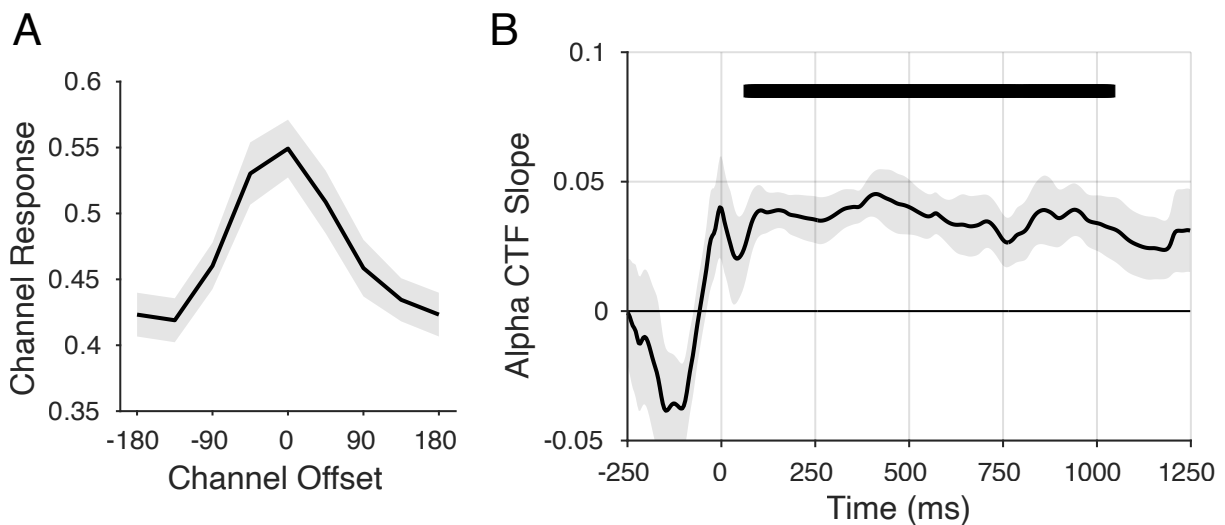


Figure 1-10. CTF and CTF slope across research groups.

Averaged CTFs reconstructed from alpha power after stimulus onset (0 – 1250 msec) after training on 146 participants and testing on Experiment 2 of Bae and Luck, 2018 (A). Slopes of averaged CTFs across time (B). Black markers indicate the time points at which CTF slope was reliably above chance.

Discussion

We compiled EEG recordings from 146 participants while they completed tasks that guided their attention to different locations around fixation. Our goal was to investigate whether alpha power topography varied with attended position systematically across participants. Training inverted encoding models (IEMs) on the alpha power topography of one group of participants allowed the tracking of attended positions in a new group of participants. These cross-training results persisted even when the groups of participants completed different tasks and came from different sites with different EEG systems and data collection procedures. Together, the current work shows that the relationship between alpha power topography and spatial attention is stereotyped across observers and that a database of EEG data characterizing this relationship allows for the time-resolved tracking of spatial attention in novel subjects.

It should be noted that there have been several claims of perfect performance for biometric recognition using EEG signals, indicating idiosyncratic information in EEG (e.g., Chen et al., 2016; Ruiz-Blondet et al., 2016). These demonstrations, however, often have small sample sizes or contain data from only one session, which could inflate recognition accuracy given session-specific conditions. Nevertheless, recognition rates using data across sessions still provide recognition rates that are higher than chance (Armstrong et al., 2015). EEG alpha power in particular shows high test-retest reliability and showed high heritability in a meta-analysis with 79% of variance explained by genetic factors (Mathewson et al., 2015; Salinsky et al., 1991; van Beijsterveldt & van Baal, 2002). Here, we show that there exists a generalizable component of alpha power despite the unique differences across individuals.

The logistics of data collection can make it a costly endeavor given the monetary expenses and time demands. Considerable advancements have been made over the last decade to

make structural and functional MRI data publicly available through data sharing projects or repositories (Eickhoff et al., 2016; Poldrack et al., 2013). Despite the established record of publicly available MRI datasets, EEG data sharing has seen a fraction of that progress although various researchers have adopted the practice of making data publicly available upon publication. One early example provides ~1,500 one- and two-minute EEG recordings from 109 participants completing motor and imagery tasks, while another dataset contains recordings from 36 participants both before and during the completion of a mental arithmetic task (Goldberger et al., 2000; Schalk et al., 2004; Zyma et al., 2019). Another dataset contains EEG data from 54 participants completing three different tasks across multiple sessions held days apart (Lee et al., 2019). In the clinical realm, a noteworthy endeavor by Temple University has made available a corpus of almost 17,000 EEG sessions from 10,874 unique subjects (Obeid & Picone, 2016).

Although most publicly available EEG datasets exist in isolation, there are efforts to document and even consolidate these datasets in a centralized space. For instance, the neuroimaging platform OpenNeuro began allowing the storage and sharing of EEG data in 2019 after EEG was incorporated into a common standardized data format for neuroimaging known as the Brain Imaging Data Structure (BIDS) (Markiewicz et al., 2021; Pernet et al., 2019). As of now, the OpenNeuro repository provides access to almost 100 public EEG datasets with over 3,500 participants, while MRI datasets exceed 550 with over 21,000 participants.

We provide the current dataset to the community at large to encourage re-analysis by researchers looking to answer related research questions using a wide array of techniques. Large datasets specifically provide more statistical power for detecting nuanced effects and decrease susceptibility to spurious effects (Yarkoni et al., 2010). Data sharing in general exploits already available data, leading to reduced costs in data collection and advances in the field through the

discovery of corroborative or novel findings (Wicherts et al., 2006). It is unsurprising that there have been growing efforts to make data publicly available after publication (King, 2007; Schofield et al., 2009; Wicherts, 2011; Wicherts et al., 2006). After all, data sharing contributes to increased transparency, accountability, and reproducibility, all of which lend credibility to scientific findings (Nosek et al., 2012; Vision, 2010).

PART II.

Tracking individuated items in mind

CHAPTER 2.

Distinct signals track items and locations in working memory.

Visual working memory (WM) is an online memory system that enables the rapid access and updating of information in the service of other cognitive tasks. Capacity limits in WM exhibit robust correlations with broad measures of intellectual ability (Vogel & Awh, 2008), suggesting that WM is integral to complex cognition. Thus, there has been strong interest in delineating the behavioral and neural processes that determine these capacity limits. Multiple neural correlates have been found to monotonically scale with the number of items in visual WM up to measured storage limits (~3 items), showing sustained activity over the delay period (Fukuda et al., 2015; Todd & Marois, 2004; Vogel & Machizawa, 2004). Moreover, these signals predict individual differences in visual WM capacity (Fukuda et al., 2015; Luria et al., 2016; Vogel & Machizawa, 2004). Using scalp-EEG and the event-related potential (ERP) technique, past work has routinely found that a contralateral and sustained negative voltage deflection, or the contralateral delay activity (CDA), shows reliable decreases in amplitude as items are added into WM (Luria et al., 2016; Vogel & Machizawa, 2004). Similarly, oscillatory activity in the alpha-band has also been shown to track the number of items in WM (Busch & Herrmann, 2003; Sauseng et al., 2009). Previous work often focused on lateralized signals that required lateralized task designs and distractors to balance visual stimulation for both hemispheres. More recently, Fukuda, Mance, and Vogel (2015) used a whole-field design that did not require distractors to examine whether whole-field signals were also sensitive to the number of items in WM. They found that parieto-occipital alpha power (7–9 Hz) and a sustained negative voltage deflection over parieto-occipital electrodes both tracked the number of items in working memory.

Indeed, it has been suggested that these two signals are manifestations of the same neural process, such that amplitude modulations of oscillatory activity can explain the generation of slow evoked components, like the CDA (Mazaheri & Jensen, 2008; van Dijk et al., 2010). In fact, lateralized sustained event-related fields (ERFs) in MEG and lateralized alpha power were found to be strongly correlated, both spatially and temporally, in a working memory task (van Dijk et al., 2010). However, the whole-field signals in Fukuda et al. (2015) showed distinct time courses and explained unique variance in predicting individual differences in WM capacity, suggesting that sustained potentials and alpha-band oscillations reflect distinct aspects of storage in visual WM (Bae & Luck, 2018; Hakim et al., 2019; Wang et al., 2019, 2020).

Here, we examine the hypothesis that each signal provides unique information about the contents of visual WM. Specifically, the negative slow wave can track the number of individuated items in WM, while parieto-occipital alpha power tracks the number of relevant locations in working memory. The distinction between items and locations is a subtle distinction that has been examined before using the CDA, which was found to be sensitive to the number of items in working memory regardless of the number of locations (Ikkai et al., 2010). On the other hand, oscillatory activity in the alpha frequency band (8-12 Hz) is known to track the deployment of spatial attention during storage in visual WM (Foster et al., 2016; Rihs et al., 2007), but extant work has not discriminated between item-based and location-based explanations of this oscillatory activity. According to our hypothesis, parieto-occipital alpha power will be primarily sensitive to the number of attended locations within the visual display rather than the number of individuated items.

We used a perceptual grouping manipulation to discriminate between item-based and location-based neural activity. Past work has shown that perceptual grouping can yield strong

improvements in WM performance such that a larger number of elements can be remembered when they are grouped. This effect has been observed using a variety of grouping cues, including proximity and connectedness cues (Jiang et al., 2004; Woodman et al., 2003; Xu, 2006), color similarity (Brady & Tenenbaum, 2013; Gao et al., 2016; Morey, 2019; Morey et al., 2015; Peterson & Berryhill, 2013; Quinlan & Cohen, 2012; Shen et al., 2012), shape similarity (Mate & Baqués, 2009), amodal completion (Walker & Davies, 2003), depth cues (Kristjánsson, 2006), and collinearity and closure cues (Gao et al., 2016). Here, we used collinearity cues to manipulate the number of individuated items, while holding constant the number of relevant locations. Thus, this manipulation provided traction for distinguishing between item-based and location-based neural signals.

In Experiment 2-1, we first establish that parieto-occipital alpha power and the negative slow wave track the number of items stored with both color and spatial memoranda, replicating the findings of Fukuda et al. (2015) and also extending them to a spatial working memory task. Our primary motivation for Experiment 2-1 was to ensure that both signals, especially parieto-occipital alpha power, were not disrupted by a spatial task. This anticipatory analysis was done to ensure a fair comparison between the two signals in the second experiment, where the number of items was manipulated, but the number of spatial locations was not. Specifically, we use collinearity cues to perceptually group elements in Experiment 2-2, while holding constant the number of elements in the display. Pairs of stimuli were either aligned to create the percept of a single item or misaligned to create the percept of two items. To anticipate the results, alpha power was sensitive to the number of locations that were attended regardless of whether the items were grouped or not, while the negative slow wave tracked the number of individuated

items stored, revealing a smaller number of items with grouping by collinearity. These findings provide further evidence for the distinction between spatial and item-based signals in visual WM.

Materials and Methods

Participants

Participants were recruited from the University of Chicago and the surrounding community. Overall, 16 (6 females, mean age = 22) and 23 (14 females, mean age = 22) participants were run in Experiment 2-1 and 2-2, respectively. For Experiment 2-1, data from two participants were excluded due to the participants' voluntary withdrawal during the experiment session. Additionally, data from two participants in Experiment 2-1 and three participants in Experiment 2-2 were excluded due to excessive EEG artifacts (< 150 trials remaining per condition).

Experimental procedures were approved by the Institutional Review Board at the University of Chicago. All participants gave informed consent and were compensated for their participation at a rate of \$15 per hour. Participants reported normal color vision and normal or corrected-to-normal visual acuity.

For Experiment 2-1, our intended sample size was 16 participants given previous research showing that this is a sufficient number of participants to observe the time course of set size effects on both the negative slow wave and parieto-occipital alpha power (Fukuda et al., 2015). However, data from 4 of our 16 participants were excluded from analyses given insufficient number of trials after artifact rejection or due to participants' voluntary withdrawal from the experiment session. For Experiment 2-2, our intended sample size was 20 participants given that our aim was to go one step further and investigate whether our set size effects were sensitive to perceptual grouping. Data from 3 of our initial 20 participants were excluded from

our analyses after artifact rejection. We replaced these 3 participants and collected data from an additional 3 participants to complete our intended sample size.

Apparatus

Participants were tested in a dimly lit, electrically shielded chamber. Stimuli were generated using MATLAB (The Mathworks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were presented on a 24 in. LCD monitor (refresh rate: 120 Hz, resolution: 1080 x 1920 pixels) at a viewing distance of approximately 75 cm and against a dark gray background.

EEG Acquisition

We recorded EEG activity using 30 active Ag/AgCl electrodes mounted in an elastic cap (Brain Products actiCHamp, Munich, Germany). We recorded from International 10/20 sites Fp1, Fp2, F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, Cz, C4, CP5, CP1, CP2, CP6, P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, and O2. Two additional electrodes were placed on the left and right mastoids, and a ground electrode was placed at position Fpz. All sites were recorded with a right-mastoid reference and were re-referenced offline to the algebraic average of the left and right mastoids. We recorded electrooculogram (EOG) using passive electrodes with a ground electrode placed on the left cheek. Horizontal EOG was recorded with a bipolar pair of electrodes placed ~1cm from the external canthus of each eye, and vertical EOG with a bipolar pair of electrodes placed above and below the right eye. Data were filtered online (low cut-off = 0.01 Hz, high cut-off = 80 Hz, slope from low-to-high cut-off = 12 dB/octave) and were digitized at 500 Hz using BrainVision Recorder (Brain Products, Munich, Germany) running on a PC. During preparation, impedances were set to be below 10 k Ω .

Eyetracking

We recorded gaze position using a desk-mounted infrared eye-tracking system (EyeLink 1000 Plus, SR Research, Ontario, Canada). Gaze position was sampled at 1000 Hz. Stable head position was maintained during the task using a chin rest. The eye tracker was re-calibrated as needed throughout the session, including whenever participants removed their chin from the chin rest.

Artifact Rejection

For artifact rejection, each trial was segmented into -400 msec pretrial and 1,750 msec post-stimulus array onset epochs. We used an automated procedure to flag trials that were contaminated by ocular or EEG artifacts. Next, we used this procedure as a guideline during manual visual inspection where it was ultimately determined which trials were to be rejected. Experimenters were blind to condition when inspecting the data for artifacts. Trials contaminated by artifacts were excluded from EEG analyses but not from behavioral analyses. Participants were excluded from the final sample if they had fewer than 150 artifact-free trials per condition.

An automated artifact detection procedure was used to detect eye movements, blinks, and EEG artifacts. Trials were flagged as containing a saccade if the Euclidean vector between the mean gaze positions in the first and second halves of an 80-msec sliding window (advanced in 10-msec increments) was greater than 0.5° of visual angle. When eye tracking data were not available, we used horizontal EOG to detect saccades. Trials were flagged as containing a saccade if the mean voltage during the first and second halves of a 150-msec sliding window (advanced in 10-msec steps) exceeded $20 \mu\text{V}$.

For blinks, trials were flagged as containing a blink if the eye tracker could not detect the pupil at any point during the trial. When eye tracking data were not available, we used vertical

EOG to detect blinks. Trials were flagged as containing a blink if the mean voltage during the first and second halves of a 150-msec sliding window (advanced in 10-msec steps) exceeded 30 μV .

For EEG artifacts, we flagged trials as containing voltage drifts (e.g., skin potentials) if the absolute change in voltage from the first quarter of the trial to the last quarter of the trial exceeded 100 μV . We flagged trials as including a sudden step in voltage (which can occur when an electrode is damaged) if the mean voltage during the first and second halves of a 250-msec sliding window (advanced in 20-msec increments) differed by more than 100 μV . We marked trials as containing high-frequency noise (e.g., muscle artifacts) if any electrode had a peak-to-peak amplitude greater than 150 μV within a 15-msec sliding window (advanced in 50-msec increments). Finally, we flagged trials as containing amplifier saturation if any electrode had 60 time-points within a 200-msec sliding window (advanced in 50-msec increments) that were within 1 μV of each other.

Negative Slow Wave Analysis

EEG activity was calculated using a baseline from -400 msec to 0 msec relative to the onset of the stimulus array. The mean baseline amplitude was subtracted from EEG amplitude at each time point in the trial. The baselined trials were then averaged for each condition to create event-related potentials (ERPs) for each condition. We included twelve parieto-occipital electrodes chosen a priori and based on previous findings: P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, and O2. Statistical analyses were performed on data that were not filtered beyond the .01- to 80-Hz on-line data-acquisition filter. We low-pass-filtered the data (30 Hz) for illustrative purposes in the figures.

Parieto-Occipital Alpha Power Analysis

EEG signal processing was performed in MATLAB (The Mathworks, Natick, MA). We band-pass filtered the raw EEG data using a filter from the FieldTrip toolbox (ft_preproc_bandpassfilter.m; Oostenveld et al., 2011) and then extracted instantaneous power values for the alpha band (8-12 Hz) by applying a Hilbert transform (hilbert.m) to the filtered data. We calculated alpha power for the same parieto-occipital electrodes as in the ERP analysis: P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, and O2. For illustrative purpose in the figures, we subtracted the mean baseline (-400 msec to 0 msec) at each time point in the trial for each condition and converted to percent change from baseline.

Stimuli

For both experiments, a black fixation dot (diameter = 0.20°) was presented at the center of a dark gray background and remained visible for the entire trial. The stimuli were presented within a predetermined area (exp. 2-1: $19.65^\circ \times 14.90^\circ$; exp. 2-2: $15.80^\circ \times 15.80^\circ$) and at least 0.75° (exp. 2-1) or 3° (exp. 2-2) away from fixation. During the delay interval, only the fixation dot remained on the screen. After the delay, a single probe stimulus reappeared.

For Experiment 2-1, the stimuli array consisted of either one or three circles (0.75°) placed randomly within a quadrant, without replacement and at least 2.25° away from each other if set size three. For the color change detection task, the circles were each randomly rendered in one color from 7 possible colors (red, green, blue, yellow, magenta, white, and black) without replacement. For the spatial change detection task, the circles were rendered in black. After the delay interval, a single probe stimulus reappeared that was either the ‘same’ as one of the originally presented stimuli or ‘different’ from any of the originally presented stimuli. For the color task, a change could occur in the color of the stimulus, while a change in spatial location could occur in the spatial task. Specifically, the change in spatial location could occur in any

angular direction ($0^\circ - 359^\circ$) and at a randomly determined distance (range: 3° to 3.75°) relative to the original stimulus. The probe stimulus was constrained to its original quadrant and had to maintain the same minimum distance requirement (2.25°) from the original stimuli locations.

For Experiment 2-2, the stimuli array consisted of either two or four black circles (3°) with rectangular gaps. The locations of the stimuli were assigned in a serial manner. The location of the first stimulus was randomly selected to fall within the predetermined bounds. In relation to the first stimulus, the location of the second stimulus could occur in a randomly determined angular direction ($0^\circ - 359^\circ$) and 5.65° away from the first stimulus, while maintaining the requirements of the first stimulus (i.e., remaining within the predetermined bounds), and effectively forming a pair. For set size four, the third and fourth stimuli locations were determined in a similar manner with the additional constraints that stimuli be at least 6° away from stimuli belonging to a different pair and that pairs of stimuli did not intersect. On half of the trials, the gaps of each pair of stimuli were oriented towards each other to form the percept of either one (set size two) or two (set size four) items (grouped condition). On the remaining trials, the gaps were misaligned (ungrouped condition). The same location configurations were used across conditions (grouped and ungrouped) but were unique for each participant. To avoid accidental pairs in the ungrouped condition, the orientation of the second stimulus in each pair had to be more than 5° clockwise or counterclockwise from the orientation it would be in the grouped condition.

Experiment 2-1 Procedure

Participants performed a whole-field change detection task (Figure 2-1). The trial began with a fixation dot presented at the center of the screen for a randomly determined duration between 600 – 1500 msec. A stimuli array followed and was presented for 250 msec. The stimuli

array consisted of either one or three circles that were either rendered in color (color change detection task) or in black (spatial change detection task). Participants were asked to remember as many of the stimuli as possible over a 1,500 msec blank delay interval. After the delay, a single stimulus reappeared. Participants used a keyboard button press to indicate whether this probe stimulus was presented at one of the previously occupied locations. For the color task (Figure 2-1A), a change could occur in the color of the stimulus, while a change in spatial location could occur in the spatial task (Figure 2-1B). Participants pressed the 'z' or '/' key to indicate whether the probe stimulus was the “same” or “different”, respectively, from the stimuli display. There were no practice trials given before the formal experiment. Participants were given verbal and written task instructions with the aid of an example trial image similar to Figure 2-1.

Within a block, half of the trials were “same” trials, and the remaining half were “different” trials. Similarly, half of the trials were set size one, and the remaining half were set size three. Participants completed 20 blocks with each containing 60 trials of either the color or spatial change detection task. There were an equal number of color and spatial change detection blocks, which were completed in an alternating order. The task for the first block was counterbalanced across participants. Participants self-initiated each block by pressing the spacebar key. The experiment session was scheduled to take 3 hours, but the actual duration of the session (~2.5 –3.5 hrs.) depended on each participant’s pace since they initiated each block and decided when (and if) to take breaks between blocks.

Experiment 2-2 Procedure

The procedure was similar to Experiment 2-1 with the following exceptions. The stimuli array consisted of either two or four black circles with rectangular gaps (Figure 2-1C). On half of

the trials, the gaps of each pair of stimuli were oriented towards each other to form the percept of either one (set size two) or two (set size four) items. On the remaining trials, the gaps were misaligned (Figure 2-1D). After the delay, participants indicated whether the orientation of the probe stimulus that reappeared was the “same” or “different”. Half of the trials were set size two, and the remaining half were set size four. Participants completed 30 blocks of 40 trials each. The experiment session was scheduled to take 3.5 hours (~3 – 4 hrs.).

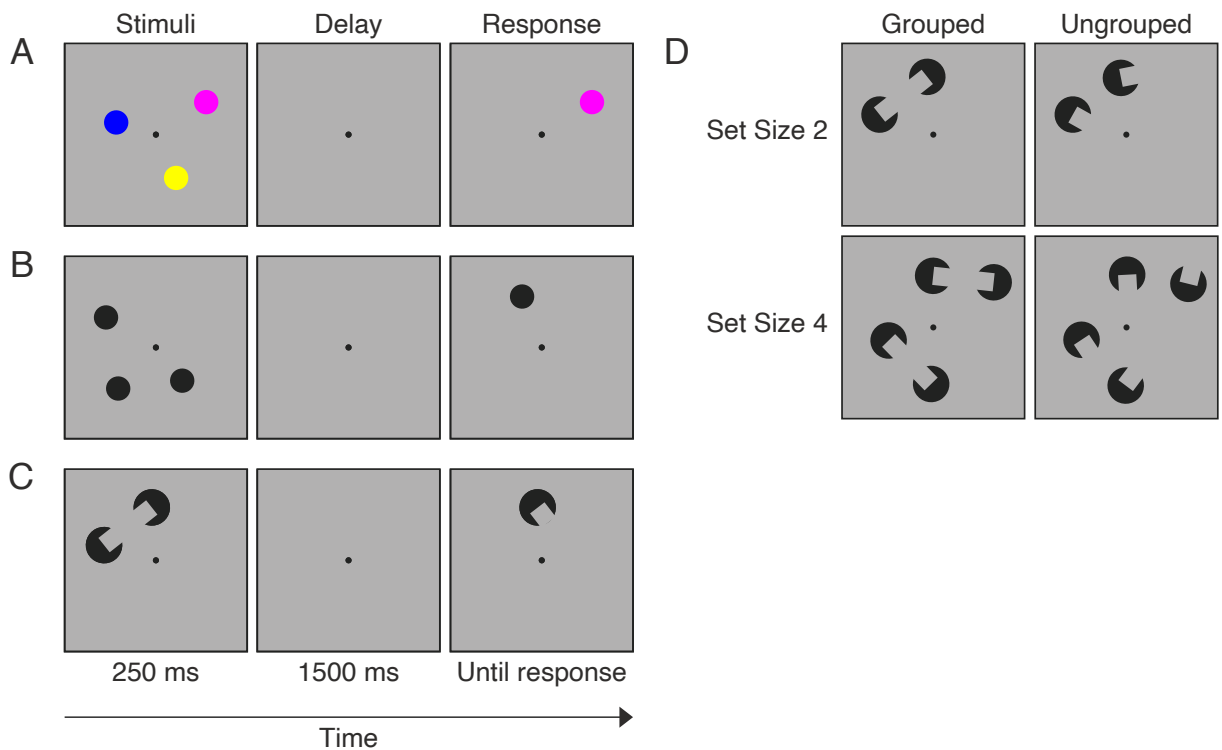


Figure 2-1. Change detection tasks and perceptual grouping displays.

Color (A) and spatial (B) change detection tasks from Experiment 2-1. Perceptual grouping change detection task from Experiment 2-2 (C) with example stimuli displays for each of the four conditions (D).

Experimental Design and Statistical Analysis

Both experiments used a 2×2 within-subject design. For Experiment 2-1, the factors were set size (1 or 3) and task type (color or spatial). The type of task alternated with each block. For Experiment 2-2, the factors were set size (2 or 4) and grouping condition (grouped or ungrouped).

Behavioral data were analyzed using a repeated measures ANOVA. Neural data were analyzed using repeated measures ANOVAs on averaged delay activity data and cluster-based permutation tests (Maris and Oostenveld, 2007; Sassenhagen and Draschkow, 2019) on data averaged over parieto-occipital electrodes. In Experiment 2-1, our focus was on ANOVAs given that we were investigating robust set size effects that had been observed before. In Experiment 2-2, we focused on cluster-based permutation tests to search for novel grouping effects in a data-driven way. For consistency, we present results using both approaches for each experiment. For the cluster-based permutation tests, we first reduced our high-dimensional data (# of time-points $\times n$ subjects; per condition) to a single value. Repeated measures t -tests were calculated at each time-point to assess the difference between conditions (e.g., set size one vs set size three), which resulted in a map of t -scores across time. Time-points were thresholded according to an a priori defined criterion (which corresponded to a p -value of 0.05, two-sided) and adjacent time-points with t -scores that exceeded this value were grouped together to form a cluster. Clusters were summarized into a single number by summing the t -values, which produced a single value for each cluster. It is important to note that the extent of the cluster became fixed at the end of this first step and individual time-points were not visible to the next inference step. Specifically, the cluster structure was our only test statistic and no statistical inference was made about individual time-points. Nevertheless, we next calculated the probability that these values came from a null distribution. Permutation tests were used to establish the probability of our data under the null

hypothesis given that it is unclear what distribution of t -value sums would be expected under the null hypothesis. The number of permutations was 10,000 or the maximum possible, whichever was lowest. On each iteration and for each time point, it was randomly determined if the first condition was subtracted from the second condition or vice versa for each subject. Then for each iteration, the cluster formation step was repeated. The cluster with the highest sum of t -values was identified, and the sum of its t -values was stored and became our surrogate-null value for that permutation. After all iterations, the cumulative density of these surrogate-null values was our approximation of the values under the null hypothesis. The p -value then was calculated as the percentage of surrogate-null values that the observed data exceeded.

Experiment 2-1 Results

There was a main effect of set size ($F(1,11) = 17.99, p = 0.001, \eta^2 = 0.62$) and task ($F(1,11) = 8.18, p = 0.016, \eta^2 = 0.43$) on accuracy, such that accuracy was higher for set size one ($M = 0.95, SD = 0.05$) than set size three ($M = 0.90, SD = 0.08$) and for the color task ($M = 0.94, SD = 0.05$) than the spatial task ($M = 0.90, SD = 0.08$). There was no significant interaction between set size and task on accuracy ($F(1,11) = 0.05, p = 0.84, \eta^2 = 0.004$).

We used repeated measures ANOVAs on averaged delay activity from parieto-occipital electrodes to analyze the effects of set size and task type on both the negative slow wave and parieto-occipital alpha power. Moreover, we used nonparametric cluster-based permutation tests (Maris & Oostenveld, 2007) to corroborate the effects of set size and task type in a more nuanced manner. Given previous research (Fukuda et al., 2015), we predicted that both signals would show characteristic set size effects. Specifically, Fukuda et al. (2015) found a monotonic increase in slow wave negativity and a monotonic decrease in parieto-occipital alpha power with increases in set size up to typical capacity (~3 items).

The negative slow wave and parieto-occipital alpha power were analyzed during the delay period (250 msec – 1750 msec). As predicted, there was a main effect of set size on voltage at parieto-occipital electrodes (Figure 2-2A; $F(1,11) = 16.99, p = 0.002, \eta^2 = 0.61$), such that the amplitude was more negative for set size three than set size one. There was no main effect of task ($F(1,11) = 0.47, p = 0.51, \eta^2 = 0.04$) and no significant interaction ($F(1,11) = 0.29, p = 0.60, \eta^2 = 0.03$). Cluster-based permutation tests echoed these findings. Three clusters in the voltage at parieto-occipital electrodes were found to be sensitive to set size and extended from ~130 to ~210 msec, ~490 to ~1020 msec, ~1020 to ~1190 msec. The permutation test indicated that the effects of set size were significant ($p = 0.048, p = 0.007, p = 0.044$). There were no clusters that were sensitive to task type. In short, we replicated and extended previous findings that the negative slow wave is sensitive to set size in a color change detection task (Fukuda et al., 2015), as well as in a spatial change detection task.

The main effect of set size on parieto-occipital alpha power did not reach significance when averaging over the entire delay interval (Figure 2-2B; $F(1,11) = 3.06, p = 0.11, \eta^2 = 0.22$), though an early transient effect of set size was evident with more alpha power suppression for set size three than set size one. There was a main effect of task ($F(1,11) = 6.83, p = 0.024, \eta^2 = 0.38$), which was driven by the rapid return to baseline in the color task for set size three trials. Importantly, the interaction did not reach significance ($F(1,11) = 1.03, p = 0.33, \eta^2 = 0.09$). Cluster-based permutation tests revealed more nuanced results. There was a cluster sensitive to set size that extended from ~-30 to ~870 msec. The permutation test indicated that there was a significant effect of set size ($p = 0.006$). There was also a cluster that was sensitive to task, extending from ~680 to ~1560 msec. The permutation test indicated a significant effect of task ($p = 0.003$). Thus, we again replicated previous findings from Fukuda et al. (2015) that parieto-

occipital alpha power shows a monotonic decrease with an increase in set size for both a color and spatial change detection task.

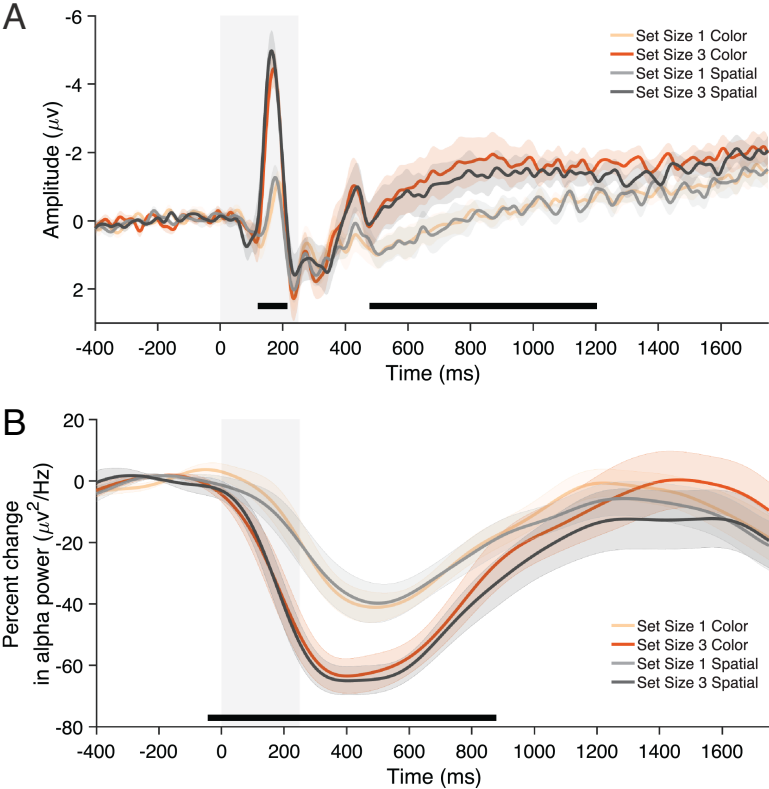


Figure 2-2. Negative slow wave and alpha power over time in Experiment 2-1.

Averaged negative slow wave (A) and averaged alpha power suppression (B) observed at parieto-occipital electrodes in Experiment 2-1. Shaded regions indicate duration of stimuli display. Black bars indicate clusters showing a significant ($p < .05$) set size effect from cluster-based permutations tests.

In summary, our findings from Experiment 2-1 replicate those of Fukuda et al. (2015) and extend them to spatial memoranda. For both color and spatial memoranda, both the negative slow wave and parieto-occipital alpha power are sensitive to set size. In Experiment 2-2, we investigate whether these set size effects are further shaped by collinearity cues. Do grouping

cues that compel the perception of fewer individuated objects also affect the magnitude of these storage-related signals?

Experiment 2-2 Results

There was a main effect of set size ($F(1,19) = 110.3, p < 0.001, \eta^2 = 0.85$) and grouping ($F(1,19) = 146.7, p < 0.001, \eta^2 = 0.89$) on accuracy, such that accuracy was higher for set size two ($M = 0.94, SD = 0.05$) than set size four ($M = 0.86, SD = 0.10$) and for grouped trials ($M = 0.95, SD = 0.05$) compared to ungrouped trials ($M = 0.85, SD = 0.10$). There was also a significant interaction between set size and grouping on accuracy ($F(1,19) = 116.1, p < 0.001, \eta^2 = 0.86$), such that there was a greater benefit of grouping for set size four than set size two.

Nonparametric cluster-based permutation tests were used to analyze the effects of set size and grouping on both the negative slow wave and parieto-occipital alpha power. In comparison to Experiment 2-1, the grouping effects of interest in Experiment 2-2 were more novel and the timing and duration of any grouping effect was not known. Accordingly, choosing an objective window of interest for ANOVAs was not possible. Instead we focused on cluster-based permutation tests, however we also report the results of repeated measures ANOVAs using averaged delay activity from three equally sized windows of 500 msec each (early, middle, and late delay).

Given previous research and our own data, we again predicted that both signals would show characteristic set size effects, with higher slow wave negativity and reduced alpha power for set size four compared to set size two. The central question of the study, however, was to investigate whether these signals were sensitive to grouping when set size effects were present. To this end, we first identified clusters that were sensitive to set size and then examined whether those clusters were also sensitive to grouping.

First, there was a main effect of set size and grouping on voltage at parieto-occipital electrodes during the early time window (250 – 750 msec; Table 2-1). A cluster in the voltage at parieto-occipital electrodes extended from ~330 to ~880 msec. The cluster-based permutation test indicated that there was a significant effect of set size (Figure 2-3A; $p = 0.005$), such that amplitude was more negative for set size four than set size two. Within this cluster, another cluster was identified that extended from ~370 to ~670 msec. The permutation test indicated that there was a significant effect of grouping ($p < 0.001$), such that amplitude was more negative when stimuli were ungrouped relative to grouped. To assess the evidence for grouping during this set size window, estimated Bayes factors were calculated ($BF_{01} = 0.89$) and indicated that the data were 1.13 times more likely to have occurred under the alternative hypothesis.

Terms	F-statistics	df1	df2	p-value	η^2
<i>Negative Slow Wave</i>					
Early (250 – 750 msec)					
Set Size (2,4)	22.21	1	19	< 0.001	0.54
Condition (Grouped, Ungrouped)	9.11	1	19	0.007	0.32
Set Size * Condition	3.07	1	19	0.10	0.14
Middle (750 – 1250 msec)					
Set Size (2,4)	5.11	1	19	0.036	0.21
Condition (Grouped, Ungrouped)	3.49	1	19	0.077	0.16
Set Size * Condition	2.39	1	19	0.14	0.11
Late (1250 – 1750 msec)					
Set Size (2,4)	3.33	1	19	0.084	0.15
Condition (Grouped, Ungrouped)	0.64	1	19	0.43	0.03
Set Size * Condition	2.76	1	19	0.11	0.13
<i>Alpha Power</i>					
Early (250 – 750 msec)					
Set Size (2,4)	14.51	1	19	0.001	0.43
Condition (Grouped, Ungrouped)	2.97	1	19	0.10	0.14
Set Size * Condition	0.08	1	19	0.78	0.00
Middle (750 – 1250 msec)					
Set Size (2,4)	10.32	1	19	0.005	0.35
Condition (Grouped, Ungrouped)	3.01	1	19	0.10	0.14
Set Size * Condition	4.89	1	19	0.04	0.21
Late (1250 – 1750 msec)					
Set Size (2,4)	2.61	1	19	0.12	0.12
Condition (Grouped, Ungrouped)	0.58	1	19	0.46	0.03
Set Size * Condition	3.84	1	19	0.065	0.17

Table 2-1. Repeated measures ANOVAs for neural data in Experiment 2-2.

We applied 2 (Set Size 2, 4) \times 2 (Grouped, Ungrouped) repeated measures ANOVAs to voltage and alpha power at parieto-occipital electrodes averaging over delay activity in three equally sized windows of 500 msec each. Significant ($p < 0.0167$; Bonferroni-corrected for $\alpha = 0.05$) effects in bold.

Thus, we first replicated our finding from Experiment 2-1 that the negative slow wave was modulated by set size. Moreover, we extended those findings by demonstrating that the negative slow wave was also modulated by grouping. The reduction in amplitude of the negative slow wave for grouped stimuli is consistent with our hypothesis that the negative slow wave tracks the number of individuated items.

Halfway through the delay, we observed an unexpected drop in amplitude in the negative slow wave for Set Size 4 Ungrouped, such that the set size effect was eliminated. Although the

mechanisms involved are unknown, it is possible that this pattern is due to the strategic reorienting of attention to the remembered items within capacity. In line with this speculation, we found that the drop in the negative slow wave was restricted to participants with high Set Size 4 Ungrouped performance, which is consistent with the idea of strategic refocusing. Additional work is needed to bolster the findings of this post-hoc analysis.

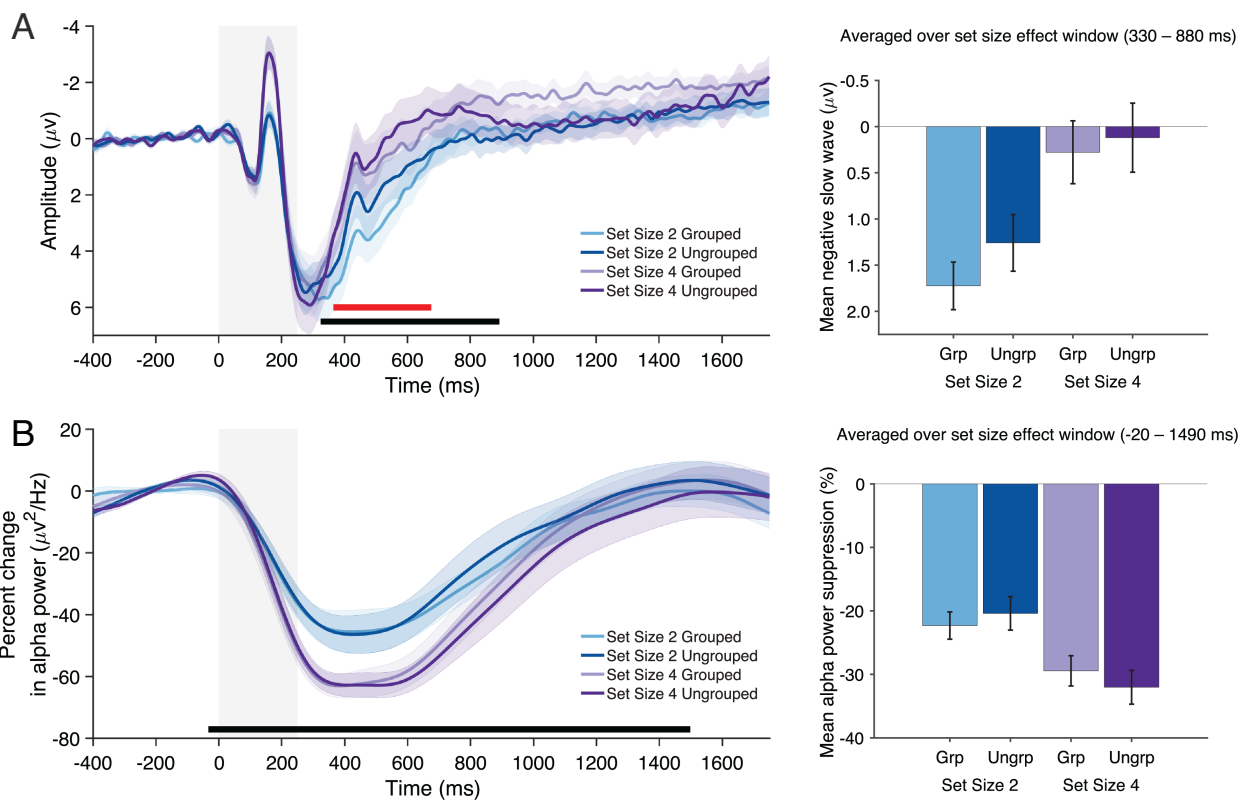


Figure 2-3. Negative slow wave and alpha power over time in Experiment 2-2.

Averaged negative slow wave (A) and averaged alpha power suppression (B) observed at parieto-occipital electrodes in Experiment 2-2. Shaded regions indicate duration of stimuli display. Black and red bars indicate clusters showing a significant ($p < .05$) set size or grouping effect, respectively, from cluster-based permutation tests.

In alpha-power, there was a main effect of set size during the early and middle time windows (250 – 750 msec, 750 – 1250 msec; Table 2-1), but no main effect of grouping. A cluster extended from ~-20 to ~1490 msec with the permutation test indicating that there was a significant effect of set size (Figure 2-3B; $p < 0.001$), such that alpha power was lower for set size four than set size two. Despite this robust effect of the number of positions, however, no clusters showing this set size effect were sensitive to the effects of perceptual grouping. Estimated Bayes factors during this set size window were calculated ($BF_{01} = 1.51$) and indicated that the data were 1.5 times more likely to have occurred under the null hypothesis. This suggests that participants maintained their attention on the same number of positions, even though the number of individuated items stored was reduced in the grouped condition.

Discussion

In line with previous research, perceptual grouping cues enhanced visual working memory (VWM) performance and also led to distinct changes in EEG signals that track WM storage. Specifically, the negative slow wave – a signal known to track the number of items stored in visual WM (Fukuda et al., 2015) – was reduced in amplitude for grouped stimuli, consistent with the idea that grouping reduces the number of individuated items. By contrast, parieto-occipital alpha power reliably tracked the number of locations that were attended, but this signal was completely unaffected by perceptual grouping. This pattern of results suggests that the negative slow wave is sensitive to the number of individuated items in WM, while parieto-occipital alpha power tracks the number of visually attended locations.

These diverging effects of perceptual grouping dovetail with recent evidence showing that the EEG signals that track WM storage fall into at least two distinct classes. Although both parieto-occipital slow waves and alpha power track the number of relevant items in a working

memory task, they have distinct temporal profiles, and explain distinct between-subject variance in WM capacity (Fukuda et al., 2015). Moreover, these two EEG signals respond distinctly to manipulations of attentional priority in WM (Günseli et al., 2019) and to whether the task motivates the maintenance of item-based information (Hakim et al., 2019).

The perceptual grouping effect on the negative slow wave is similar to the effects reported in studies that examined contralateral delay activity (CDA), a lateralized event-related potential (ERP) component that is observed at similar electrode sites as the negative slow wave and is highly sensitive to the number of items in VWM (Luria et al., 2016; Vogel & Machizawa, 2004). Previous research characterizing the negative slow wave was driven by the goal of finding a whole-field alternative to the CDA that would not require the filtering of distractors (Fukuda et al., 2015). The findings so far are consistent with the hypothesis that the CDA and negative slow wave are indexing the same neural operation, though further work is needed to determine how these signals are related. That said, previous research shows that the CDA tracks the number of items regardless of the number of relevant locations and even as the items undergo physical transformations. For instance, the CDA reliably tracks additional items that are added into VWM regardless of whether these additional items are presented in the same locations as previously encoded memoranda or not, which suggests that it ultimately tracks the number of items rather than the number of attended positions (Ikkai et al., 2010). Moreover, the CDA is sensitive to common fate cues, such that there is a reduction in the CDA when common motion cues encourage the perception of multiple elements as a single item (Luria & Vogel, 2014). There is also a reduction in the CDA for identical colors relative to distinct colors, which suggests that identical colors might be maintained as a group in VWM (Gao et al., 2011; Peterson et al., 2015).

Importantly, the current work extends these findings by providing a direct measure of information maintained about each of the grouped elements, specifically spatial information. Alpha oscillations provided a sensitive index of the number of attended positions and revealed that perceptual grouping yielded no reduction in the number of attended positions. Specifically, parieto-occipital alpha power continued to index the same number of spatial locations even when the stimuli were grouped, which suggests that both constituents of the group were maintained. This finding rules out the possibility that participants maintained only half of the group in grouped trials and used that information to inform their decision at the time of response. Instead, perceptual grouping boosted orientation WM performance by reducing the number of individuated representations stored without sacrificing spatial information about each of the grouped elements. Our findings are broadly consistent with past work that found a reduced CDA for identical colors presumably maintained as a group (Gao et al., 2011; Peterson et al., 2015). Future research could extend the current work to include these similarity cues or other grouping cues that are especially compelling, such as common onset or common fate, or cues that have been found to boost working memory performance, such as proximity (e.g., Woodman et al., 2003).

Finally, the distinct effects of perceptual grouping on item-based and spatial EEG signals are consistent with a broad class of models that distinguish between the number of individuated objects maintained in working memory and the specific featural information that is stored about each of those objects (Awh et al., 2007; Fukuda et al., 2010; Xu & Chun, 2006). For example, Xu and Chun (2009) described a neural object file theory that distinguishes between an initial stage of object individuation and a subsequent stage in which the details of the individuated objects are apprehended. This account is supported by their finding that grouped objects elicited

lower fMRI responses in inferior intraparietal sulcus (IPS) (a brain region thought to track the number of individuated items in WM) but higher responses in superior IPS and the lateral occipital complex (LOC), regions that are thought to be sensitive to the information load or complexity of the stored items (Xu & Chun, 2007). According to the neural object file account, grouping reduced the number of individuated items in a capacity-limited individuation stage, allowing more information to be relayed to higher visual areas in the identification stage. Our results bolster this interpretation by showing that distinct neural signals track the number of individuated items in WM and the details of the specific locations associated with each item. Thus, while neural signals tracking WM storage are sometimes viewed in a monolithic fashion, there is growing evidence that distinct delay signals with distinct computational roles work in parallel to maintain information in an online state (Hakim et al., 2020).

CHAPTER 3.

Evidence for an item-based pointer system.

Working memory is an online memory system that allows us to store, manipulate, and rapidly access information, paving the way for complex behaviors like arithmetic and reading comprehension. There is great interest in understanding the neural basis of working memory given its fundamental role in everyday life and its implications in clinical disorders, like attention-deficit hyperactive disorder (Karatekin & Asarnow, 1998; Kuntsi et al., 2001; Westerberg et al., 2004).

Using scalp electroencephalography (EEG), previous work has described neural signals that track the number of individuated items in mind (Vogel & Machizawa, 2004; Fukuda et al., 2015; Luria et al., 2016; Diaz et al., 2021). For example, Vogel and Machizawa (2004) observed a contralateral and sustained negative slow wave at posterior electrodes that decreased in amplitude as more items occupied working memory but only until capacity was reached (~3 items). Similar signals tracking the number of items in working memory, or WM load, have also been found using functional MRI (Todd & Marois, 2004, 2005; Xu & Chun, 2006).

Recent work has proposed that signals purported to track WM load might actually be tracking the product of a more fundamental cognitive operation, a spatiotemporal “pointer system” that enables continuous tracking of items in both space and time (Thyer et al., 2022). This perspective falls in line with past proposals that have distinguished between mechanisms for the spatiotemporal tracking of items, on the one hand, and the representation of the specific features associated with each tracked item on the other hand (e.g., Kahneman et al., 1992; Pylyshyn, 2009). Moreover, past work has found that distinct cortical regions carry information about the *number* of items stored and the *content* of those memory representations (Xu & Chun,

2006). Thyer et al. (2022) bolstered this proposal by using multivariate analysis of voltage topography (Adam et al., 2020) – hereafter referred to as “mvLoad” – to reveal a common signature of load across variations in both the type and amount of information stored about each item in working memory. Specifically, the pattern of voltage activity for any given number of items looked the same regardless of whether participants maintained color information, orientation information, or both. Additionally, this load-sensitive signal was sensitive to perceptual grouping such that it tracked the number of *perceived* items and not the number of constituent parts or the number of attended positions (Thyer et al., 2022). At the same time, univariate changes in posterior alpha power (8 – 12 Hz) remained sensitive to the spatial extent of the attended regions in these memory trials (Diaz et al., 2021). Thus, the initial work suggests that this multivariate signature of WM load cannot be explained by variations in the spatial extent of covert attention, but instead appears to reflect the number of individuated *items* in visual working memory.

Those findings notwithstanding, one limitation of the extant work is that it has employed stimulus displays that distinguish items from each other based only on their spatial positions. But if the proposed pointer operation tracks items through both space and *time*, then the mvLoad signal should track the total number of stored items even when the items are separated by temporal positions alone. Here, we tested this prediction. To this end, participants stored two successive arrays of items in working memory while EEG activity was recorded. In one condition, the second array of items appeared in novel positions. Meanwhile in the other condition, the second array of items appeared in precisely the same positions as in the first array so that only temporal position separated the items in the first and second arrays.

To preface the results, we demonstrate that oscillatory activity in the alpha frequency band (8 – 12 Hz) at posterior electrodes differed between conditions and tracked the number of relevant locations in mind as demonstrated in previous work (Diaz et al., 2021). In contrast, multivariate analysis of EEG voltage enabled robust decoding of the total number of items stored in both conditions, and load-specific patterns of activity were equivalent between conditions. These findings reinforce our proposal that load-sensitive signals in EEG activity tap into a pointer operation that tracks stored items based on both spatial and temporal coordinates.

Materials and Methods

Participants

Participants were recruited from the University of Chicago and the surrounding community. In total, 12 (9 female; mean age = 25 years, S.D. = 3.7) and 27 (18 female, mean age = 25 years, S.D. = 4.0) participants were run in Experiments 3-1 and 3-2, respectively. Data from two participants in Experiment 3-1 and seven participants in Experiment 3-2 were excluded because of excessive EEG artifacts (< 140 (Exp. 3-1) and 200 (Exp. 3-2) trials remaining per condition).

Experimental procedures were approved by the institutional review board at the University of Chicago. All participants gave informed consent and were compensated for their participation at a rate of \$15 per hour. Participants reported normal color vision and normal or corrected-to-normal visual acuity.

Apparatus

Participants were tested in a dimly lit, electrically shielded chamber. Stimuli were generated using MATLAB (The Mathworks) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were presented on a 24-in. LCD monitor (refresh rate: 120 Hz, resolution:

1080 × 1920 pixels) at a viewing distance of approximately 75 cm and against a dark gray background.

EEG Acquisition

We recorded EEG activity using 30 active Ag/AgCl electrodes mounted in an elastic cap (Brain Products actiCHamp). We recorded from International 10–20 sites Fp1, Fp2, F7, F3, Fz, F4, F8, FC5, FC1, FC2, FC6, C3, Cz, C4, CP5, CP1, CP2, CP6, P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, and O2. Two additional electrodes were placed on the left and right mastoids, and a ground electrode was placed at position Fpz. All sites were recorded with a right-mastoid reference and were rereferenced off-line to the algebraic average of the left and right mastoids. We recorded EOG using passive electrodes with a ground electrode placed on the left cheek. HEOG was recorded with a bipolar pair of electrodes placed ~1 cm from the external canthus of each eye, and vertical EOG with a bipolar pair of electrodes placed above and below the right eye. Data were filtered on-line (low cutoff = 0.01 Hz, high cutoff = 80 Hz, slope from low-to-high cutoff = 12 dB/octave) and were digitized at 500 Hz using BrainVision Recorder (Brain Products) running on a PC. During preparation, impedances were set to be below 10 k Ω .

Eyetracking

We recorded gaze position using a desk-mounted infrared eye-tracking system (EyeLink 1000 Plus, SR Research). Gaze position was sampled at 1000 Hz. Stable head position was maintained during the task using a chin rest. The eye tracker was recalibrated as needed throughout the session, including whenever participants removed their chin from the chin rest.

Artifact Rejection

For artifact rejection, each trial was segmented into –400 msec pretrial and 1550 msec poststimulus array onset epochs. We used an automated procedure to flag trials that were

contaminated by ocular or EEG artifacts. Next, we used this procedure as a guideline during manual visual inspection where it was ultimately determined which trials were to be rejected. Experimenters were blind to condition when inspecting the data for artifacts. Trials contaminated by artifacts were excluded from EEG analyses but not from behavioral analyses. Participants were excluded from the final sample if they had fewer than 200 artifact-free trials per condition in Experiment 3-1 (average trials per condition = ~ 277 out of 320) and fewer than 140 artifact-free trials per condition in Experiment 3-2 (average = ~ 198 out of 240).

An automated artifact detection procedure was used to detect eye movements, blinks, and EEG artifacts. Trials were flagged as containing a saccade if the Euclidean vector between the mean gaze positions in the first and second halves of an 80-msec sliding window (advanced in 10-msec increments) was greater than 0.5° of visual angle. When eye tracking data were not available, we used HEOG to detect saccades. Trials were flagged as containing a saccade if the mean voltage during the first and second halves of a 150-msec sliding window (advanced in 10-msec steps) exceeded $20 \mu\text{V}$.

For blinks, trials were flagged as containing a blink if the eye tracker could not detect the pupil at any point during the trial. When eye tracking data were not available, we used vertical EOG to detect blinks. Trials were flagged as containing a blink if the mean voltage during the first and second halves of a 150-msec sliding window (advanced in 10-msec steps) exceeded $30 \mu\text{V}$.

For EEG artifacts, we flagged trials as containing voltage drifts (e.g., skin potentials) if the absolute change in voltage from the first quarter of the trial to the last quarter of the trial exceeded $100 \mu\text{V}$. We flagged trials as including a sudden step in voltage (which can occur when an electrode is damaged) if the mean voltage during the first and second halves of a 250-msec

sliding window (advanced in 20-msec increments) differed by more than 100 μV . We marked trials as containing high-frequency noise (e.g., muscle artifacts) if any electrode had a peak-to-peak amplitude greater than 150 μV within a 15-msec sliding window (advanced in 50-msec increments). Finally, we flagged trials as containing amplifier saturation if any electrode had 60 time points within a 200-msec sliding window (advanced in 50-msec increments) that were within 1 μV of each other.

Experiment 3-1 Procedure

Participants performed a change detection task (Figure 3-1). The trial began with a black fixation dot (diameter = 0.20°) presented at the center of a dark gray background for a randomly determined duration between 600 msec and 1500 msec. The fixation dot remained visible throughout the trial. A stimuli array followed consisting of one or two colored squares (set size two or four, respectively; length = 2°) presented for 150 msec. The stimuli were presented within a predetermined area (exp. 3-1: $9.90^\circ \times 9.90^\circ$; exp. 3-2: $12.5^\circ \times 12.5^\circ$) and at least 2.1° (exp. 3-1) or 3° (exp. 3-2) away from fixation. The predetermined area was divided into a 4×4 grid each of which could contain a single stimulus. The positions needed for the first array were randomly chosen from the 16 possible locations without replacement. In the same location condition, the same locations were used for the second array. Otherwise, new locations were randomly selected from the set of possible locations remaining that were not already used in the first array. Jitter ($\sim 0.25^\circ$) was added to the actual locations occupied by stimuli.

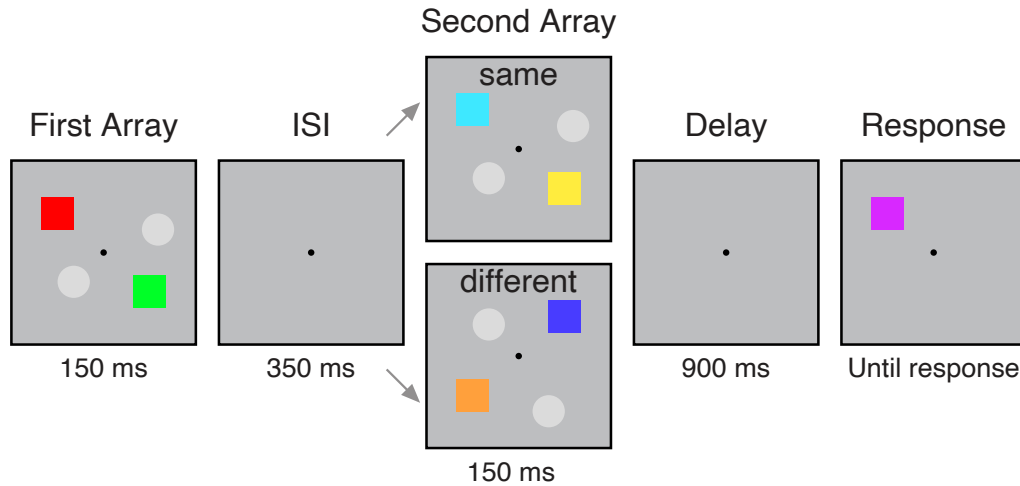


Figure 3-1. Sequential change detection task.

Example trials for sequential change detection task. Gray placeholders were not presented in Experiment 3-1. For both Experiment 3-1 and 3-2, two stimuli arrays were separated by a brief interstimulus interval. The first stimuli array contained either one or two colored squares depending on the Set Size, two or four respectively. The second stimuli array also contained one or two squares that were either presented in the same locations occupied in the first array or in different locations.

The stimuli in the first array were rendered in a color drawn randomly from nine possible colors without replacement (red, green, blue, yellow, magenta, cyan, white, black, and orange). Next, an interstimulus interval followed in which only the fixation dot remained on the screen for 350 msec. Then, the second stimuli array followed consisting of either one or two squares to complete the set size to be maintained for that trial (set size two or four, respectively). The stimuli in this second array were either presented in the same locations that were previously occupied or in new locations (same or different locations, respectively).

The colors for the stimuli presented in the second array were determined randomly from the colors remaining and excluded the colors used in the first array. Participants were asked to remember the color and locations of the stimuli over a 900 msec blank delay interval where only the fixation dot remained on the screen. After the delay, a single probe stimulus reappeared in

one of the locations that was previously occupied and was either rendered in the same color as one of the original stimuli presented there or in a different color drawn randomly from the entire set of colors, excluding the color that was actually presented there. This meant that when the locations of the stimuli were the same between the first and second array (different location condition), the set of possible “change” colors excluded both colors presented in the probed location. Participants used a keyboard button press to indicate whether this probe stimulus was the same or not. Participants pressed the “z” or “/” key to indicate whether the color of the probe stimulus was the “same” or “different”, respectively. There were no practice trials given before the formal experiment. Participants were given verbal and written task instructions with the aid of an example trial image similar to Figure 3-1.

Participants completed twenty blocks with each containing 64 trials. Within a block, half of the trials were “no change” trials, and the remaining half were “change” trials where the probe stimulus was rendered in a different color than the one (or any) presented there. Similarly, half of the trials were Set Size 2 such that the first array contained one square and the second array contained an additional square. The remaining half were Set Size 4 with two squares presented in each the first and second array. Within each block, there were also an equal number of same and different location condition trials that determined whether the stimuli in the second array were presented in the same locations as the first array or not. Finally, trials were equally split on whether the probe stimulus was from the first array or the second array.

Participants self-initiated each block by pressing the spacebar key. The experiment session was scheduled to take 3 hr, but the actual duration of the session depended on each participant's pace because they initiated each block and decided when (and if) to take breaks between blocks.

Experiment 3-2 Procedure

The procedure was similar to Experiment 3-1 with the following exceptions. First, an additional condition was included that simultaneously presented two or four squares in the first array, though the data from this condition are not analyzed further. Additionally, gray placeholder circles (diameter = $\sim 2.26^\circ$) were presented in both the first and second arrays so that four items were always presented. For Set Size two, this meant that each array contained three placeholders, while Set Size four trials contained two placeholders in each array. On each trial, four positions were randomly chosen from the set of possible locations. Depending on the set size, positions were assigned to stimuli and placeholders. In same location condition trials, the locations of the gray placeholders (and stimuli) were the same for both arrays. The locations were switched on the different location condition trials, such that the stimuli were placed in the placeholder locations from the first array and the placeholders were placed in the stimuli locations from the first array. Participants completed twenty blocks with each containing 72 trials.

The experiment session was scheduled to take 3.5 hr, but the actual duration of the session depended on each participant's pace because they initiated each block and decided when (and if) to take breaks between blocks.

Experimental Design

Both experiments used a 2×2 within-subject design. The factors were set size (2 or 4) and location condition (same or different locations). Behavioral data (i.e., accuracy) were analyzed using a repeated-measures ANOVA.

Parieto-Occipital Alpha Power Analysis

EEG signal processing was performed in MATLAB. We band-pass filtered the raw EEG data using a filter from the FieldTrip toolbox (`ft_preproc_bandpassfilter.m`; Oostenveld et al., 2011), and then extracted instantaneous power values for the alpha band (8 – 12 Hz) by applying a Hilbert transform (`hilbert.m`) to the filtered data. We calculated alpha power for the parieto-occipital electrodes: P7, P3, Pz, P4, P8, PO7, PO3, PO4, PO8, O1, Oz, O2. For illustrative purposes in the figures, we subtracted the mean baseline (-400 msec to 0 msec) at each time point in the trial for each condition and converted to percent change from baseline.

Multivariate Classification Analysis

For the mvLoad analysis, we used a logistic regression model to classify the number of items in working memory (i.e., WM load) using baselined EEG (Thyer et al., 2022). EEG activity was calculated using a baseline from -500 msec to -100 msec relative to the onset of the stimulus array. The mean baseline amplitude was subtracted from EEG amplitude at each time point in the trial. To improve our signal-to-noise ratio, we randomly selected trials within each condition of interest to create groups of 20 trials and then averaged across the trials in each group (Figure 3-2A). The classification procedure was performed by averaging voltage for each electrode across a 50 msec time window that advanced in 25 msec steps. At each time point, the training data were standardized, and the testing data were standardized using the mean and standard deviation of the training data ('StandardScaler' in scikit-learn, Pedregosa et al., 2011).

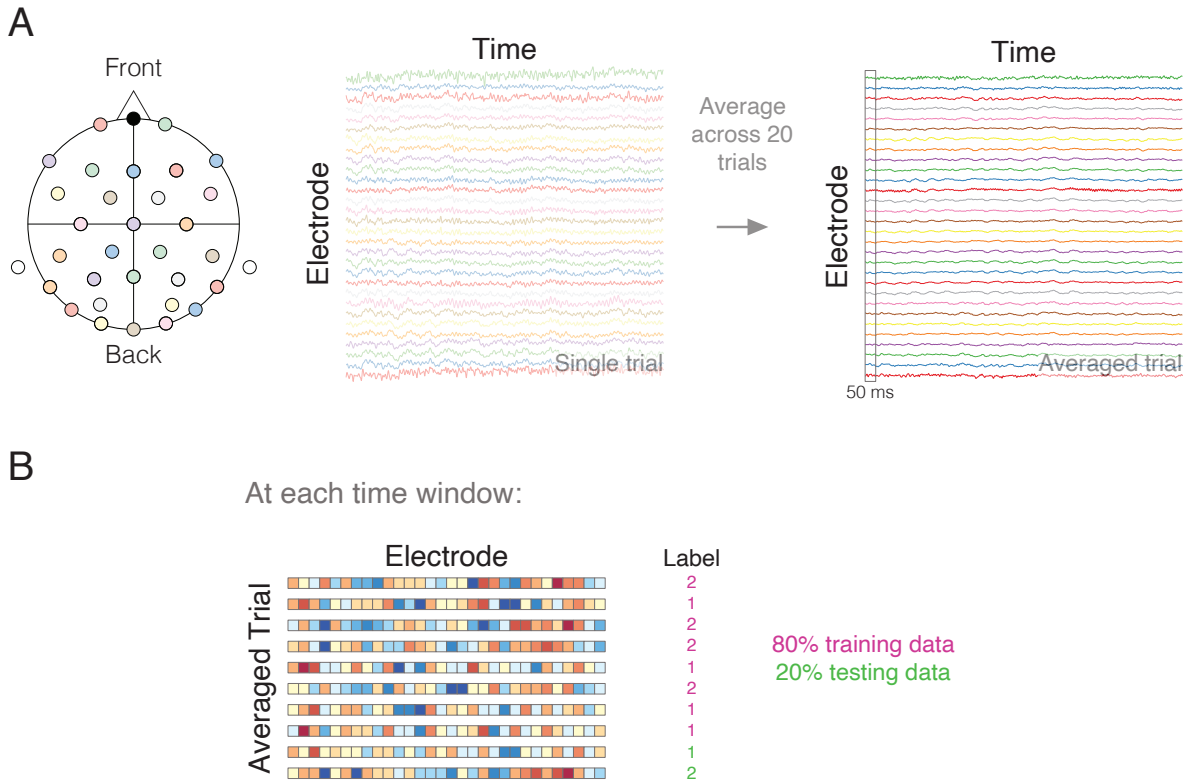


Figure 3-2. mvLoad approach.

Overview of the multivariate classification analysis using voltage. Within each condition of interest, trials were randomly selected to create groups of 20 trials and then averaged (A). The classification procedure was performed by averaging voltage for each electrode across a 50 msec time window. The classifiers were trained to discriminate between our conditions of interest with 80% of the data used in training and the remaining 20% of data used in testing (B).

The classifiers were trained to discriminate between our conditions of interest, and then tested on a held out set of data ('StratifiedShuffleSplit' in scikit-learn, Figure 3-2B). In this cross-validation procedure, the data for any given condition were split so that 80% of the data were used in training and the remaining 20% of data were used in testing. This procedure

preserved the percentage of samples for each condition of interest and was repeated 1000 times with results averaged across all iterations.

Statistical Analysis

Behavioral data were analyzed using a repeated-measures ANOVA. Cluster-based permutation tests were used for alpha power and mvLoad analyses (Maris & Oostenveld, 2007). For alpha power analyses, cluster-based permutation tests were used to test for differences between conditions of interest across time. For mvLoad analyses, cluster-based permutations tests were used in one of two ways. In the first kind of mvLoad analysis, permutations tests were used to test if classification accuracy was significantly above chance with empirical chance accuracy obtained by testing the trained models on randomly shuffled trial labels. In the second variation of mvLoad analysis, permutations tests were used to test for differences in the classifiers' confidence between conditions of interest at each time window. Confidence scores refer to the signed distance of the test sample to the hyperplane in arbitrary units ('decision_function' in scikit-learn).

For cluster-based permutation tests, repeated-measures t tests were calculated at each time point (or window) to assess the difference between conditions of interest, which produced a map of t scores across time. Time points were thresholded according to an a priori defined criterion (p value < .05, two-sided), and adjacent time points with t scores that exceeded this value were grouped together to form a cluster. Clusters were summarized into a single number by summing the t values. At this point, the extent of the cluster became fixed, so the individual time windows were not visible to the next inference step (Sassenhagen & Draschkow, 2019). We next calculated the probability that these values came from a null distribution. Permutation tests were used to establish the probability of our data under the null hypothesis. The number of

permutations was 10,000 or the maximum possible, whichever was lowest. On each iteration and for each time point, it was randomly determined if the first condition was subtracted from the second condition or vice versa for each participant. Then for each iteration, the cluster formation step was repeated. The cluster with the highest sum of t values was identified, and the sum of its t values was stored and became our surrogate-null value for that permutation. After all iterations, the cumulative density of these surrogate-null values was our approximation of the values under the null hypothesis. The p value was then calculated as the percentage of surrogate-null values that the observed data exceeded.

Results

Behavior

For Experiment 3-1, there was a main effect of Set Size, $F(1,9) = 61.90, p < .001, \eta^2 = .87$, and Location Condition, $F(1,9) = 8.80, p = .016, \eta^2 = .49$, on accuracy, such that accuracy was higher for Set Size 2 ($M = 0.98, SD = 0.01$) than Set Size 4 ($M = 0.87, SD = 0.06$) and when different locations were occupied in the second array ($M = 0.93, SD = 0.07$) compared to the same locations ($M = 0.92, SD = 0.07$). There was a significant interaction between Set Size and Location Condition on accuracy, $F(1,9) = 8.17, p = .019, \eta^2 = .48$, such that the benefit of appearing in different locations was greater for Set Size 4 than Set Size 2, perhaps due to a ceiling effect in Set Size 2.

The pattern of behavioral results was replicated in Experiment 3-2. There was a main effect of Set Size, $F(1,19) = 69.47, p < .001, \eta^2 = .79$, and a trending effect of Location Condition, $F(1,19) = 4.09, p = .057, \eta^2 = .18$, such that accuracy was higher for Set Size 2 ($M = 0.95, SD = 0.06$) than Set Size 4 ($M = 0.82, SD = 0.10$) and when different locations were occupied in the second array ($M = 0.89, SD = 0.11$) compared to the same locations ($M = 0.88,$

$SD = 0.10$). Similar to Experiment 3-1, there was a significant interaction between Set Size and Location Condition on accuracy, $F(1,19) = 9.49, p = .006, \eta^2 = .33$, in which the benefit of appearing in different locations was only present for Set Size 4 compared to Set Size 2.

Across both experiments, there was a reliable set size effect, such that performance was higher for Set Size 2 than Set Size 4. For Set Size 4, there was also a small benefit when the stimuli locations occupied in the second array were different than those used in the first array.

Parieto-Occipital Alpha Power

We used nonparametric cluster-based permutations tests to examine the effects of Set Size and Location Condition on alpha power at parieto-occipital electrodes (Maris & Oostenveld, 2007). Based on previous research, we predicted that alpha power would be sensitive to both the number of items held in mind and the number of relevant locations (Diaz et al., 2021; Fukuda et al., 2015). To investigate whether alpha power was sensitive to Set Size, cluster-based permutation tests were performed on data averaged across Location Condition for each Set Size. To detect whether alpha power was sensitive to Location Condition, cluster-based permutation tests were performed separately for each set size (e.g., Set Size 2 Same vs Set Size 2 Different).

In Experiment 3-1, cluster-based permutation tests revealed a cluster sensitive to Set Size that extended from ~130 msec to ~1220 msec (Figure 3-3A). For each Set Size, there was one cluster sensitive to Location Condition (i.e., same or different locations). For Set Size 2, the cluster extended from ~540 msec to the end of the delay (~1550 msec), while the cluster for Set Size 4 extended from ~640 msec to the end of the delay. All permutations tests indicated that the effects of Set Size and Location Condition were significant ($p = .006, p = .011, p = .005$, respectively).

In Experiment 3-2, cluster-based permutation tests revealed two clusters sensitive to Set Size that extended from ~-280 msec to 170 msec and from ~350 msec to ~1270 msec (Figure 3-3B). The permutation tests indicated that the effects of set size were significant ($p = .04$, $p = .007$, respectively). For Set Size 4, there was a cluster sensitive to Location Condition that extended from ~810 msec to ~1380 msec. The permutation test indicated that the effect of Location Condition was significant ($p = .021$).

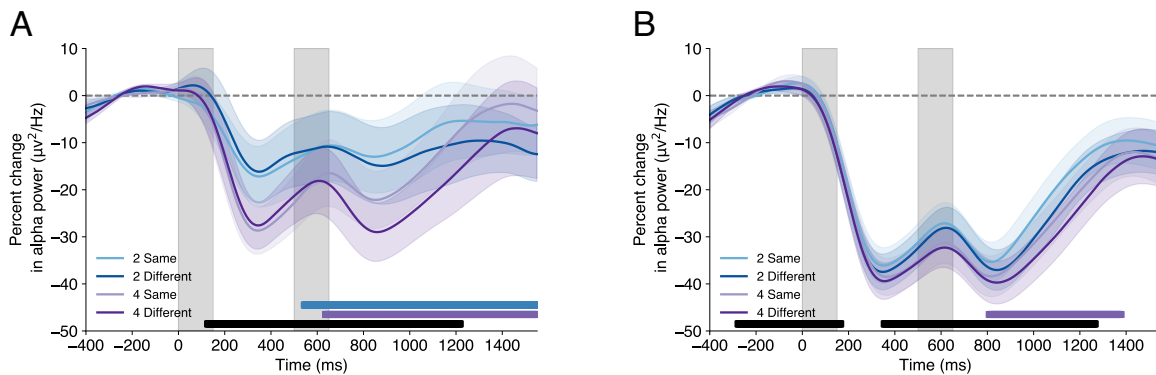


Figure 3-3. Alpha power over time.

Averaged alpha power suppression observed at parieto-occipital electrodes in Experiment 3-1 (A) and 3-2 (B). Shaded regions indicate duration of stimuli arrays. Black bars indicate clusters showing a significant ($p < .05$) Set Size effect from cluster-based permutations tests, while purple and blue bars indicate clusters showing a significant Location Condition effect for Set Size 2 and 4, respectively.

As predicted, alpha power at parieto-occipital electrodes was sensitive to the number of items in mind. For both experiments, there was greater alpha suppression for two items than one item after the first array and for four items than two items after the second array. Alpha power was also sensitive to whether items were added to the same locations or different locations. Specifically, there was greater alpha suppression when items were added to different locations for both Set Size 2 (Experiment 3-1) and Set Size 4 (Experiment 3-1 and 3-2). One possible

interpretation for this finding is that alpha power tracked the shift in spatial attention required in the Different Location trials. However, we propose that alpha power instead tracked the number of locations occupied by the memoranda. This latter interpretation is in line with previous work suggesting that alpha power at posterior electrodes tracks the number of relevant locations when all items are presented simultaneously (Diaz et al., 2021).

Multivariate Analysis of Voltage

Besides parieto-occipital alpha power, the number of items in working memory (i.e., WM load) can be decoded using multivariate analysis of EEG voltage, or mvLoad (Adam et al., 2020; Thyer et al., 2022). However, the extant work with this approach has always relied on spatial position to differentiate each item from the others. Thus, our goal in the present work was to test whether the mvLoad approach is sensitive to the total number of items when only *temporal* rather than spatial position separates one half of the display from the other. We used a logistic regression model to classify WM load using baselined EEG. For each condition of interest, trials were divided into groups of 20 and then averaged with the resulting matrix (electrodes \times time points) subsequently referred to as a trial for ease of reference. The trials (-400 msec to 1550 msec) were divided into 50 msec time windows with a sliding window of 25 msec. Data were averaged within each time window, so that each trial was represented by a 30×80 matrix (electrodes \times time windows). Finally, classification analyses were performed at each time window for each participant.

First, we investigated whether mvLoad could decode the number of items in mind during both Location Conditions (i.e., same or different locations). To this end, we trained and tested classification models to assess whether Set Size 2 Same and Set Size 4 Same trials were reliably and accurately classified, which would suggest load-specific patterns of activity between the

conditions. The same procedure was carried out for Set Size 2 Different and Set Size 4 Different trials.

For Experiment 3-1, we could classify Set Size shortly after the first array in both Location Conditions (Figure 3-4A). For both Location Conditions, cluster-based permutation tests revealed that there was a cluster sensitive to Set Size, extending from the time window beginning at ~80 msec to the end of the delay. The permutation test indicated above chance classification ($p = .005$) for both conditions. The results from Experiment 3-2 followed a similar pattern (Figure 3-4B). Specifically, there was a cluster sensitive to Set Size, extending from ~180 msec to the end of the delay, for trials in which the additional items occupied the same locations used in the first array (i.e., same locations). Similarly, there was a cluster sensitive to Set Size, extending from ~270 msec to ~1300 msec, for trials in which the additional items occupied novel locations (i.e., different locations). For both of these clusters, the permutation tests indicated above chance classification accuracy ($p < .001$).

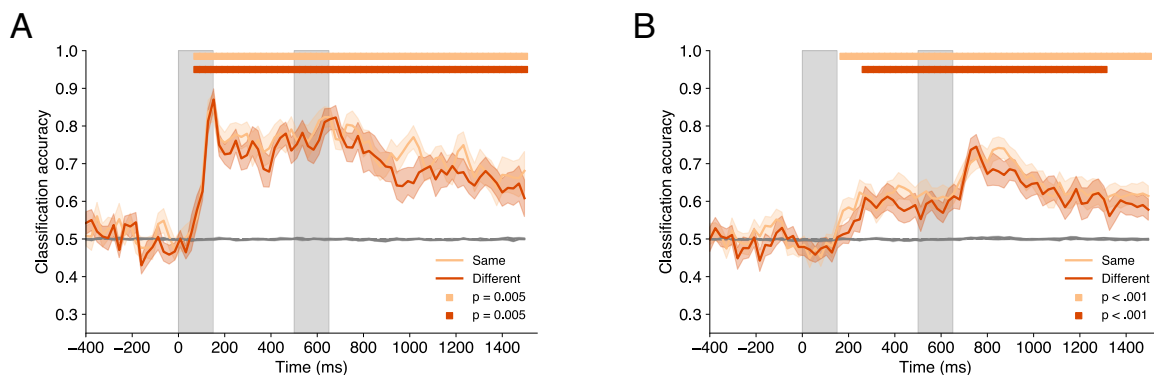


Figure 3-4. Classification accuracy over time.

Classification accuracy over time for Set Size for Same Locations (light orange) and Different Locations (dark orange) in Experiment 3-1 (A) and 3-2 (B). Shaded regions indicate duration of stimuli arrays. Light and dark orange bars indicate clusters showing significant ($p < .05$) above-chance classification from cluster-based permutation tests.

To summarize, above chance classification accuracy after the first array and during the ISI indicated that the patterns of activity were distinct between one item and two items. After the second array when the additional items were presented, above chance classification accuracy indicated that the patterns of activity were distinct between two items and four items. This distinction was evident in both Location Conditions. Multivariate analysis of EEG voltage tracked the number of items regardless of whether the additional items were added to novel positions or the same positions used in the first array, suggesting that items can be individuated on the basis of temporal position alone.

Instead of tracking the number of items, it is possible that mvLoad was instead sensitive to the number of locations in memory. For the Same condition, the patterns of activity that were characterized during the delay might instead correspond to one location and two locations. For the Different condition, the multivariate analysis might be picking up on distinct patterns between two locations and four locations. This alternative account can be addressed by investigating whether the EEG signature for a given set size is consistent across the same and different location conditions. Accordingly, we carried out a direct comparison of EEG voltage topography between Location Conditions. If mvLoad is tracking the number of items in mind, then the pattern of activity should look similar for a given Set Size regardless of whether the additional items in the second array occupied the same locations from the first array or novel locations. In other words, a load-sensitive signal should not be affected by the number of occupied locations.

To assess the similarities in voltage topography between Location Conditions, we asked if there was a signal tracking the number of individuated items regardless of the number of attended positions. To this end, we first characterized voltage topography for two items and four

items as in the above analyses. Classifiers were trained on either data from Set Size 2 Different and Set Size 4 Different or data from Set Size 2 Same and Set Size 4 Same.

As in the previous analyses, classifiers were then tested on held out data from the conditions that were used for training. The previous analysis demonstrated that mvLoad could reliably and accurately distinguish between these two conditions. Unlike the previous analyses, classifiers were also tested using data from a third condition that was not used in training. In the analysis that used training data from the Different conditions, the third condition that was tested was Set Size 4 Same, which shared the same number of items as Set Size 4 Different but occupied the same number of positions as Set Size 2 Different. In the analysis that used training data from the Same conditions, the third condition that was tested was Set Size 2 Different, which shared the same number of items as Set Size 2 Same but occupied the same number of positions as Set Size 4 Same. The key question was whether the mvLoad classification would be determined by the number of individuated items stored, or by the total number of relevant positions in the display. Incidentally, training the classifiers within Location Conditions controlled for the presence (or absence) of shifts of attention, which might also be reflected in voltage topography.

For all three test conditions, we obtained a measure of the classifiers' confidence for each decision made about the test sample ('decision_function' in scikit-learn). This confidence score reflects the signed distance of the test sample to the hyperplane in arbitrary units. Positive scores indicated that the trial was classified as Set Size 4 (Same or Different depending on the analysis) with higher scores reflecting stronger evidence for this decision. Meanwhile, negative scores indicated a Set Size 2 (Same or Different) classification with lower scores reflecting stronger evidence. Cluster-based permutation tests were used to test for differences in confidence scores

between conditions (see *Statistical Analysis*). Our time window of interest was the delay period (starting at 650 msec after the initial sample onset) given that our comparisons of interest relied on the total set size.

The results for the analysis that used training data from the Different condition are discussed first. As demonstrated in the previous analysis, multivariate analysis could distinguish the patterns of activity between Set Size 2 Different and Set Size 4 Different. The confidence scores in these classification decisions reveal that there was strong evidence that supported above chance classification accuracy for both experiments (Figure 3-5). Stronger evidence for Set Size 4 is plotted in the positive direction, while stronger evidence for Set Size 2 is plotted in the negative direction. Shortly after stimulus onset, there was strong evidence that supported the classifier's decisions to classify Set Size 2 Different and Set Size 4 Different trials accordingly.

When the same classifiers were tested on Set Size 4 Same trials, the trials were consistently classified in the same category as Set Size 4 Different trials throughout the entire trial. There were no clusters that were sensitive to differences in evidence between Set Size 4 Same and Set Size 4 Different for both Experiment 3-1 and 3-2, suggesting that the patterns of voltage activity were equivalent for four items in two locations and four items in four locations (delay).

On the other hand, cluster-based permutation tests revealed a cluster that was sensitive to differences in evidence between Set Size 4 Same and Set Size 2 Different. For Experiment 3-1, the cluster extended from the time window beginning at ~80 msec to the end of the delay ($p = .005$). For Experiment 3-2, the cluster extended from the time window beginning at ~250 msec to the end of the delay ($p < .001$). The patterns of activity between two items in two locations and four items in two locations (delay) were discernibly and reliably different.

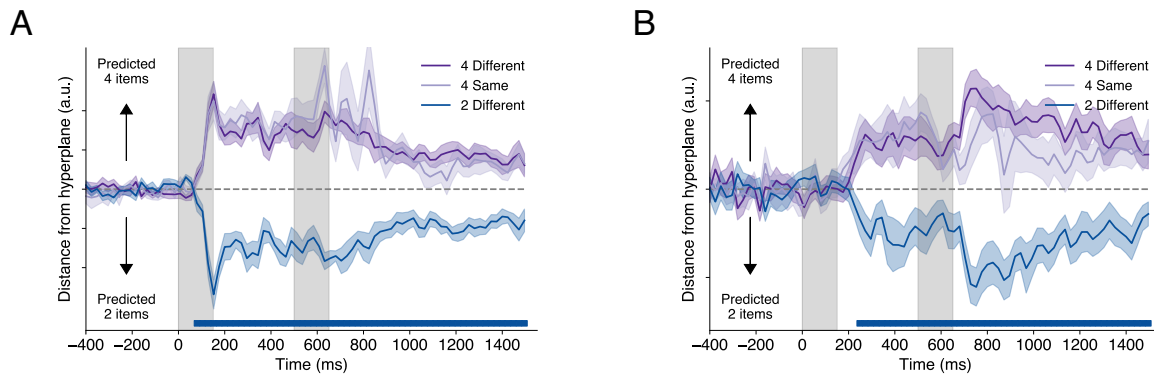


Figure 3-5. Classifier evidence over time when novel positions were occupied.

Distance from the classification hyperplane for Set Size 2 Different, Set Size 4 Different, and Set Size 4 Same trials across time for Experiment 3-1 (A) and 3-2 (B). Classification is trained on Set Size 2 and 4 Different trials and tested on all three conditions. Dashed gray line indicates hyperplane. Blue bar indicates clusters showing a significant ($p < .05$) difference between Set Size 2 Different and Set Size 4 Same from cluster-based permutation tests.

After the first array, there was strong evidence to classify the two items in the Set Size 4 Same condition as two items (Set Size 4 Different) rather than one item (Set Size 2 Different) as expected. Remarkably, after the second array, there was strong evidence to classify four items in two locations (Set Size 4 Same) in the same category as four items in four locations (Set Size 4 Different) rather than two items in two locations (Set Size 2 Different). Moreover, the classifiers were just as “confident” in making these decisions as they were in making the decision that Set Size 4 Different trials were in fact Set Size 4 Different trials. This finding suggests that mvLoad can distinguish between the number of items in mind regardless of the number of positions occupied by the items.

The results for the analysis that used training data from the Same condition largely mirrored this finding. As demonstrated before, multivariate analysis could distinguish the patterns of activity between Set Size 2 Same and Set Size 4 Same. Again, the confidence scores

in these classification decisions reveal that there was strong evidence that supported above chance classification accuracy (Figure 3-6). Shortly after stimulus onset, there was strong evidence that supported the classifier's decisions to classify Set Size 2 Same and Set Size 4 Same trials accordingly.

When the same classifiers were tested on Set Size 2 Different trials, the trials were more likely to be classified as Set Size 2 Same trials throughout the entire trial in Experiment 3-1 (Figure 3-6A). There were no clusters that were sensitive to differences in evidence between Set Size 2 Different and Set Size 2 Same for Experiment 3-1, suggesting that the patterns of voltage activity were equivalent for two items in two locations and two items in one location (delay). In Experiment 3-2, Set Size 2 Different trials were more likely to be classified as Set Size 2 Same trials until shortly after the second array (Figure 3-6B). Cluster-based permutation tests revealed a cluster that was sensitive to differences in evidence between Set Size 2 Different and Set Size 2 Same in Experiment 3-2. The cluster extended from the time window beginning at ~940 msec to ~1350 msec ($p = .001$).

On the other hand, cluster-based permutation tests revealed a more stable cluster that was sensitive to differences in evidence between Set Size 4 Same and Set Size 2 Different. For Experiment 3-1, the cluster extended from the time window beginning at ~80 msec to the end of the delay ($p = .003$). For Experiment 3-2, the cluster extended from the time window beginning at ~150 msec to ~1110 msec ($p < .001$), suggesting that the patterns of activity between four items in two locations and two items in two locations (delay) were reliably different.

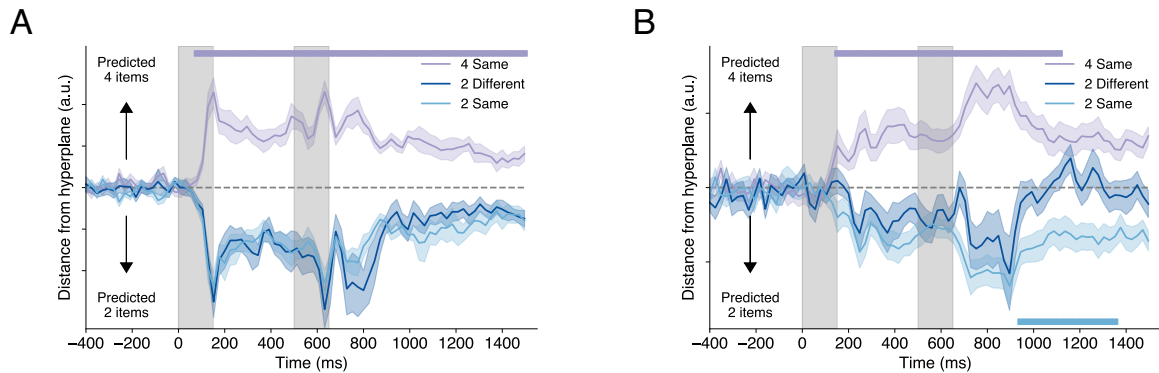


Figure 3-6. Classifier evidence over time when the same positions were occupied.

Distance from the classification hyperplane for Set Size 2 Same, Set Size 4 Same, and Set Size 2 Different trials across time for Experiment 3-1 (A) and 3-2 (B). Classification is trained on Set Size 2 and 4 Same trials and tested on all three conditions. Dashed gray line indicates hyperplane. Purple bar indicates clusters showing a significant ($p < .05$) difference between Set Size 4 Same and Set Size 2 Different from cluster-based permutation tests, while blue bar indicates cluster showing a significant difference between Set Size 2 Same and Set Size 2 Different.

After the first array, there was strong evidence to classify the single item in the Set Size 2 Different condition as one item (Set Size 2 Same) rather than two items (Set Size 4 Same) as expected. After the second array in Experiment 3-1, there was strong evidence to classify two items in two locations (Set Size 2 Different) as two items in one location (Set Size 2 Same) rather than four items in two locations (Set Size 4 Same). The lack of any reliable difference between Set Size 2 Same and Set Size 2 Different suggests that the pattern of voltage was equivalent between the conditions. In Experiment 3-2, the evidence to classify two items in two locations as two items in one location was also present but short-lived. Here, evidence for classifying Set Size 2 Same trials as either training condition weakened over time with test cases falling closer to the hyperplane. By the end of the delay, Set Size 2 Different evidence fell to chance levels and did not reliably differ from either Set Size 2 Same or Set Size 4 Same. These findings largely corroborate the results found using training data from the Different conditions.

Specifically, multivariate analysis of EEG voltage can distinguish between the number of items in mind regardless of the number of positions.

Discussion

In summary, multivariate analysis of EEG voltage, or mvLoad, enabled the robust decoding of the total number of items in mind across sequential arrays even when items were distinguished from each other only on the basis of time (i.e., shared positions). The pattern of voltage activity across electrodes appeared similar for any given number of items regardless of the number of occupied locations. In other words, mvLoad was primarily sensitive to the number of items in mind rather than the number of relevant locations. On the contrary, alpha power activity (8 – 12 Hz) at posterior electrodes tracked the number of locations in mind, thereby ensuring that participants did, in fact, maintain the relevant locations.

In addition to this load-sensitive signal (Adam et al., 2020; Thyer et al., 2022), there have been other demonstrations of neural signals tracking the number of items in working memory using EEG (Diaz et al., 2021; Fukuda et al., 2015; Vogel & Machizawa, 2004). For instance, the contralateral delay activity (CDA), a contralateral and sustained negative voltage deflection at posterior electrodes, shows reliable decreases in amplitude as items are added into working memory until capacity is reached. The CDA tracked the total number of items in mind across sequential arrays regardless of whether items occupied the same positions or novel positions, which is consistent with the current findings using mvLoad (Vogel & Machizawa, 2004). The current work further demonstrates that the patterns of voltage activity across electrodes were equivalent for any given set size regardless of the number of occupied locations.

The similarities between the CDA and mvLoad hint at the possibility that the signals reflect the same underlying neural mechanism given that both signals are derived from EEG

voltage. In fact, there is another related signal that might be contributing to the current multivariate analysis. Namely, a sustained negative voltage deflection at posterior electrodes, sometimes referred to as the negative slow wave, also tracks the number of items in mind (Diaz et al., 2021; Fukuda et al., 2015). There are, however, critical differences in how these signals use EEG voltage information. First, the CDA is a lateralized signal that makes use of the contralateral organization of the human visual system. This means that detecting the CDA requires the use of lateralized arrays in which participants are cued to attend and remember items in one hemifield and ignore identical visual stimuli presented to the other hemifield. Meanwhile, the negative slow wave and mvLoad are not lateralized components and can therefore be used with whole-field designs. Second, the CDA and the negative slow wave are detected at posterior electrodes, while mvLoad uses EEG voltage across all electrodes. Finally, the CDA and the negative slow wave are both univariate signals obtained by averaging across hundreds of trials and across electrodes of interest, while the current multivariate analysis of EEG voltage works at the level of a single trial (Adam et al., 2020). It is likely that these load-sensitive signals share underlying information, although further research is required to determine the extent of their similarities and to clearly delineate any differences.

The current work provides additional support for the notion that mvLoad might reflect the operation of a “pointer system” that supports object individuation (Thyer et al., 2022). The current findings are the first demonstration that this signal is sensitive to the individuation of items across time and that it can use time alone as a basis for individuation. Previous ideas that a “pointer system” supports object individuation have made clear the requirement that the system must track items continuously in both space and time. For instance, Kahneman et al. (1992) introduced the concept of an “object file” that represents a temporary episodic memory for

individual visual objects and accumulates new information about the object as it changes through time and space. Further, Pylyshyn (2009) proposed that objects in a scene are assigned individual indices that continuously track the objects' positions in time and space, as well as their positions relative to one another.

The current work also highlights a dissociation between item-based and spatial-based EEG signals. Specifically, we demonstrate that while multivariate analysis of EEG voltage tracked the number of items in mind, alpha power at posterior electrodes tracked the number of relevant locations with alpha power suppression increasing as more locations were occupied. This finding is consistent with previous work demonstrating that alpha power at posterior electrodes is sensitive to the number of locations in mind, while voltage at posterior electrodes is mostly sensitive to the number of individuated items (Diaz et al., 2021). It is also consistent with previous demonstrations that sustained potentials and alpha-band oscillations reflect distinct aspects of storage in working memory (Bae & Luck, 2018; de Vries et al., 2017; Hakim et al., 2019; Wang et al., 2019, 2020).

GENERAL DISCUSSION

The successful interplay between attention and working memory enables routine behavior from holding a conversation to grocery shopping. Accordingly, there has been substantial interest and progress in understanding the neural mechanisms underlying working memory and attention. One prominent idea is that the active maintenance of information relies on several complementary but distinct cognitive processes. In line with this view, I provide evidence that neural activity during simple working memory tasks can be dissected to reveal a wealth of information about what participants are holding in mind. This dissertation explores multiple electroencephalography (EEG) signals that track distinct aspects of working memory, including its spatial content and the number of individuated items in mind.

In Chapter 1, I focused on the well-known relationship between oscillatory activity in the alpha frequency band (8 – 12 Hz) and spatial attention. Previous work has shown that alpha power topography, or the pattern of alpha power across the scalp, systematically varies as participants attend to positions around fixation (e.g., Rihs et al., 2007). Alpha power topography can be used to decode the spatial position held in mind using inverted encoding models, or IEMs, (e.g., Foster et al., 2016). Here, I extended this finding by demonstrating that alpha power topography is a generalizable human signature of spatial attention. To do this, I built a database comprised of EEG recordings from 146 subjects completing simple working memory tasks. Critically, the tasks required participants to attend to stimuli that could appear in various spatial positions around fixation. Using IEMs, I demonstrated that the alpha power topography of a group of subjects could be used to decode the spatial position in the mind of a novel subject. This generalizability persisted across different tasks, research sites, EEG systems, and research

groups. Alpha power topography, thus, is a reliable indicator of spatial information in working memory.

In Chapter 2, I further investigated the role of alpha power in the maintenance of spatial information, while also considering how it differed from other traditional working memory signals, namely event-related potentials. I examined a sustained negative voltage deflection, or negative slow wave, at posterior electrodes purported to track the number of items in mind (Fukuda et al., 2015). First, I replicated the previous demonstration that this load-sensitive signal tracked the number of items in working memory (WM load) when participants maintained both color and spatial information about the items. Next, I established that this negative slow wave also tracked WM load even when participants maintained only spatial information. Alpha power measured at those same electrodes similarly tracked WM load in both conditions. Importantly, each item in these displays was associated with a single spatial position, so it was unclear whether the negative slow wave and alpha power were tracking the number of items in mind or the allocated spatial attention. Motivated by the strong relationship between alpha power and spatial attention, I investigated the possibility that alpha power at posterior electrodes is instead primarily sensitive to the attended spatial positions rather than WM load.

To this end, I disentangled the number of individuated items from spatial attention using perceptual grouping cues. When aligned, these collinearity cues made two stimuli appear as one individuated item. When misaligned, the stimuli were perceived as two distinct items. In both conditions, the same number of spatial positions were occupied. Alpha power at posterior electrodes tracked the number of occupied spatial positions, but it was not sensitive to whether the stimuli were aligned or not. Conversely, the negative slow wave was sensitive to the grouping manipulation and tracked the number of perceived items. Both alpha power and the

negative slow wave tracked WM load when items were associated with a single spatial position. Yet, the perceptual grouping results suggest that these EEG signals reflect distinct aspects of maintenance with one closely tracking allocated spatial attention and the other primarily sensitive to WM load. For load-sensitive signals, the question then arises as to whether the patterns of neural activity were comparable between one item comprised of one stimulus and one *perceived* item comprised of two stimuli.

In Chapter 3, I investigated this open question and further explored the dissociation between alpha power at posterior electrodes and load-sensitive signals. Here, I focused on a load-sensitive signal that would enable the comparison of voltage activity patterns across conditions, namely multivariate analysis of EEG voltage across the scalp (Adam et al., 2020; Thyer et al., 2022). Participants were asked to remember items across two sequential arrays with the items in the second array either occupying the same positions as in the first array or novel positions. Accordingly, the number of spatial positions was manipulated, while the number of items was held constant; this was a direct complement to the perceptual grouping cues where the number of individuated items varied while spatial attention was held constant. Consistent with the perceptual grouping findings, alpha power at posterior electrodes was sensitive to the number of spatial positions in mind, while multivariate analysis of EEG voltage primarily tracked WM load. Notably, the pattern of voltage activity for any given WM load was stable regardless of variations in spatial attention. This dissociation between alpha power at posterior electrodes and multivariate analysis of EEG voltage serves as another demonstration that active maintenance is supported by distinct subcomponents.

Multivariate analysis of voltage is just one of many load-sensitive signals that can be detected during working memory maintenance. Here, I demonstrated that voltage topography for

any given WM load remains stable regardless of variations in spatial attention. This finding supports the proposal that load-sensitive signals might be indexing the product of a “pointer” system that enables continuous tracking of items in space and time rather than tracking the item representations themselves (Thyer et al., 2022). First, voltage topography tracked WM load even when items were individuated on the basis of time alone, which is expected behavior from a “pointer” purported to track items in space and *time*. Second, voltage topography for any given WM load was equivalent regardless of the spatial information maintained, which more closely describes a nonspecific index that points to the item representation rather than the item representation itself.

This dissertation highlights several neural signals that support the active maintenance of information. These signals can be broadly classified along many dimensions. For instance, the signals discussed reflected either oscillatory (i.e., alpha power) or voltage activity. These signals considered either neural activity across the scalp (i.e., topography) or from a subset of posterior electrodes. Finally, the analytic tools differed and relied on either multivariate or univariate analyses. Altogether, these concurrent signals tracked distinct aspects of active maintenance, including the spatial content, the number of individuated items, and the spatial attention allocated to those items. The coordination of these subcomponents ensures that information admitted into working memory is effectively maintained in the service of ongoing behavior.

Future work will be necessary to fully understand the mechanisms underlying attention and working memory, as well as their relationship with higher order cognitive processes. The current work advances our understanding by demonstrating the multifaceted nature of neural activity during working memory. These incremental steps bring us closer to fully appreciating the intricate coordination that enables human behavior from the mundane to the extraordinary.

REFERENCES

- Adam, K. C. S., Vogel, E. K., & Awh, E. (2020). Multivariate analysis reveals a generalizable human electrophysiological signature of working memory load. *Psychophysiology*, *57*(12), e13691. <https://doi.org/10.1111/psyp.13691>
- Armstrong, B. C., Ruiz-Blondet, M. V., Khalifian, N., Kurtz, K. J., Jin, Z., & Laszlo, S. (2015). Brainprint: Assessing the uniqueness, collectability, and permanence of a novel method for ERP biometrics. *Neurocomputing*, *166*, 59–67. <https://doi.org/10.1016/j.neucom.2015.04.025>
- Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, *18*(7), 622–628. <https://doi.org/10.1111/j.1467-9280.2007.01949.x>
- Bae, G.-Y., & Luck, S. J. (2018). Dissociable Decoding of Spatial Attention and Working Memory from EEG Oscillations and Sustained Potentials. *The Journal of Neuroscience*, *38*(2), 409–422. <https://doi.org/10.1523/JNEUROSCI.2860-17.2017>
- Brady, T. F., & Tenenbaum, J. B. (2013). A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, *120*(1), 85–109. <https://doi.org/10.1037/a0030779>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *Journal of Neuroscience*, *29*(44), 13992–14003. <https://doi.org/10.1523/JNEUROSCI.3577-09.2009>
- Busch, N. A., & Herrmann, C. S. (2003). Object-load and feature-load modulate EEG in a short-term memory task. *NeuroReport*, *14*(13), 1721–1724. <https://doi.org/10.1097/00001756-200309150-00013>
- Carrasco, M., & McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences*, *98*(9), 5363–5367. <https://doi.org/10.1073/pnas.081074098>

- Chen, Y., Atnafu, A. D., Schlattner, I., Weldtsadik, W. T., Roh, M.-C., Kim, H. J., Lee, S.-W., Blankertz, B., & Fazli, S. (2016). A High-Security EEG-Based Login System with RSVP Stimuli and Dry Electrodes. *IEEE Transactions on Information Forensics and Security*, *11*(12), 2635–2647. <https://doi.org/10.1109/TIFS.2016.2577551>
- Christophel, T. B., Hebart, M. N., & Haynes, J.-D. (2012). Decoding the Contents of Visual Short-Term Memory from Human Visual and Parietal Cortex. *The Journal of Neuroscience*, *32*(38), 12983–12989. <https://doi.org/10.1523/JNEUROSCI.0184-12.2012>
- Conway, A. R. A., Kane, M. J., & Engle, R. W. (2003). Working memory capacity and its relation to general intelligence. *Trends in Cognitive Sciences*, *7*(12), 547–552. <https://doi.org/10.1016/j.tics.2003.10.005>
- Courtney, S. M., Ungerleider, L. G., Keil, K., & Haxby, J. V. (1997). Transient and sustained activity in a distributed neural system for human working memory. *Nature*, *386*(6625), 608–611. <https://doi.org/10.1038/386608a0>
- de Vries, I. E. J., van Driel, J., & Olivers, C. N. L. (2017). Posterior α EEG Dynamics Dissociate Current from Future Goals in Working Memory-Guided Visual Search. *The Journal of Neuroscience*, *37*(6), 1591–1603. <https://doi.org/10.1523/JNEUROSCI.2945-16.2016>
- 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. <https://doi.org/10.1016/j.jneumeth.2003.10.009>
- Diaz, G. K., Vogel, E. K., & Awh, E. (2021). *Perceptual Grouping Reveals Distinct Roles for Sustained Slow Wave Activity and Alpha Oscillations in Working Memory*. *33*(7), 11.
- Eickhoff, S., Nichols, T. E., Van Horn, J. D., & Turner, J. A. (2016). Sharing the wealth: Neuroimaging data repositories. *NeuroImage*, *124*(Pt B), 1065–1068. <https://doi.org/10.1016/j.neuroimage.2015.10.079>
- Epstein, J. N., Conners, C. K., Erhardt, D., March, J. S., & Swanson, J. M. (1997). Asymmetrical hemispheric control of visual-spatial attention in adults with attention deficit hyperactivity disorder. *Neuropsychology*, *11*(4), 467–473. <https://doi.org/10.1037/0894-4105.11.4.467>

- Ester, E. F., Anderson, D. E., Serences, J. T., & Awh, E. (2013). A neural measure of precision in visual working memory. *Journal of Cognitive Neuroscience*, *25*(5), 754–761.
- 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.e6. <https://doi.org/10.1016/j.cub.2017.09.031>
- Foster, J. J., Sutterer, D. W., Serences, J. T., Vogel, E. K., & Awh, E. (2016). The topography of alpha-band activity tracks the content of spatial working memory. *Journal of Neurophysiology*, *115*(1), 168–177. <https://doi.org/10.1152/jn.00860.2015>
- Fukuda, K., Mance, I., & Vogel, E. K. (2015). α Power Modulation and Event-Related Slow Wave Provide Dissociable Correlates of Visual Working Memory. *The Journal of Neuroscience*, *35*(41), 14009–14016. <https://doi.org/10.1523/JNEUROSCI.5003-14.2015>
- Fukuda, K., Vogel, E., Mayr, U., & Awh, E. (2010). Quantity, not quality: The relationship between fluid intelligence and working memory capacity. *Psychonomic Bulletin & Review*, *17*(5), 673–679. <https://doi.org/10.3758/17.5.673>
- Fuster, J. M., & Alexander, G. E. (1971). Neuron Activity Related to Short-Term Memory. *Science*, *173*(3997), 652–654.
- Fuster, J. M., & Jervey, J. P. (1982). Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. *Journal of Neuroscience*, *2*(3), 361–375. <https://doi.org/10.1523/JNEUROSCI.02-03-00361.1982>
- Gao, Z., Gao, Q., Tang, N., Shui, R., & Shen, M. (2016). Organization principles in visual working memory: Evidence from sequential stimulus display. *Cognition*, *146*, 277–288. <https://doi.org/10.1016/j.cognition.2015.10.005>
- Gao, Z., Xu, X., Chen, Z., Yin, J., Shen, M., & Shui, R. (2011). Contralateral delay activity tracks object identity information in visual short term memory. *Brain Research*, *1406*, 30–42. <https://doi.org/10.1016/j.brainres.2011.06.049>
- Goldberger, A. L., Amaral, L. A. N., Glass, L., Hausdorff, J. M., Ivanov, P. Ch., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., & Stanley, H. E. (2000). PhysioBank,

PhysioToolkit, and PhysioNet. *Circulation*, 101(23), e215–e220.
<https://doi.org/10.1161/01.CIR.101.23.e215>

Günseli, E., Fahrenfort, J. J., van Moorselaar, D., Daoultzis, K. C., Meeter, M., & Olivers, C. N. L. (2019). EEG dynamics reveal a dissociation between storage and selective attention within working memory. *Scientific Reports*, 9(1). <https://doi.org/10.1038/s41598-019-49577-0>

Hakim, N., Adam, K. C. S., Günseli, E., Awh, E., & Vogel, E. K. (2019). Dissecting the Neural Focus of Attention Reveals Distinct Processes for Spatial Attention and Object-Based Storage in Visual Working Memory. *Psychological Science*, 30(4), 526–540.
<https://doi.org/10.1177/0956797619830384>

Hakim, N., Awh, E., & Vogel, E. K. (2020). Manifold visual working memory. In R. Logie, V. Camos, & N. Cowan (Eds.), *Working Memory: The State of the Science*. Oxford University Press.

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, 458(7238), 632–635. <https://doi.org/10.1038/nature07832>

Ikkai, A., McCollough, A. W., & Vogel, E. K. (2010). Contralateral delay activity provides a neural measure of the number of representations in visual working memory. *Journal of Neurophysiology*, 103(4), 1963–1968. <https://doi.org/10.1152/jn.00978.2009>

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

Jerde, T. A., Merriam, E. P., Riggall, A. C., Hedges, J. H., & Curtis, C. E. (2012). Prioritized Maps of Space in Human Frontoparietal Cortex. *Journal of Neuroscience*, 32(48), 17382–17390. <https://doi.org/10.1523/JNEUROSCI.3810-12.2012>

Jiang, Y., Chun, M. M., & Olson, I. R. (2004). Perceptual grouping in change detection. *Perception & Psychophysics*, 66(3), 446–453. <https://doi.org/10.3758/BF03194892>

- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24(2), 175–219. [https://doi.org/10.1016/0010-0285\(92\)90007-O](https://doi.org/10.1016/0010-0285(92)90007-O)
- Karatekin, C., & Asarnow, R. F. (1998). Working memory in childhood-onset schizophrenia and attention-deficit/hyperactivity disorder. *Psychiatry Research*, 80(2), 165–176. [https://doi.org/10.1016/S0165-1781\(98\)00061-4](https://doi.org/10.1016/S0165-1781(98)00061-4)
- Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in Alpha Oscillatory Power Reflect an Active Retinotopic Mechanism for Distracter Suppression During Sustained Visuospatial Attention. *Journal of Neurophysiology*, 95(6), 3844–3851. <https://doi.org/10.1152/jn.01234.2005>
- King, G. (2007). An Introduction to the Dataverse Network as an Infrastructure for Data Sharing. *Sociological Methods & Research*, 36(2), 173–199. <https://doi.org/10.1177/0049124107306660>
- Kristjánsson, Á. (2006). Surface Assignment Modulates Object Formation for Visual Short-Term Memory. *Perception*, 35(7), 865–881. <https://doi.org/10.1068/p5526>
- Kuntsi, J., Oosterlaan, J., & Stevenson, J. (2001). Psychological Mechanisms in Hyperactivity: I Response Inhibition Deficit, Working Memory Impairment, Delay Aversion, or Something Else? *Journal of Child Psychology and Psychiatry*, 42(2), 199–210. <https://doi.org/10.1111/1469-7610.00711>
- Lee, M.-H., Kwon, O.-Y., Kim, Y.-J., Kim, H.-K., Lee, Y.-E., Williamson, J., Fazli, S., & Lee, S.-W. (2019). EEG dataset and OpenBMI toolbox for three BCI paradigms: An investigation into BCI illiteracy. *GigaScience*, 8(5), giz002. <https://doi.org/10.1093/gigascience/giz002>
- 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. https://doi.org/10.1162/jocn_a_00140
- Luria, R., Balaban, H., Awh, E., & Vogel, E. K. (2016). The contralateral delay activity as a neural measure of visual working memory. *Neuroscience and Biobehavioral Reviews*, 62, 100–108. <https://doi.org/10.1016/j.neubiorev.2016.01.003>

- Luria, R., & Vogel, E. K. (2014). Come Together, Right Now: Dynamic Overwriting of an Object's History through Common Fate. *Journal of Cognitive Neuroscience*, 26(8), 1819–1828. https://doi.org/10.1162/jocn_a_00584
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190. <https://doi.org/10.1016/j.jneumeth.2007.03.024>
- Markiewicz, C. J., Gorgolewski, K. J., Feingold, F., Blair, R., Halchenko, Y. O., Miller, E., Hardcastle, N., Wexler, J., Esteban, O., Goncavles, M., Jwa, A., & Poldrack, R. (2021). The OpenNeuro resource for sharing of neuroscience data. *ELife*, 10, e71774. <https://doi.org/10.7554/eLife.71774>
- Mate, J., & Baqués, J. (2009). Short article: Visual similarity at encoding and retrieval in an item recognition task. *Quarterly Journal of Experimental Psychology*, 62(7), 1277–1284. <https://doi.org/10.1080/17470210802680769>
- Mathewson, K., Hashemi, A., Sheng, B., Sekuler, A., Bennett, P., & Schmidt, L. (2015). Regional electroencephalogram (EEG) alpha power and asymmetry in older adults: A study of short-term test–retest reliability. *Frontiers in Aging Neuroscience*, 7. <https://www.frontiersin.org/article/10.3389/fnagi.2015.00177>
- Mazaheri, A., & Jensen, O. (2008). Asymmetric Amplitude Modulations of Brain Oscillations Generate Slow Evoked Responses. *The Journal of Neuroscience*, 28(31), 7781–7787. <https://doi.org/10.1523/JNEUROSCI.1631-08.2008>
- McDonald, S., Bennett, K. M. B., Chambers, H., & Castiello, U. (1999). Covert orienting and focusing of attention in children with attention deficit hyperactivity disorder. *Neuropsychologia*, 37(3), 345–356. [https://doi.org/10.1016/S0028-3932\(98\)00078-5](https://doi.org/10.1016/S0028-3932(98)00078-5)
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. *The Journal of Neuroscience*, 16(16), 5154–5167. <https://doi.org/10.1523/JNEUROSCI.16-16-05154.1996>
- Morey, C. C. (2019). Perceptual grouping boosts visual working memory capacity and reduces effort during retention. *British Journal of Psychology*, 110(2), 306–327. <https://doi.org/10.1111/bjop.12355>

- Morey, C. C., Cong, Y., Zheng, Y., Price, M., & Morey, R. D. (2015). The color-sharing bonus: Roles of perceptual organization and attentive processes in visual working memory. *Archives of Scientific Psychology*, 3(1), 18–29. <https://doi.org/10.1037/arc0000014>
- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific Utopia: II. Restructuring Incentives and Practices to Promote Truth Over Publishability. *Perspectives on Psychological Science*, 7(6), 615–631. <https://doi.org/10.1177/1745691612459058>
- Obeid, I., & Picone, J. (2016). The Temple University Hospital EEG Data Corpus. *Frontiers in Neuroscience*, 10. <https://www.frontiersin.org/article/10.3389/fnins.2016.00196>
- Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience*, 2011, 1–9. <https://doi.org/10.1155/2011/156869>
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442. <https://doi.org/10.1163/156856897X00366>
- Pernet, C. R., Appelhoff, S., Gorgolewski, K. J., Flandin, G., Phillips, C., Delorme, A., & Oostenveld, R. (2019). EEG-BIDS, an extension to the brain imaging data structure for electroencephalography. *Scientific Data*, 6(1), 103. <https://doi.org/10.1038/s41597-019-0104-8>
- Pessoa, L., Gutierrez, E., Bandettini, P. A., & Ungerleider, L. G. (2002). Neural Correlates of Visual Working Memory: FMRI Amplitude Predicts Task Performance. *Neuron*, 35(5), 975–987. [https://doi.org/10.1016/S0896-6273\(02\)00817-6](https://doi.org/10.1016/S0896-6273(02)00817-6)
- Peterson, D. J., & Berryhill, M. E. (2013). The Gestalt principle of similarity benefits visual working memory. *Psychonomic Bulletin & Review*, 20(6), 1282–1289. <https://doi.org/10.3758/s13423-013-0460-x>

- Peterson, D. J., Gözenman, F., Arciniega, H., & Berryhill, M. E. (2015). Contralateral delay activity tracks the influence of Gestalt grouping principles on active visual working memory representations. *Attention, Perception, & Psychophysics*, 77(7), 2270–2283. <https://doi.org/10.3758/s13414-015-0929-y>
- Poldrack, R., Barch, D., Mitchell, J., Wager, T., Wagner, A., Devlin, J., Cumba, C., Koyejo, O., & Milham, M. (2013). Toward open sharing of task-based fMRI data: The OpenfMRI project. *Frontiers in Neuroinformatics*, 7. <https://www.frontiersin.org/article/10.3389/fninf.2013.00012>
- Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32(1), 3–25.
- Pylyshyn, Z. W. (2009). Perception, representation, and the world: The FINST that binds. In *Computation, cognition, and Pylyshyn* (pp. 3–48). MIT Press.
- Quinlan, P. T., & Cohen, D. J. (2012). Grouping and binding in visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(5), 1432–1438. <https://doi.org/10.1037/a0027866>
- Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by α -band EEG synchronization. *European Journal of Neuroscience*, 25(2), 603–610. <https://doi.org/10.1111/j.1460-9568.2007.05278.x>
- 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. <https://doi.org/10.1126/science.aah7011>
- Ruiz-Blondet, M. V., Jin, Z., & Laszlo, S. (2016). CEREBRE: A Novel Method for Very High Accuracy Event-Related Potential Biometric Identification. *IEEE Transactions on Information Forensics and Security*, 11(7), 1618–1629. <https://doi.org/10.1109/TIFS.2016.2543524>
- Salinsky, M. C., Oken, B. S., & Morehead, L. (1991). Test-retest reliability in EEG frequency analysis. *Electroencephalography and Clinical Neurophysiology*, 79(5), 382–392. [https://doi.org/10.1016/0013-4694\(91\)90203-G](https://doi.org/10.1016/0013-4694(91)90203-G)

- Sassenhagen, J., & Draschkow, D. (2019). Cluster-based permutation tests of MEG/EEG data do not establish significance of effect latency or location. *Psychophysiology*, *56*(6), e13335. <https://doi.org/10.1111/psyp.13335>
- Sauseng, P., Klimesch, W., Heise, K. F., Gruber, W. R., Holz, E., Karim, A. A., Glennon, M., Gerloff, C., Birbaumer, N., & Hummel, F. C. (2009). Brain Oscillatory Substrates of Visual Short-Term Memory Capacity. *Current Biology*, *19*(21), 1846–1852. <https://doi.org/10.1016/j.cub.2009.08.062>
- Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., Gruber, W. R., & 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. <https://doi.org/10.1111/j.1460-9568.2005.04482.x>
- Schalk, G., McFarland, D. J., Hinterberger, T., Birbaumer, N., & Wolpaw, J. R. (2004). BCI2000: A general-purpose brain-computer interface (BCI) system. *IEEE Transactions on Biomedical Engineering*, *51*(6), 1034–1043. <https://doi.org/10.1109/TBME.2004.827072>
- Schofield, P. N., Bubela, T., Weaver, T., Portilla, L., Brown, S. D., Hancock, J. M., Einhorn, D., Tocchini-Valentini, G., Hrabe de Angelis, M., & Rosenthal, N. (2009). Post-publication sharing of data and tools. *Nature*, *461*(7261), 171–173. <https://doi.org/10.1038/461171a>
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-Specific Delay Activity in Human Primary Visual Cortex. *Psychological Science*, *20*(2), 207–214. <https://doi.org/10.1111/j.1467-9280.2009.02276.x>
- Shen, M., Yu, W., Xu, X., & Gao, Z. (2012). Building Blocks of Visual Working Memory: Objects or Boolean Maps? *Journal of Cognitive Neuroscience*, *25*(5), 743–753. https://doi.org/10.1162/jocn_a_00348
- Shipstead, Z., Redick, T. S., & Engle, R. W. (2012). Is working memory training effective? *Psychological Bulletin*, *138*(4), 628–654. <https://doi.org/10.1037/a0027473>
- 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. <https://doi.org/10.1016/j.cub.2014.07.066>

- Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*(12), 1879–1887. <https://doi.org/10.1038/nn.3574>
- Srimal, R., & Curtis, C. E. (2008). Persistent neural activity during the maintenance of spatial position in working memory. *NeuroImage*, *39*(1), 455–468. <https://doi.org/10.1016/j.neuroimage.2007.08.040>
- Stokes, M. G. (2015). ‘Activity-silent’ working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405. <https://doi.org/10.1016/j.tics.2015.05.004>
- Sutterer, D. W., Foster, J. J., Adam, K. C. S., Vogel, E. K., & Awh, E. (2019). Item-specific delay activity demonstrates concurrent storage of multiple active neural representations in working memory. *PLOS Biology*, *17*(4), e3000239. <https://doi.org/10.1371/journal.pbio.3000239>
- Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). α -Band Electroencephalographic Activity over Occipital Cortex Indexes Visuospatial Attention Bias and Predicts Visual Target Detection. *The Journal of Neuroscience*, *26*(37), 9494–9502. <https://doi.org/10.1523/JNEUROSCI.0875-06.2006>
- Thyer, W., Adam, K. C. S., Diaz, G. K., Velázquez Sánchez, I. N., Vogel, E. K., & Awh, E. (2022). Storage in visual working memory recruits a content-independent pointer system. *Psychological Science*.
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, *428*(6984), 751–754. <https://doi.org/10.1038/nature02466>
- Todd, J. J., & Marois, R. (2005). Posterior parietal cortex activity predicts individual differences in visual short-term memory capacity. *Cognitive, Affective, & Behavioral Neuroscience*, *5*(2), 144–155. <https://doi.org/10.3758/CABN.5.2.144>
- Unsworth, N., Fukuda, K., Awh, E., & Vogel, E. K. (2014). Working memory and fluid intelligence: Capacity, attention control, and secondary memory retrieval. *Cognitive Psychology*, *71*, 1–26. <https://doi.org/10.1016/j.cogpsych.2014.01.003>

- van Beijsterveldt, C. E. M., & van Baal, G. C. M. (2002). Twin and family studies of the human electroencephalogram: A review and a meta-analysis. *Biological Psychology*, *61*(1), 111–138. [https://doi.org/10.1016/S0301-0511\(02\)00055-8](https://doi.org/10.1016/S0301-0511(02)00055-8)
- 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*, *107*(2), 900–905. <https://doi.org/10.1073/pnas.0908821107>
- Vision, T. J. (2010). Open Data and the Social Contract of Scientific Publishing. *BioScience*, *60*(5), 330–331. <https://doi.org/10.1525/bio.2010.60.5.2>
- Vogel, E. K., & Awh, E. (2008). How to Exploit Diversity for Scientific Gain: Using Individual Differences to Constrain Cognitive Theory. *Current Directions in Psychological Science*, *17*(2), 171–176. <https://doi.org/10.1111/j.1467-8721.2008.00569.x>
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*(6984), 748–751. <https://doi.org/10.1038/nature02447>
- Walker, P., & Davies, S. J. (2003). Perceptual completion and object-based representations in short-term visual memory. *Memory & Cognition*, *31*(5), 746–760. <https://doi.org/10.3758/BF03196113>
- Wang, S., Megla, E. E., & Woodman, G. F. (2020). Stimulus-induced Alpha Suppression Tracks the Difficulty of Attentional Selection, Not Visual Working Memory Storage. *Journal of Cognitive Neuroscience*, 1–27. https://doi.org/10.1162/jocn_a_01637
- Wang, S., Rajsic, J., & Woodman, G. F. (2019). The Contralateral Delay Activity Tracks the Sequential Loading of Objects into Visual Working Memory, Unlike Lateralized Alpha Oscillations. *Journal of Cognitive Neuroscience*, *31*(11), 1689–1698. https://doi.org/10.1162/jocn_a_01446
- Westerberg, H., Hirvikoski, T., Forsberg, H., & Klingberg, T. (2004). Visuo-Spatial Working Memory Span: A Sensitive Measure of Cognitive Deficits in Children With ADHD. *Child Neuropsychology*, *10*(3), 155–161. <https://doi.org/10.1080/09297040409609806>

- Wicherts, J. M. (2011). Psychology must learn a lesson from fraud case. *Nature*, *480*(7375), 7–7. <https://doi.org/10.1038/480007a>
- Wicherts, J. M., Borsboom, D., Kats, J., & Molenaar, D. (2006). The poor availability of psychological research data for reanalysis. *American Psychologist*, *61*(7), 726–728. <https://doi.org/10.1037/0003-066X.61.7.726>
- Woodman, G. F., Vecera, S. P., & Luck, S. J. (2003). Perceptual organization influences visual working memory. *Psychonomic Bulletin & Review*, *10*(1), 80–87. <https://doi.org/10.3758/BF03196470>
- Xu, Y. (2006). Understanding the object benefit in visual short-term memory: The roles of feature proximity and connectedness. *Perception & Psychophysics*, *68*(5), 815–828. <https://doi.org/10.3758/BF03193704>
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, *440*(7080), 91–95. <https://doi.org/10.1038/nature04262>
- Xu, Y., & Chun, M. M. (2007). Visual grouping in human parietal cortex. *Proceedings of the National Academy of Sciences*, *104*(47), 18766–18771. <https://doi.org/10.1073/pnas.0705618104>
- Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in Cognitive Sciences*, *13*(4), 167–174. <https://doi.org/10.1016/j.tics.2009.01.008>
- Yarkoni, T., Poldrack, R. A., Van Essen, D. C., & Wager, T. D. (2010). Cognitive neuroscience 2.0: Building a cumulative science of human brain function. *Trends in Cognitive Sciences*, *14*(11), 489–496. <https://doi.org/10.1016/j.tics.2010.08.004>
- Zyma, I., Tukaev, S., Seleznev, I., Kiyono, K., Popov, A., Chernykh, M., & Shpenkov, O. (2019). Electroencephalograms during Mental Arithmetic Task Performance. *Data*, *4*(1), 14. <https://doi.org/10.3390/data4010014>