

Memory-Based Attentional Guidance:
A Window to the Relationship between Working Memory and Attention

by

Emma Wu Dowd

Department of Psychology & Neuroscience
Duke University

Date: _____

Approved:

Tobias Egner, Supervisor

Stephen R. Mitroff

Scott A. Huettel

Michael L. Platt

Dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy in the Department of
Psychology & Neuroscience in the Graduate School of
Duke University

2016

ABSTRACT

Memory-Based Attentional Guidance:

A Window to the Relationship between Working Memory and Attention

by

Emma Wu Dowd

Department of Psychology & Neuroscience
Duke University

Date: _____

Approved:

Tobias Egner, Supervisor

Stephen R. Mitroff

Scott A. Huettel

Michael L. Platt

An abstract of a dissertation submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy in the Department of
Psychology & Neuroscience in the Graduate School of
Duke University

2016

Copyright by
Emma Wu Dowd
2016

Abstract

Attention, the cognitive means by which we prioritize the processing of a subset of information, is necessary for operating efficiently and effectively in the world. Thus, a critical theoretical question is how information is selected. In the visual domain, working memory (WM)—which refers to the short-term maintenance and manipulation of information that is no longer accessible by the senses—has been highlighted as an important determinant of what is selected by visual attention. Furthermore, although WM and attention have traditionally been conceived as separate cognitive constructs, an abundance of behavioral and neural evidence indicates that these two domains are in fact intertwined and overlapping. The aim of this dissertation is to better understand the nature of WM and attention, primarily through the phenomenon of memory-based attentional guidance, whereby the active maintenance of items in visual WM reliably biases the deployment of attention to memory-matching items in the visual environment. The research presented here employs a combination of behavioral, functional imaging, and computational modeling techniques that address: (1) WM guidance effects with respect to the traditional dichotomy of top-down versus bottom-up attentional control; (2) under what circumstances the contents of WM impact visual attention; and (3) the broader hypothesis of a predictive and competitive interaction between WM and attention. Collectively, these empirical findings reveal the importance of WM as a distinct factor in attentional control and support current models of multiple-state WM, which may have broader implications for how we select and maintain information.

Dedication

For my beloved Maisie.

Contents

Abstract.....	iv
List of Tables	xi
List of Figures	xii
Acknowledgements.....	xvi
1. General Introduction	1
1.1 A Traditional Dichotomy of Attentional Selection.....	3
1.2 Working Memory and Attention.....	6
1.2.1 The Role of Attention in Working Memory	6
1.2.2 The Role of Working Memory in Attention	8
1.2.3 A Reconceptualization of Working Memory and Attention.....	10
1.3 Working Memory Guides Visual Attention.....	13
1.4 The Limits of Attentional Guidance by Working Memory	17
1.5 Aims of the Present Thesis	21
2. Working Memory Guidance Overrides Bottom-Up Salience Cues.....	23
2.1 Introduction.....	23
2.2 Current Goals	26
2.3 General Method	28
2.3.1 Apparatus	28
2.3.2 Stimuli.....	28
2.3.3 Procedure	29
2.3.4 Trial Conditions	31

2.3.5 Data Analysis	32
2.4 Experiment 1A	36
2.4.1 Results	36
2.4.2 Discussion	37
2.5 Experiment 1B	39
2.5.1 Results	39
2.5.2 Discussion	40
2.6 Experiment 2	41
2.6.1 Verification of “Pop-Out”-Like Search	42
2.6.2 Results	43
2.6.3 Discussion	44
2.7 Experiment 3	45
2.7.1 Results	46
2.7.2 Discussion	46
2.8 General Discussion	47
2.8.1 Working Memory Guidance or Strategy?	49
2.8.2 Conclusion	52
3. Top-Down Control Over Working Memory Guidance Operates Via Working Memory Accessibility	54
3.1 Introduction	55
3.2 General Methods	61
3.2.1 Apparatus	61
3.2.2 Procedure	62

3.2.3 Design	64
3.2.4 Data Analysis	65
3.3 Experiment 1	68
3.3.1 Behavioral Results	69
3.3.2 Model Fitting and Comparison	70
3.3.3 Discussion	71
3.4 Experiment 2	73
3.4.1 Behavioral Results	74
3.4.2 Model Fitting and Comparison	75
3.4.3 Discussion	76
3.5 Comparison of Model Families	76
3.6 General Discussion	79
4 Individual Differences in Working Memory Guidance	86
4.1 Introduction	87
4.2 Methods	89
4.2.1 Participants	89
4.2.2 Behavioral Testing	91
4.2.2.1 Dual-Task Paradigms	91
4.2.2.2 Visual Working Memory Capacity	93
4.2.2.3 Operation Span	94
4.2.3 Individual-Differences Assessments	94
4.3 Results	95

4.3.1 Binary-Stimulus Paradigm.....	95
4.3.2 Unitary-Stimulus Paradigm	96
4.3.3 Visual Working Memory Capacity	98
4.3.4 Operation Span.....	98
4.3.5 Regression Analyses	98
4.4 Discussion.....	102
5 Decoding Working Memory Content from Attentional Biases	106
5.1 Introduction.....	106
5.2 Method	109
5.2.1 Participants.....	109
5.2.2 Stimuli and Procedure.....	109
5.2.3 Classification Analysis.....	112
5.3 Results.....	115
5.3.1 Mean Performance Data	115
5.3.2 Classification Results.....	117
5.4 Discussion.....	121
6 Neural Working Memory Representations are Modulated by Attentional Demands..	127
6.1 Introduction.....	128
6.2 Methods.....	130
6.2.1 Participants.....	130
6.2.2 Design	131
6.2.3 Dual-Task WM/Visual Search Procedure.....	132

6.2.4 Functional Localizer Procedure	135
6.2.5 Image Acquisition	135
6.2.6 fMRI Analyses	135
6.2.7 Univariate Analyses	136
6.2.8 ROI Definition	137
6.2.9 Multivariate Analyses	138
6.2.9.1 Searchlight MVPA	139
6.2.9.2 Time-Resolved ROI-Based MVPA	140
6.3 Results	141
6.3.1 Behavioral Results	141
6.3.2 Univariate fMRI Results	143
6.3.3 Multivariate fMRI Results	145
6.4 Discussion	148
7 General Discussion	153
7.1 Working Memory as a Distinct Influence on Attention, neither “Top-Down” nor “Bottom-Up”	154
7.2 Flexible, Multiple-State Models of Working Memory and Attention	158
7.3 Shared Resources and Representations of Working Memory and Attention	160
7.4 Conclusion	165
Appendix A	166
References	167
Biography	194

List of Tables

Table 1: Trial Types by Number of Search Targets, Saliency of Search Targets, and Memory-to-Search Target Relationship	32
Table 2: Overall Accuracy and Search Times to Find a Target First in Dual-Target Searches	34
Table 3: Descriptions of Model Families and Model Prior Probability Distributions for Each Parameter across Block Types	68
Table 4: Mean Accuracy and Response Times for Visual Search and Memory Tasks for Each Trial Type.....	72
Table 5: Mean <i>sd</i> and <i>g</i> Values for Continuous-Report Memory Task for Each Block Type	76

List of Figures

Figure 1: Example trial sequence of Experiment 1A showing a memory task with an intervening search task. In this example, the search task includes two targets, a high-salience target and a low-salience target. This trial depicts a valid, high-match trial because the memory cue matches the high-salience target in both luminance and orientation. The memory probe differs from the memory cue in this example. 31

Figure 2: Data from Experiment 1A demonstrate that WM guidance both enhances and overrides a baseline salience effect. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors. 38

Figure 3: Data from Experiment 1B replicate the WM guidance effects from Experiment 1A. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors. Note the values may not total to 100% due to rounding. 41

Figure 4: Data from Experiment 2 verify the “pop-out” nature of the high-salience target. Search times as a function of set size when searching for low-salience versus high-salience targets among distractors, with slope values provided. Errors bars represent standard errors. 43

Figure 5: Data from Experiment 2 show that WM guidance modulates an amplified visual salience effect. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors. 45

Figure 6: Data from Experiment 3 reveal a minimal influence of visual priming. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors. 47

Figure 7: Example trial sequence. Participants were to remember a cue color while performing an intervening visual search task, in which they reported the direction of the tilted line. Then participants were given one memory task: either a continuous-report task (*left*), in which they recalled the memory cue from a continuous spectrum of responses, or a change-detection task (*right*), in which they responded whether the memory probe was same or different from the memory cue. 64

Figure 8: Bayesian model comparison between individual models for each participant. Each column represents a subject, divided by Experiment 1 (A) and Experiment 2 (B), and each row represents a mixture model as identified by model number (see Table 3) and prior probability distributions over two parameters: the standard deviation of the

Gaussian component (i.e., precision; sd) and the probability of drawing from the uniform component (i.e., guessing; g). Prior probabilities were applied across three block types: Valid (V), Invalid (I), and Mixed (M). Curly brackets indicate variable precision or guessing across the three block types. Cell color indicates a model's deviance information criterion (DIC) relative to that of the subject's best-fitting model (a higher value means a worse fit). In general, the best-fitting models (redder cells) specify constant precision, whereas the worst-fitting models (bluer cells) specify variable or ordered precision..... 73

Figure 9: Comparison of individual models (A) and of four major families of models (B). Bars indicate the probability that a given model family generated the data of a randomly chosen subject. P^* values indicate exceedance probability (i.e., posterior probability that a given model is more likely than any other model). The best-fitting model (Model 7) specifies constant precision but ordered guessing (Invalid > Mixed > Valid), and the best-fitting model family specifies constant precision but variable/ordered guessing across block types. 78

Figure 10: Trial sequences for the four behavioral tasks. (A) Binary-stimulus dual-task paradigm that combines a memory task with an intervening search task; participants searched for the tilted line target among vertical or horizontal line distractors. (B) Unitary-stimulus dual-task paradigm; participants searched for the top or bottom gap target among right or left gap distractors. In both dual-task paradigms, the memory task tested whether the memory probe was same or different from the memory cue. (C) Visual working memory capacity task; participants reported whether memory probe color was same or different from the item at that location in the sample array. (D) Automated operation span task; participants alternated between remembering letters and solving arithmetic problems. 90

Figure 11: Memory-based capture results. (A) In both the binary-stimulus and unitary-stimulus paradigms, search times were slower for invalid trials compared to neutral trials. Error bars represent 95% confidence intervals. (B) Individual differences in normalized difference scores (i.e., the amount of attentional capture by memory-matching distractors) did not correlate between the binary-stimulus and unitary-stimulus paradigms. 97

Figure 12: Fitted linear regression models predicting memory-based capture in each dual-task paradigm by selected, significant variables of working memory and attention. (A) Normalized difference scores (i.e., the amount of attentional capture by memory-matching distractors) in the binary-stimulus paradigm were significantly predicted by individual differences in visual working memory capacity. (B) Normalized difference scores in the unitary-stimulus paradigm were significantly predicted by self-report assessments of attention deficit-hyperactivity disorder and two sub-categories of the autism spectrum quotient. Lines depict 95% confidence intervals around mean

standardized regression coefficients (black dots) in predicting normalized difference scores in a multiple regression model. 101

Figure 13: Example trial sequence. Participants were to remember a cue color (red, blue, green, or yellow) while performing a series of 12 intervening visual searches, in which they reported the direction of the tilted line. Then participants were given a four-alternative forced-choice memory probe array and asked to report which color matched the initial memory cue. In this example, the cue color is red; in the series of 12 searches, 3 are valid (top row; target is red), 3 are invalid (second row; distractor is red), and 6 are neutral (bottom two rows; neither target nor distractor is red). 112

Figure 14: The multidimensional feature space was organized according to the 12 possible two-color (6) and target-distractor (2) arrangements (*top*). Classifier inputs were search response times for each feature. Thus, each trial contributed a vector of 12 search response times, labeled by the corresponding memory color; four example input vectors are shown in the grid. Response times were detrended across blocks and rescaled within each feature column to [0 1]. For within-subject classification analyses, we adopted a leave-one-trial-out cross-validation scheme; within each participant's data, the classifier was trained on the vectors corresponding to all but one trial (*training set*), then tested on its ability to decipher the memory label on that one trial (*testing set*). 115

Figure 15: Mean search response times revealed significant validity effects. (A) Response times were overall faster on valid searches and slower on invalid searches, compared to neutral searches. (B) Validity effects were still significant after splitting the series of 12 searches into first and second halves, indicating that a single instance of WM maintenance could impact attentional guidance across multiple subsequent searches. The diameter of each dot represents 95% confidence intervals. *** $p < .001$ 117

Figure 16: Within-subject classification results. (A) Across three classifiers, within-subject decoding accuracy was significantly above chance (dashed line at 25%). For each classifier, a violin plot illustrates the distribution of that classifier's accuracy over the 100 participants. The boxes mark the middle two quartiles of classification performance, and the whiskers represent the full range of classification performance (outliers are marked with a dot). (B) Individual differences in classification accuracy for all three approaches were significantly predicted by validity effect (i.e., magnitude of attentional guidance), suggesting that trained classifiers capitalized on patterns of search validity effects that emerged across each trial. *** $p < .001$ 120

Figure 17: Between-subject decoding accuracy was also significantly above chance for all three classifiers, as evaluated by binomial tests against chance (dashed line at 25%). For each classifier, a violin plot illustrates the distribution of that classifier's accuracy over the 100 participants. The boxes mark the middle two quartiles of classification

performance, and the whiskers represent the full range of classification performance (outliers are marked with a dot). *** $p < .001$ 121

Figure 18: Behavioral task design. (A) During the delay interval of a match-to-sample WM task, participants completed a series of four visual searches. (B) WM load conditions: Participants maintained either 1 (low) or 2 (high) face or house stimuli in WM. (C) Visual search load conditions: Participants searched for the vertical body or tool target among horizontal (easy) or tilted (hard) distractors. 134

Figure 19: Behavioral results. (A) Mean visual search response time (RT) was faster for easy searches compared to hard searches. (B) WM probe recognition accuracy was higher for low load (1 item) compared to high load (2 item) WM. (C) Load effect scores (i.e., the impact of higher demand on search RT or WM accuracy) on the visual search and WM portions of the dual-task paradigm were significantly correlated. 143

Figure 20: Univariate fMRI results. Rendered data are whole-brain FDR-corrected, $p < .05$. (A) A large bilateral network of frontal, parietal, and occipital regions showed elevated activation for hard versus easy visual searches. (B) Lateral prefrontal, parietal, and posterior temporal clusters were sensitive to the interaction between WM and search load levels. Representative beta values are displayed for an ROI surrounding the peak of the interaction contrast in the middle frontal gyrus (MFG). 145

Figure 21: Multivariate fMRI results. (A) Clusters of searchlights across the ventral visual stream classified WM category significantly above chance during the visual search sequence. (B) In functionally-defined parahippocampal place area (PPA), at 2 s timepoints across the trial, WM category classification displayed a U-shaped pattern, suggesting an allocation of attention away from WM during the concurrent search task. WM decoding accuracy remained above chance across easy (gray) visual searches, but dropped below chance later on in the search sequence for hard (orange) visual searches. Ribbons represent ± 1 SEM. C = WM cue; D = delay; S = search; P = probe. 148

Acknowledgements

I am deeply grateful for the brilliant and generous mentorship of Tobias Egner, who has offered invaluable wisdom and support at every stage in this process. I am equally thankful to Steve Mitroff for his unwavering guidance and support. Many thanks also to my committee members Michael Platt and Scott Huettel, as well as Liz Brannon and Kevin LaBar, for all of their knowledge and help. I also wish to acknowledge my collaborators Jeff Beck, John Pearson, and especially Anastasia Kiyonaga, volleyball partner extraordinaire. Thank you to all of my amazing friends, mentors, and colleagues at Duke and beyond—your generosity, inspiration, and kindness during the best and the worst of times will never be forgotten. And thank you to my husband, Jason, for always believing in me.

1. General Introduction

We are constantly bombarded with a wealth of sensory information that overloads our limited cognitive capacities; thus, this information must be sorted, filtered, and processed for the purpose of interacting with the world. For example, the act of finding your gate at the airport involves: searching for the correct signage, monitoring pre-boarding announcements, navigating through crowds of people, and keeping track of your own belongings—all of which require separating relevant information from irrelevant distraction. This crucial selection mechanism is broadly referred to as “attention,” the cognitive means by which we prioritize the processing of a subset of information. To operate efficiently and effectively in the world, selective attention is an integral property of all stages of cognition and behavior. Thus, a critical theoretical question is how limited processing resources are allocated—specifically, what are the factors that drive attention?

According to the biased-competition model of visual attention (Beck & Kastner, 2009; Desimone & Duncan, 1995), simultaneously-presented visual items activate populations of neurons that compete with each other for attentional selection; various factors can bias neural processing in favor of one item over the others, resulting in the deployment of attention to a particular location in the visual field. Selective attention can also be framed as an activation or priority map that incorporates the weighted contributions of different attentional influences, such that the location with the highest current activation is selected (e.g., Koch & Ullman, 1985; Treue, 2003; Wolfe, 2007;

Zelinsky & Bisley, 2015). In both of these types of models, control over visual attention has been traditionally conceived as a dichotomy of “top-down” factors (e.g., an item’s task-relevance; see Egeth & Yantis, 1997) and “bottom-up” factors (e.g., an item’s physical distinctiveness; see Theeuwes, 2010); however, this top-down/bottom-up dichotomy of attentional control has been criticized for failing to address other types of factors, such as memory (Awh, Belopolsky, & Theeuwes, 2012; Hutchinson & Turk-Browne, 2012).

Specifically, working memory (WM)—which refers to the short-term maintenance and manipulation of information that is no longer accessible by the senses—has been highlighted as an important determinant of what is selected by visual attention (see Awh et al., 2012; Awh, Vogel, & Oh, 2006; Hutchinson & Turk-Browne, 2012). The primary aim of this thesis is to situate WM with respect to the traditional dichotomy of attentional control and elaborate on the nature of memory-based¹ attentional guidance from novel perspectives. More broadly, this thesis also aims to contribute to a better understanding of the shared mechanisms underlying WM and visual attention.

¹ Here, “memory-based” guidance refers to non-volitional attentional biasing toward the contents of working memory. While other types of memory (e.g., priming, associative, long-term) have also been shown to impact visual attention, in this thesis, I use “memory-based” to specifically refer to the active maintenance of information in working memory.

1.1 A Traditional Dichotomy of Attentional Selection

Traditional models of visual attention (e.g., Bundesen, 1990; Duncan & Humphreys, 1989; Wolfe, 2007) account for two major drivers of attentional selection: “top-down” influences (e.g., when attention is driven by internal task goals) and “bottom-up” considerations (e.g., when attention is captured by the physical properties of a stimulus). While much debate exists regarding the relative importance and time courses of these two contributors (see Egeth & Yantis, 1997; Folk, Remington, & Johnston, 1992; Theeuwes, 2010), it has been generally agreed that top-down and bottom-up factors interact to ultimately drive visual selection (e.g., Serences et al., 2005; Wolfe, 2007).

To examine how different factors influence visual attention, researchers often use visual search paradigms (see Eckstein, 2011), in which participants look for targets among distractors—a task that requires participants to selectively attend to particular locations across the search array. In these experiments, the time that it takes to respond to the target item is interpreted as a proxy for selecting (and rejecting) some number of non-target distractors before selecting and processing the target; different manipulations of task demands and target/non-target features then inform theories about how visual attention shifts across time and space (e.g., Duncan & Humphreys, 1989; Eimer, 2014; Eriksen & Hoffman, 1973).

In terms of visual search, top-down and goal-driven attentional selection is an intuitive concept—knowledge about the search target forms an “attentional template” or “task set” that is compared to information in the visual field, such that attention is guided

to locations with features that match the target template (e.g., Duncan & Humphreys, 1989; Vickery, King, & Jiang, 2005). For example, given the goal of finding bananas in a grocery store, attentional processes would favor yellow items. Evidence for top-down attention has also been demonstrated with contingent capture paradigms, in which participants are faster to find a particularly colored target when a same-colored cue appears immediately before the target display at the target location; this benefit, however, is limited to cues that match the target color, indicating that attention is drawn only toward template-matching, task-relevant features (Folk et al., 1992).

Bottom-up attentional selection, on the other hand, is determined by the physical characteristics of the stimulus itself (see Egeth & Yantis, 1997; Theeuwes, 2010; Yantis, 2000). Again, we can intuitively think about physically salient items that “pop-out” of a visual scene, such as a Granny Smith in a bushel of Red Delicious apples (Treisman, 1986). Visual salience is thus defined as differences in visual cues (e.g., luminance, color, motion, orientation, depth) between an item and the rest of the visual field (Nothdurft, 2002). Therefore, if visual stimuli are physically distinctive enough, these items involuntarily capture attention. Much of the evidence for bottom-up attentional capture comes from the additional singleton paradigm (Theeuwes, 1992)—when participants are instructed to find a shape singleton (e.g., a diamond shape among circles), search performance is nevertheless impeded if one of the distractors is a color singleton (e.g., red among green), suggesting that the task-irrelevant color singleton is inadvertently selected even when it does not match the target template (see Theeuwes, 2010).

More recently, the dichotomy of top-down and bottom-up attentional control has been criticized for overshadowing the contributions of other influences on attention (Awh et al., 2012; Hutchinson & Turk-Browne, 2012). For instance, considerable research has shown that attentional deployment is influenced by the priming of visual stimuli across trials, such that repeated targets are found more easily than novel ones (e.g., Belopolsky, Schreij, & Theeuwes, 2010; Maljkovic & Nakayama, 1994; see Kristjánsson & Campana, 2010; Lamy & Kristjánsson, 2013). Likewise, previous experience with particular task sets or search strategies across longer time periods encourages and enhances the implementation of those attentional sets (e.g., Leber & Egeth, 2006; Leber, Kawahara, & Gabari, 2009; Vatterott & Vecera, 2012). Visual attention is also biased by previously learned associations with contextual cues (e.g., Chun & Jiang, 1998; Summerfield, Lepsien, Gitelman, Mesulam, & Nobre, 2006; see Chun & Turk-Browne, 2008), reward (e.g., Libera & Chelazzi, 2006; see Anderson, 2015), and emotion (see Vuilleumier, 2005; 2015). Additionally, long-term conceptual associations between items within a visual array can influence attentional selection—for instance, visual search for a real-world target object is more negatively impacted by the presence of conceptually-related distractors than conceptually-unrelated ones (Moore, Laiti, & Chelazzi, 2003). Finally, as the focus of this thesis, the contents of WM have also been demonstrated to impact the deployment of visual attention (see Olivers, Peters, Houtkamp, & Roelfsema, 2011; Soto, Hodsoll, Rotshtein, & Humphreys, 2008; Woodman, Carlisle, & Reinhart, 2013). While all of these effects tap into various forms of prior experience or memory, it has been

argued that they are conceptually distinct from traditional top-down or bottom-up labels (Awh et al., 2012; Hutchinson & Turk-Browne, 2012). In the first part of this thesis, I argue that WM is a distinct and important factor in attentional selection that functions independently of and interacts with traditional top-down and bottom-up mechanisms.

1.2 Working Memory and Attention

Working memory (WM) is a temporary and capacity-limited storage system that involves the active maintenance and manipulation of information necessary for ongoing tasks (Baddeley, 2003). Specifically, WM operates on “internal” information that is not immediately available to the senses (Chun, Golomb, & Turk-Browne, 2011); correspondingly, visual WM is the short-term maintenance of visual information without perceptual input. Although traditionally conceptualized as a separate cognitive construct, WM is recognized to be tightly interconnected with attention (see Awh et al., 2006; Awh & Jonides, 2001; Gazzaley & Nobre, 2012). In this section, I provide a brief review of the bidirectional interactions between specifically visual WM and visual attention.

1.2.1 The Role of Attention in Working Memory

Visual attention is thought to serve as a “gatekeeper” for WM, selecting what is encoded into WM (see Awh et al., 2006; Irwin & Gordon, 1998). Across a visual field, for instance, the more attention allocated to particular item, the more likely that a perceptual representation of that item will be encoded into WM (Duncan & Humphreys, 1989). In one study, researchers cued spatial attention immediately prior to a delayed match-to-sample WM task and found that attending to a specific location increased the

probability of encoding information in that location into WM (Schmidt, Vogel, Woodman, & Luck, 2002); neural electroencephalogram (EEG) markers of selective attention toward task-relevant stimuli at encoding also predict subsequent WM accuracy (Rutman, Clapp, & Chadick, 2010), indicating that selective attention is necessary for successful WM encoding. Thus, the capacity limits of WM (see Luck & Vogel, 2013) may be a function of the limits of selective attention (see also Cowan, 2001; Engle, 2002). Conversely, the selective nature of attention may be a product of WM limits—for instance, the Neural Theory of Visual Attention (Bundesen, Habekost, & Kyllingsbæk, 2005) posits that attentional processing of a visual object by definition encodes it into visual WM, which then maintains attentional selection via neural feedback mechanisms (but see Woodman, Vogel, & Luck, 2001 for conflicting behavioral evidence).

Once encoded, the maintenance of information in WM is also hypothesized to depend on attentional mechanisms. Much of the evidence for this idea comes from “retro-cue” paradigms, in which a cue retrospectively (i.e., after initial encoding, during the delay period) indicates which specific WM representation is more task-relevant and should be prioritized (e.g., Griffin & Nobre, 2003; Makovski, Sussman, & Jiang, 2008; Sperling, 1960). The presentation of a retro-cue during the WM delay confers greater memory benefits for that cued representation, indicating that the maintenance and prioritization of internal information can be shifted much in the same way that visual attention shifts across sensory information (Griffin & Nobre, 2003). Similarly, the rehearsal of spatial information in WM is thought to be mediated by attentional

mechanisms (Awh & Jonides, 2001; but see Belopolsky & Theeuwes, 2009), as WM accuracy for spatial locations is impaired by unrelated spatial shifts of attention during memory retention (Awh, Jonides, & Reuter-Lorenz, 1998), and neural activation in posterior visual cortex is enhanced contralateral to both attended and memorized locations (Awh et al., 1999; Postle, Awh, Jonides, Smith, & D'Esposito, 2004). Neural markers of selective attention for perceptual representations are also evoked by selecting internal representations within WM, providing further evidence that WM maintenance relies on functionally and neurally analogous attentional mechanisms (Dell'Acqua, Sessa, Toffanin, Luria, & Jolicoeur, 2010; Kuo, Rao, & Lepsien, 2009; Kuo, Stokes, Murray, & Nobre, 2014).

1.2.2 The Role of Working Memory in Attention

While the above research has framed attention as a crucial cognitive process in service of WM, the relationship between attention and WM also operates in the other direction—specifically, the contents of WM are likewise necessary for determining what is selected by visual attention. In visual search, the target template is presumed to be instantiated as a memory representation that specifies the particular target features, facilitating the selection of template-matching perceptual inputs (e.g., Bundesen, 1990; Duncan & Humphreys, 1989; Treisman, 1988; Zelinsky & Bisley, 2015). Similarly, the biased-competition model of visual attention posits that active maintenance of a template representation in WM biases neural competition toward matching features in the external visual field, allowing template-matching items to better compete for attentional

processing (Desimone & Duncan, 1995). For example, one neurophysiological study recorded activity from neurons in monkey inferior temporal cortex during a delayed match-to-sample WM task (Chelazzi, Miller, Duncan, & Desimone, 1993). Monkeys remembered a target (i.e., cue) stimulus across a delay period, then were presented with two probe stimuli and had to make an eye movement to the probe stimulus that matched the initial memory cue. A subset of inferior temporal neurons demonstrated elevated firing activity across the delay period, which was interpreted as reflecting cue maintenance in WM. Importantly, after the onset of the probe array, neural activity diverged even before any eye movement response, with elevated firing rates for WM-related neurons and suppressed firing rates for non-WM-related—indicating that WM maintenance biases neural competition for cells that selectively respond to the target (Chelazzi et al., 1993).

Humans also demonstrate electrophysiological markers of holding an attentional template in WM. In EEG, an event-related component called contralateral delay activity (CDA) is a relative negativity at parietal, occipital, and temporal locations that indexes WM maintenance (Vogel & Machizawa, 2004). In a recent study, not only did the presentation of a target cue elicit a CDA that continued through a visual search array, but the amplitude of the CDA was also correlated to search response time performance, indicating that the robustness of a target template in WM impacted the guidance of attention in subsequent search (Woodman & Arita, 2011; see also Woodman et al., 2013).

Evidence that search templates are maintained in WM supports the idea that top-down, goal-driven influences on visual attention are instantiated through WM. However, it is also possible that attentional biasing by WM goes beyond the strict definition of top-down labeling. Thus, while many models of visual attention include the intentional usage of WM representations as attentional templates (e.g., Bundesen, 1990; Desimone & Duncan, 1995; Duncan & Humphreys, 1989), a stronger test of the hypothesis that attentional templates are stored in WM would be to show that the maintenance of *any* information (i.e., not goal-driven and unrelated to a search target) in WM is sufficient to guide visual attention to memory-matching items in the environment (see Olivers et al., 2011; Soto et al., 2008; Woodman et al., 2013). For instance, one early study asked participants to remember a face, then perform a speeded decision task at one of two locations; critically, during the delay, two faces were presented, one of which matched the memory stimulus (Downing, 2000). Participants were faster to respond at the location where the memory-matching face had previously appeared, indicating that holding a visual object in WM induced an attentional shift toward the location of the memory-matching item. As the primary focus of this thesis, further relevant supporting literature for (unintentional) memory-based attentional guidance (i.e., WM guidance) is reviewed more thoroughly in the following two sections.

1.2.3 A Reconceptualization of Working Memory and Attention

Given the evidence for much overlap and interaction between WM and attention, it has recently been proposed that the close reciprocity between WM and attention might

be thought of as two sides of the same coin—specifically, that WM and attention could draw from, and thus compete for, a common cognitive resource (see Chun et al., 2011; Kiyonaga & Egner, 2013). Such competition has been previously demonstrated in terms of dual-task interference effects. For example, concurrent load in WM modulates the magnitude with which distractors interfere with an attentional task (de Fockert, Rees, Frith, & Lavie, 2001). In a dual-task paradigm that combined visual selective attention (i.e., a face-Stroop task: classifying written names as pop stars or politicians while ignoring distracting congruent or incongruent faces of pop stars and politicians) with WM load (i.e., remembering a series of digits), participants were less able to ignore distractors when WM load was high compared to low, indicating that WM availability was necessary for successful attentional control (de Fockert et al., 2001).

More specific to the visual domain, two studies found that maintaining spatial locations (but not object identities; see Woodman et al., 2001; Woodman, Luck, & Schall, 2007) in WM while simultaneously performing unrelated visual search tasks interfered with search efficiency (Oh & Kim, 2004; Woodman & Luck, 2004). Visual WM maintenance also competes with dynamic attentional tracking of moving objects (i.e., multiple object tracking; MOT)—in dual-task paradigms combining WM and MOT tasks, maintaining visual objects in WM interfered with MOT performance, and vice versa, tracking multiple objects interfered with WM accuracy (Fougnie & Marois, 2006). These reciprocal impairments also tracked with load, such that dual-task interference was more severe as the number of items in WM or the number of moving targets to track

increased. Taken together, these dual-task interference effects point to a closely intertwined relationship between WM and attention, such that the two cognitive constructs may overlap and share some—but not all—common processing mechanisms (see Fougne, 2008 for a review; see also Hedge, Oberauer, & Leonards, 2015).

In fact, the construct of WM has recently been reconceptualized as internally-directed attention that activates perceptual representations (and thus sensory cortical regions) much in the same way that externally-directed attention to external stimuli would (e.g., Chun, 2011; Kiyonaga & Egner, 2013). For example, recent studies have found internal representations in WM to exhibit the same spatial and feature-based attributes that selectively attended visual representations do, such as spatial gradient models of processing (Sahan, Verguts, Boehler, & Pourtois, 2015), center-surround feature inhibition (Kiyonaga & Egner, 2015), object-based processing benefits (Peters, Kaiser, Rahm, & Bledowski, 2015), and even the role of oculomotor eye movements (Theeuwes, Belopolsky, & Olivers, 2009). The neural networks that control WM and visual attention are also closely intertwined, especially in several frontal and parietal regions (e.g., Cabeza & Nyberg, 2000; Corbetta & Shulman, 2002; LaBar, Gitelman, Parrish, & Mesulam, 1999). More recently, these fronto-parietal regions have been taken to reflect attentional control over sensory regions that represent the information content themselves (Gazzaley & Nobre, 2012; Sreenivasan, Curtis, & D'Esposito, 2014a), such that object and feature representations encoded by visual cortex are activated both when externally attending to those visual stimuli and when internally maintaining them in WM

(e.g., Albers, Kok, Toni, Dijkerman, & de Lange, 2013; Christophel, Hebart, & Haynes, 2012; Ester, Anderson, Serences, & Awh, 2013; Ester, Sprague, & Serences, 2015; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009). This line of research continues to be a fruitful avenue for exploring the shared mechanisms and shared representations between WM and attention.

1.3 Working Memory Guides Visual Attention

A specific example of the close link between WM and attention is the role that WM plays in guiding visual attention. While many models of visual attention presume that top-down deployment operates via an attentional template held in WM (Bundesen, 1990; Desimone & Duncan, 1995; Duncan & Humphreys, 1989), many recent studies have shown that attentional guidance by WM is not limited to the intentional usage of WM content as a template (see Olivers et al., 2011; Soto et al., 2008; Woodman et al., 2013). Specifically, the biased-competition model (Desimone & Duncan, 1995) predicts that the maintenance of *any* information in WM will activate memory-matching sensory features, providing an advantage for further attentional processing in the visual field. Thus, one strong interpretation of this model is that the active maintenance of information in WM *automatically* (i.e., non-volitionally) biases visual attention to memory-matching items in the environment.

To test this hypothesis, studies of (unintentional) memory-based attentional guidance typically use dual-task paradigms that combine WM and visual attention tasks. For example, in one influential paradigm, participants remember an item (e.g., a colored

shape) while performing an intervening but unrelated visual search task during the delay period between the WM cue and subsequent probe (e.g., Olivers, Meijer, & Theeuwes, 2006; Soto, Heinke, Humphreys, & Blanco, 2005). Critically, the memory item can reappear in the search display, either coinciding with the location of a target (i.e., “valid” search) or with the location of a distractor (i.e., “invalid” search), or fail to reappear in the display at all (i.e., “neutral” search). Note that the reappearance of the memory item in the search display is technically orthogonal and unrelated to the search task. The canonical finding is that search is speeded when the memory item validly cues the target location, and slowed when the memory item invalidly cues a distractor location (relative to the neutral condition; see Soto et al., 2008; but see also Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Woodman & Luck, 2007). These response time validity effects have been bolstered by similar patterns of eye movements, such that more first saccades land on memory-matching items in the search display (Olivers et al., 2006; Soto et al., 2005; but see Houtkamp & Roelfsema, 2006). Electrophysiological measures of selective attention also support the early deployment of attention toward memory-matching locations (but see Carlisle & Woodman, 2011b; Peters, Goebel, & Roelfsema, 2009). Recent EEG studies have found that the N2pc (i.e., a posterior negative deflection occurring 200–300 ms after stimulus onset that reflects the efficiency of target selection; see Kiss, Van Velzen, & Eimer, 2008) is enhanced when the memory-matching distractor appears in the same side of space as the target item (Carlisle & Woodman, 2013; Kumar, Soto, & Humphreys, 2009).

One alternative explanation suggests that attentional effects are driven not by WM but by visual priming of stimulus features. In other words, mere exposure to the memory item, even without a memory requirement, before the visual search task could be sufficient to bias attention toward primed features. Several studies have ruled out this possibility by including a control condition in which participants are presented with an item before search, but know that they will not be tested in a subsequent memory task; no attentional guidance effects were found when there was no requirement to hold the item in WM (e.g., Downing, 2000; Soto et al., 2005). In another set of studies, the order of tasks was changed such that the WM test was administered *prior* to the search task, effectively “priming” the cue representation twice (once during the cue presentation and once during the WM test; Olivers et al., 2006; Olivers & Eimer, 2011). Again, no attentional guidance effects were found, suggesting that after the memory test was completed, the memory item was released (or defocused; see Oberauer, 2002) from active WM maintenance, thus losing its effect on biasing visual attention (Olivers et al., 2006). Neuroimaging evidence also contradicts the notion that WM guidance is driven by visual priming, as the reappearance of an actively-held WM item versus mere repetition of a previously seen item elicits distinct functional magnetic resonance imaging (fMRI) responses in parahippocampal gyrus, lingual gyrus, and superior frontal gyrus (Soto, Humphreys, & Rotshtein, 2007). Similarly, applying transcranial magnetic stimulation (TMS) to occipital cortex differentially affects the impact of WM on search performance,

depending on whether the initial cue is actively maintained in WM or merely attended (Soto, Llewelyn, & Silvanto, 2012b).

A strong test of the automaticity claim of WM guidance (or “capture”) requires the memory-matching stimulus to be unhelpful or even detrimental to the current task. Thus, several studies have included conditions in which attentional biasing by WM content is reliably detrimental to search performance—for instance, in conditions where the memory-matching item always invalidly cues the location of a distractor (e.g., Kiyonaga, Egner, & Soto, 2012; Olivers et al., 2006; Soto et al., 2005). In these tasks, despite knowing that memory-matching contents in the search array were *never* predictive of the search target, participants still demonstrated attentional costs in the face of invalid distractors, supporting automatic effects of memory-based capture (see also Han, 2015; Tan et al., 2015).

Another potential confound in studies supporting WM guidance effects is that participants might have strategically attended to the memory-matching distractor for the purpose of bolstering their WM representation (i.e., the “perceptual resampling” hypothesis; Woodman & Luck, 2007). For example, while the previously described study by Downing (2000) revealed an attentional bias toward memory-matching locations, participants could have deliberately attended to the memory-matching item in order to refresh their memory for it. However, this potential confound has been refuted in several ways. One study used images of real-world objects to demonstrate that memory-based attentional guidance generalized across category exemplars (i.e., maintaining the image

of a cat in WM resulted in equally strong attentional biasing toward the same image *and* toward a new image of a different cat; Balani, Soto, & Humphreys, 2010). WM guidance effects were unlikely to have resulted from deliberately attending to the different category exemplar, as it would not have benefitted (and perhaps even disrupted) the specific memory representation for the original exemplar. In other studies, researchers manipulated the dual-task setup such that after the initial WM cue, participants responded to *either* a visual search task or a WM probe. Because memory was never probed after the search task, there was no incentive to attend to memory-matching items simply for the purpose of refreshing the memory representation; resulting WM guidance effects were thus attributed to automatic attentional capture (Kiyonaga et al., 2012; Kiyonaga, Korb, Lucas, Soto, & Egner, 2014; Soto, Greene, Kiyonaga, Rosenthal, & Egner, 2012a).

1.4 The Limits of Attentional Guidance by Working Memory

While many studies have demonstrated that the active maintenance of a visual item or feature in WM is sufficient for inducing an attentional bias toward memory-matching contents in the visual field (see Olivers et al., 2011; Soto et al., 2008; Woodman et al., 2013), another set of studies using very similar designs have found no effect of memory-based attentional guidance (e.g., Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Peters et al., 2009), or even a reversal of the effect (e.g., Woodman & Luck, 2007). In Woodman and Luck (2007), participants remembered a color for a later memory test, and during the delay performed a visual search for an outline box with a gap at the top or bottom (i.e., Landolt C stimuli) among distractor boxes with gaps on

the left or right. When the WM color reappeared in the search array as a distractor, response times were either unaffected or even faster, indicating that participants could strategically avoid memory-matching items if they were unhelpful for the current task (Woodman & Luck, 2007). Thus, some researchers have argued that WM content can be flexibly and strategically employed as either a target template, or conversely a “template for rejection” (Woodman et al., 2013; Woodman & Luck, 2007; see also Arita, Carlisle, & Woodman, 2012; Beck & Hollingworth, 2015). More recent studies have indeed found evidence for both automatic and strategic components in memory-based attentional guidance (Carlisle & Woodman, 2011a; Kiyonaga et al., 2012).

Due to these inconsistencies in the literature, much of the early research on memory-based attentional guidance focused on the conditions under which WM does or does not bias selective attention. Many studies have attributed discrepancies to methodological details, such as the difficulty of the search task (Han & Kim, 2009; Olivers, 2009), memory load (Soto & Humphreys, 2008; Zhang, Zhang, Huang, Kong, & Wang, 2011), articulatory suppression (Olivers, 2009; Soto & Humphreys, 2008), specific stimulus attributes (Carlisle & Woodman, 2013; Zhang et al., 2010), temporal demands (Dalvit & Eimer, 2011; Dombrowe, Olivers, & Donk, 2010; Tsvetanov, Arvanitis, & Humphreys, 2012), or the size of the observer’s attentional focus (Hernández, Costa, & Humphreys, 2010; Pan & Soto, 2010). One important factor that consistently affects memory-based attentional guidance is how often the search target is repeated across trials (Olivers, 2009; Woodman et al., 2007). In many of the dual-task

paradigms demonstrating WM guidance effects, while the memory item changed from trial to trial, the visual search target remained constant (e.g., the tilted line target among vertical line distractors; Soto et al., 2005); however, when both the WM item and search target varied across trials, the deployment of attention no longer favored memory-matching items (Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Olivers, 2009; Peters et al., 2009; Woodman et al., 2007).

The finding that attentional biasing by WM depends on whether the search target is repeated has motivated an influential model of WM that accounts for multiple “activation” states of representations within WM (Olivers et al., 2011). Specifically, a single, active WM representation acts as an attentional template (even when it is not an intentional target template) and biases attentional processes, whereas other WM representations are relegated to accessory states that are still remembered but have little influence on visual attention (Olivers et al., 2011; van Moorselaar, Theeuwes, & Olivers, 2014b; but see Beck, Hollingworth, & Luck, 2012). Thus, in dual-task paradigms in which the search target is variably mapped across trials, participants must maintain a highly active template for the search target in WM, resulting in the suppression of accessory memory items (including the WM item) and the lack of attentional guidance by those accessory items (Olivers et al., 2011). In contrast, when the search target is consistently mapped and repeated across trials, the search template representation may be offloaded to a less active memory representation (Carlisle, Arita, Pardo, & Woodman,

2011; Gunseli, Olivers, & Meeter, 2014b; 2015), allowing the active WM representation of that trial's memory cue to bias attentional selection.

In support of this model, a recent study used a dual-task paradigm combining an n -back WM task with unpredictable, intervening visual search tasks (Greene, Kennedy, & Soto, 2015). In the 2-back condition, participants compared the color in the current trial to the color presented two trials previously—this required the maintenance of two colors in WM (from both trial $n-1$ and trial $n-2$), such that participants had to dynamically update on each trial which color was more task-relevant (i.e., from trial $n-2$) and which was temporarily less task-relevant (i.e., from trial $n-1$), at least until the next trial. As with previous studies (e.g., Soto et al., 2005), the trial $n-1$ or trial $n-2$ color could reappear in the search display at the location of a target or a distractor. In the 2-back condition, validity effects emerged only for the trial $n-2$ memory color and not the trial $n-1$ color, indicating that only the more active and task-relevant WM representation (i.e., the trial $n-2$ color) biased selective attention. This was especially striking because the trial $n-1$ color was only temporarily suppressed for that single trial, updating to an active trial $n-2$ representation for the subsequent trial. Thus, different activation states of representations within WM can be adaptively modulated across short timescales, resulting in equally dynamic guidance effects on selective attention (Greene et al., 2015; see also van Moorselaar, Olivers, Theeuwes, Lamme, & Sligte, 2015).

Given a model of WM in which representations can adopt different levels of activation or robustness (see LaRocque, Lewis-Peacock, & Postle, 2014; Oberauer, 2002;

Olivers et al., 2011), less is known about how various activation states within WM are instantiated. In other words, what happens to the memory representations to induce this distinction between active and accessory status? One possibility is that the more activated WM item is strengthened in a way that suppresses other accessory items or allows them to deteriorate (e.g., Matsukura, Luck, & Vecera, 2007; van Moorselaar et al., 2015; but see Rerko & Oberauer, 2013), but the specifics as to how suppression or deterioration occurs are uncertain (e.g., Hollingworth & Hwang, 2013; van Moorselaar et al., 2014b). Neurally, one theory is that prefrontal areas of the brain are involved in functionally separating active, immediately task-relevant information from more suppressed, accessory information; the neural populations that subserve active representations are better able to interface with lower visual areas, impacting attentional selection via feedback channels (Olivers et al., 2011). Alternatively, a suppressed WM representation may be offloaded to cortical regions or neural storage mechanisms (e.g., weight-based synaptic modifications; see LaRocque et al., 2014) that subserve long-term memory, such that the offloaded representation cannot interact with lower visual areas and impact visual processing; this memory representation can then be retrieved again when it becomes immediately task-relevant (LaRocque et al., 2014; Lewis-Peacock & Postle, 2008; Lewis-Peacock, Drysdale, & Oberauer, 2012).

1.5 Aims of the Present Thesis

The remaining chapters of this thesis present research that delineates mechanisms and boundary conditions of the interaction between WM and attention—focusing

primarily on (unintentional) memory-based attentional guidance. Chapters 2 and 3 aim to situate the contents of WM as a distinct influence on attentional selection that functions independently of traditional bottom-up and top-down factors. Thus, Chapter 2 examines how WM representations and bottom-up physical salience interact to guide attention, while Chapter 3 investigates how top-down strategic control modulates WM guidance. Moreover, Chapter 3 presents Bayesian model comparison techniques to show that top-down modulation of WM (and perhaps whether WM guides attention) depends on the functional state of the WM representation itself. This question of why WM does or does not bias attentional deployment across different paradigms is also the focus of Chapter 4, from an individual-differences perspective. Chapter 5 offers novel evidence that the link between WM content and attentional selection is highly reliable, such that multivariate pattern classification techniques can decode what an observer is holding in mind based on behavioral WM guidance effects, both within and across individuals. Chapter 6 presents a broader perspective on the bidirectional effects of concurrent processing between WM and attention by examining competitive tradeoffs between the two domains, providing both behavioral and neural support for the hypothesis that WM maintenance is accomplished via attentional mechanisms. Finally, Chapter 7 concludes the thesis by summarizing and synthesizing these findings, and offering commentary on the current understanding of memory-based attentional guidance, and more broadly, on interactions between WM and attention.

2. Working Memory Guidance Overrides Bottom-Up Salience Cues

The initial empirical chapters focus on situating memory-based attentional guidance with respect to the traditional dichotomy of bottom-up and top-down factors on attentional control. The current chapter presents a behavioral study that directly pits the influence of WM content against the influence of bottom-up visual salience. Participants held an item in WM while simultaneously conducting a multiple-target visual search; when there were two search targets, one was high-salience (easier to find) and one was low-salience (harder to find). Importantly, the memory item could match one or neither of the two search targets. The current findings demonstrate that the impact of WM on visual attention is strong enough to modulate, and even override, a baseline salience effect, illustrating the influential role of WM in guiding visual attention, even in the face of competing bottom-up salience cues. I also present an argument for why WM guidance in these experiments was not attributable to solely top-down strategy. The data presented in this chapter have been published as a research article in the *Journal of Experimental Psychology: Human Perception & Performance* (Dowd & Mitroff, 2013).

2.1 Introduction

At any given moment, the human visual system receives more input than it can process, necessitating attentional mechanisms that filter and select a subset of information for further processing. An important question to consider is what determines the filtering and selection such that certain items are processed more so than others. In general,

attentional selection is thought to be influenced by bottom-up cues (e.g., when attention is captured by an item's physical distinctiveness) and top-down cues (e.g., when attention is directed toward a task-relevant location). Several theoretical models (e.g., Treisman, 1986; Wolfe, 2007) embrace the idea that visual attention is guided by a reciprocal interaction between both bottom-up and top-down factors, but have recently been criticized for failing to consider other factors, such as memory (Awh et al., 2012; Hutchinson & Turk-Browne, 2012). Furthermore, while there has been substantial research investigating the contributions of bottom-up and top-down factors on visual selection (e.g., Theeuwes, 2010), less work has focused on the interactions between these various influences (e.g., Awh et al., 2012; McMains & Kastner, 2011; Soto, Humphreys, & Heinke, 2006). To better understand the dynamics of attentional guidance, the current study examines the interplay between two specific factors: bottom-up visual salience and the contents of WM.

In terms of bottom-up attentional selection, visual salience is often defined as differences in visual cues (e.g., luminance, color, motion, orientation, or depth) between an item and the rest of the visual field (e.g., Nothdurft, 2002; Theeuwes, 2010). Neurophysiological studies in monkeys have demonstrated that salience is represented neurally by retinal ganglion cells that emphasize discontinuities (Treue, 2003) and cortical V4 cells that respond preferentially to the highest contrast within their receptive fields (Reynolds & Desimone, 2003). In humans, behavioral evidence has also shown that local contrast in at least one feature dimension—such as color, form, or luminance—

in the visual field can capture attention (Nothdurft, 1993; 2002; Turatto & Galfano, 2000). Furthermore, several computational models of visual attention have successfully implemented maps of visual salience values in guiding bottom-up control of attention (e.g., Harel, Koch, & Perona, 2006; Itti & Koch, 2000; Parkhurst, Law, & Niebur, 2002).

Previous research has also established a role for the contents of WM in guiding attention (see Olivers et al., 2011; Soto et al., 2008). One theory of task-relevant attentional guidance posits that active maintenance of a search template in WM biases attention toward template-matching items in the visual field (Desimone & Duncan, 1995); however, a stronger interpretation of this hypothesis claims that maintaining *any* visual information in WM should also bias attention to memory-matching visual items. To test this hypothesis, several studies have used a dual-task paradigm in which participants are asked to remember an item (e.g., a colored shape) while performing an intervening visual search task (e.g., Olivers et al., 2006; Soto et al., 2005; Woodman & Luck, 2007). Critically, the memory item might reappear in the search display at the location of the target or the location of an irrelevant distractor. Compared to a neutral condition, in which the memory item does not match any location in the array, search times are faster when the memory item matches the target location and slower when the memory item matches a distractor location (e.g., Downing, 2000; Olivers et al., 2006; Soto et al., 2005). These results have been replicated with eye-tracking, with participants making more first saccades to the location matching the memory item (Olivers et al., 2006; Soto et al., 2005).

An important consideration in the memory-based attentional guidance literature is whether attentional effects are driven by WM or by visual priming, such that mere exposure to the ‘memory’ item, without a memory requirement, would be sufficient to bias the deployment of attention. Previous studies (Downing, 2000, Experiment 3; Soto et al., 2005, Experiment 3) have addressed this possibility by including a control condition in which participants are presented with an item before a search but know they will not be tested in a subsequent memory task; neither study found attentional guidance effects when there was no requirement to maintain the item in WM. In another study, the order of tasks was changed such that the memory test was administered *prior* to the search task, effectively ‘priming’ the item representation twice (once during the memory item presentation and once during the memory item test; Olivers et al., 2006, Experiment 5). Even so, no attentional effects were found, suggesting that after the memory test was completed, the memory item was released from WM and thus lost its effect on guiding visual attention (Olivers et al., 2006). These previous studies strongly suggest that attentional guidance effects are driven by active WM maintenance and not by simply priming low-level features in the search array, and the current study extended and confirmed this conclusion.

2.2 Current Goals

A critical issue is whether there is an interaction between bottom-up salience and WM on guiding visual attention. Soto et al. (2006) approached this question by measuring WM guidance of attention in search tasks that contained a more or less salient

target. Search times were faster when the memory cue matched the target and slower when the memory cue matched a distractor, but performance benefits and costs were amplified for higher salience targets compared to lower salience targets. Furthermore, when search arrays included a ‘pop-out’ target (i.e., a target with a flat search slope over increasing set sizes; Treisman, 1986), already-efficient search times were still faster when the memory cue matched the target. Thus, WM cues are able to enhance visual salience cues, resulting in more efficient attentional guidance. While the above shows that WM contents can additively enhance the effects of visual salience, it remains unknown whether WM cues can in fact override the attentional priority of visual salience. Such results would argue for an important and distinct role for WM representations in guiding attention.

The current study examines bottom-up visual salience and WM cues as competing forces, specifically looking at how WM guidance might enhance detection of a less salient target. Three experiments investigated whether WM can modulate the detection of more salient and less salient targets within the same search array. In Experiments 1A and 1B, we used a dual-task procedure that combined a WM task with an intervening search task. However, to align our findings with previous visual search literature, our search task used larger search arrays (i.e., 25–35 items) with less discriminable search stimuli (e.g., Wolfe, 1998; Wolfe & DiMase, 2003). Critically, we adopted a multiple-target search paradigm, which can provide a different perspective on search dynamics (e.g., Fleck, Samei, & Mitroff, 2010; Horowitz & Wolfe, 2001). Our search arrays could contain up to

two targets, one more salient and one less salient, such that the influence of WM could be directly pitted against effects of visual salience. Experiment 2 amplified visual salience effects to gauge the limits of WM cues on attentional guidance. Experiment 3 tested whether search effects were attributable to WM or to visual priming alone.

2.3 General Method

All three experiments in this study used a similar paradigm, detailed here. Any differences from this paradigm are noted for each experiment.

2.3.1 Apparatus

All experiments were conducted on a Dell Dimension E510/520 computer running Windows XP and were programmed in Matlab using Psychophysics Toolbox, Version 3.0 (Brainard, 1997). Participants viewed the experimental displays on a ViewSonic G90f CRT monitor with a refresh rate of 75 Hz at an approximate distance of 60 cm.

2.3.2 Stimuli

Based on previous studies (e.g., Cain & Mitroff, 2013; Cain, Dunsmoor, LaBar, & Mitroff, 2011; Fleck et al., 2010), the stimuli used here were ‘T’s and pseudo-‘L’s. Each stimulus item consisted of two perpendicular lines slightly separated from one other (stroke width = 0.3° , subtending $1.0^\circ \times 1.0^\circ$ total), with target Ts having a crossbar aligned directly in the middle and with distractor Ls having the crossbar slid at variable distances away from the center (Figure 1).

All stimuli were presented with varying luminance (i.e., shades of gray) on a white background. For the purposes of these experiments, visual salience is operationally defined as the contrast in luminance between an item and the rest of the array. Target Ts were presented at one or both of two visibility levels: high salience (range = 65–75% black) or low salience (range = 30–40% black). Distractor Ls were presented at varying shades of gray, which ranged between the low-salience and high-salience shades for each trial. For example, if on a given trial the high-salience T was 73% black and the low-salience T was 36% black, then distractors were all less than 73% black but greater than 36% black. If a trial had only one target, for example only the 73% black high-salience T, then distractors were all less than 73% black but greater than 30% black. This setup made high-salience Ts relatively easy to detect, and low-salience Ts more difficult to detect. All stimuli were also presented at varying orientations (range = 6–356° in steps of 10°). Each stimulus was placed with a slight spatial jitter within randomly selected cells of an invisible 8 × 7 grid subtending 19.4° × 14.8°. Items did not overlap.

2.3.3 Procedure

Each trial began with the presentation of a central cross (stroke width = 0.2°, subtending 0.8° × 0.8°) for 500 ms (Figure 1). The cross was followed by a blank screen for 300 ms, then by the memory item, a single T of random luminance (range = 30%–75% black) and orientation (range = 6–356° in steps of 10°), which was presented at the center of the screen for 500 ms. Participants were instructed to remember both the color and orientation of the memory item.

The memory item was followed by a blank screen for 300 ms, then by the search array, which consisted of 25 items. There were 1 or 2 target Ts to find within each display. Participants used the mouse to click on each detected target item, regardless of luminance or orientation, then pressed the space bar to complete the search. Each search had a time limit of 15 s, after which no further clicks were accepted and a message was displayed encouraging participants to try to finish searching and press the space bar before time elapsed on subsequent trials. Responses made prior to the time limit were recorded and analyzed even if the space bar was not pressed.

The search array was followed by a blank screen for 300 ms, then by a probe item that appeared in the center of the screen until response. The probe item was a single T that was either identical to the memory item in both dimensions of luminance and orientation, or different on one or both dimensions. Participants reported via keypress ('s' or 'd', respectively) whether the probe item was exactly the same as the previously seen memory item or different on either or both features. Feedback tones were used to indicate whether the memory probe response was correct (a short, high tone) or incorrect (a long, low tone) and were meant to encourage participants to remember the memory item.

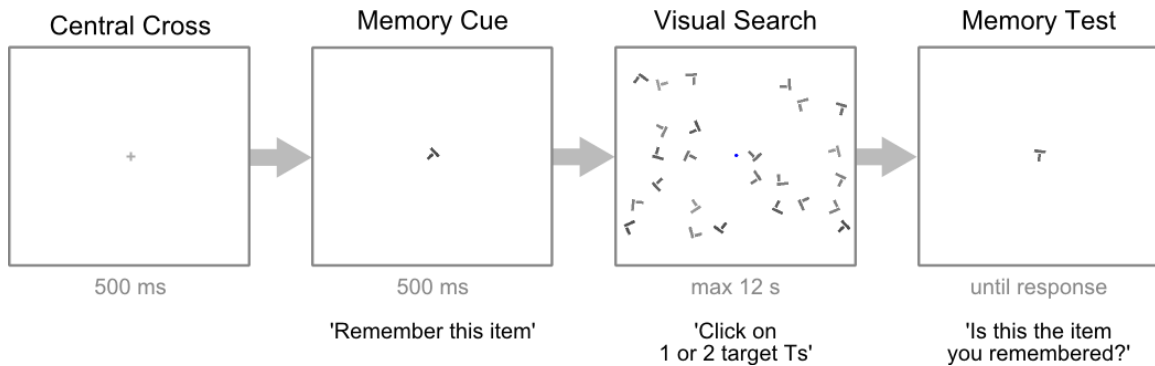


Figure 1: Example trial sequence of Experiment 1A showing a memory task with an intervening search task. In this example, the search task includes two targets, a high-salience target and a low-salience target. This trial depicts a valid, high-match trial because the memory cue matches the high-salience target in both luminance and orientation. The memory probe differs from the memory cue in this example.

2.3.4 Trial Conditions

Trials were classified across two factors: validity (valid, neutral) and search type (single-target low-salience, single-target high-salience, dual-target low- and high-salience). The two validity conditions were defined by the relationship of the search items to the WM cue. In valid trials, the memory item matched a target within the search display; if there was a match, the target was identical to the memory item on both dimensions of luminance and orientation. In neutral trials, the memory item matched neither a target nor any of the distractors on luminance or orientation. There was never an invalid condition, in which the memory item matched a distractor. There were three search types, each defined by the number and salience of the target Ts presented within the array of distractor Ls. The search array could include a single target, either a high-salience or a low-salience, or two targets, both a high-salience and a low-salience. For

valid, dual-target trials, the memory item could match either the high-saliency target ('high-match') or the low-saliency target ('low-match'). In total, there were seven different trial type combinations (Table 1).

Each experiment began with a practice phase of 12 trials that was matched to the trial-type frequency of the rest of the experiment. During practice, immediate search feedback was provided on any false positive identification or missed targets, in addition to memory feedback tones. The experimental phase consisted of 240 trials, which were counterbalanced by trial types according to validity, number of targets, and saliency of targets (Table 1). Neutral trials were equally as frequent as valid trials. No search feedback was provided during the experimental trials.

Table 1: Trial Types by Number of Search Targets, Saliency of Search Targets, and Memory-to-Search Target Relationship

Number of targets in visual search	Saliency of search targets	Memory-to-search target relationship	Memory-to-search target match condition
Single	High	Valid	High-match (20)
		Neutral	Neutral (20)
	Low	Valid	Low-match (20)
		Neutral	Neutral (20)
Dual	High and low	Valid	High-match (40)
		Valid	Low-match (40)
		Neutral	Neutral (80)

Note. Memory-to-search target validity condition refers to whether the memory cue matched a search target, and memory-to-search target match condition refers to which target the memory cue matched. Trial counts are presented in parentheses.

2.3.5 Data Analysis

Our primary measure of interest was search accuracy, as it offers a more compatible measure with previous studies (e.g., Soto et al., 2005) than response times in

the current design. The current dual-task memory and search paradigm differs from others in that participants respond with mouse clicks on targets, as opposed to forced-choice responses, and the search size arrays employed are much larger, with as many as 35 items in the current tasks compared to a maximum of eight items in Soto et al. (2005). Nevertheless, the response time data are consistent with the accuracy results and are presented in Table 2 and in Appendix A.

In all experiments reported here, a one-way analysis of variance (ANOVA) was conducted to compare first hit percentages of the high-salience target across the three memory-to-search-target match conditions (high-match, low-match, neutral).¹ Paired *t*-tests were also conducted, comparing first hit percentages of the high-salience target to first hit percentages of the low-salience target within each of the three match conditions. To analyze search response times, a two-way ANOVA was conducted across factors of memory-to-search-target relationship (i.e., neutral, valid) and target salience. All statistics are reported in Table 2.

¹ A two-way ANOVA with the main factors of memory-to-target match condition and salience was not administered because the first-hit percentage values within each condition were dependent. Because the first hit was either high salience or low salience (or neither was found at all), such an analysis over condition would sample the same data twice.

Table 2: Overall Accuracy and Search Times to Find a Target First in Dual-Target Searches

	Proportion of first hits		Average time in seconds to first hit		Statistical tests
	High-salience	Low-salience	High-salience	Low-salience	
Experiment 1A ($n = 10$)					
Neutral	.555 (.06)	.445 (.06)	4.15 (1.05)	4.20 (1.01)	$t(9) = 0.85$ $p = .849$
High-match	.741 (.11)	.259 (.11)	3.37 (.79)	4.40 (1.05)	$t(9) = 2.76$ $p = .022$
Low-match	.282 (.16)	.718 (.16)	4.12 (1.05)	3.09 (.57)	$t(9) = 3.27$ $p = .010$ $F(1, 36) = 11.57$ $p = .002$
Experiment 1B ($n = 12$)					
Neutral	.568 (.09)	.432 (.09)	5.24 (.92)	5.27 (1.4)	$t(11) = 0.09$ $p = .927$
High-match	.766 (.15)	.224 (.15)	4.09 (1.10)	5.80 (2.14)	$t(11) = 2.54$ $p = .028$
Low-match	.363 (.20)	.637 (.20)	5.61 (1.61)	4.36 (1.17)	$t(11) = 2.63$ $p = .023$ $F(1, 44) = 9.47$ $p = .004$

Experiment 2 ($n = 10$)						
Neutral	.919 (.07)	.081 (.07)	$t(9) = 18.41$ $p < .001$	2.13 (.55)	3.72 (.87)	$t(9) = 6.61$ $p < .001$
High-match	.985 (.03)	.015 (.03)	$t(9) = 57.07$ $p < .001$	1.72 (.38)	5.02 (6.75)*	$t(9) = 1.09$ $p = .302$
Low-match	.720 (.16)	.280 (.16)	$t(9) = 4.41$ $p = .002$	2.70 (.95)	3.68 (.71)	$t(9) = 2.75$ $p = .022$
			$F(2, 27) = 12.78$ $p < .001$			$F(1, 36) = 1.164$ $p = .288$
Experiment 3 ($n = 10$)						
Neutral	.651 (.10)	.349 (.10)	$t(9) = 4.97$ $p < .001$	5.38 (.77)	5.28 (.71)	$t(9) = .36$ $p = .726$
High-match	.712 (.12)	.288 (.12)	$t(9) = 5.82$ $p < .001$	4.82 (.71)	4.84 (.90)	$t(9) = -.11$ $p = .915$
Low-match	.613 (.10)	.387 (.10)	$t(9) = 3.73$ $p = .005$	5.31 (.87)	4.69 (.94)	$t(9) = 1.66$ $p = .131$
			$F(2, 27) = 2.37$ $p = .113$			$F(1, 36) = 3.33$ $p = .027$

Note. Standard deviations are presented in parentheses. Test statistics represent paired t -tests for first hits and search times between high-salience and low-salience targets within each match condition, and ANOVA test statistics for the main effect of match condition across first hits and the main effect of validity across search times.

* Seven of ten participants never found the low-salience target first in the high-match condition.

2.4 Experiment 1A

A dual-task paradigm for WM and visual search was tested to directly contrast the effects of bottom-up visual salience versus the effects of WM maintenance on the deployment of visual attention. Twelve Duke University students (six male, ages 18–21 years) participated in the study, and all signed informed consent in accordance with Duke University’s Institutional Review Board. Data from two participants were excluded from analysis, one for having a false alarm rate more than 2.5 standard deviations from the mean across all participants, and one for being more than 2.5 standard deviations from the mean across all participants for search accuracy in the neutral condition.

2.4.1 Results

Memory task. Performance on the memory task was above chance ($M = 69.4\%$, $SD = 7.7\%$), $t(9) = 8.00$, $p < .001$, and was significantly better for valid trials ($M = 71.3\%$, $SD = 9.2\%$) than for neutral trials ($M = 67.6\%$, $SD = 6.4\%$), $t(9) = 2.88$, $p = .018$. Validity benefits likely resulted from an updating of the memory representation due to the re-processing of the memory item if the valid target item was found.

Visual search task. Analyses for the search task were focused on accuracy performance for dual-target search trials; see Table 2 for dual-target response time data and Appendix A for single-target search trial data. False alarms (defined as mouse clicks that were not on target items; i.e., on distractor items or on empty space) accounted for 2.0% of all clicks, and those trials were excluded from all subsequent analyses. Participants exceeded the search time limit on 0.9% of trials. The critical measure was

“first hits,” or which of the two targets was found first. Previous research using a dual-target search task has found that given a more salient and a less salient target, the more salient target is not only more likely to be found, but also more likely to be found first (Cain & Mitroff, 2013; Fleck et al., 2010). The percentage of high-salience targets found first was compared to the percentage of low-salience targets found first, separately for the neutral and valid conditions (Table 2).

The critical comparison was whether the pattern of first hit performance between conditions was affected by the contents of WM. Even in the presence of a baseline salience effect from the neutral condition, in which high-salience targets were more likely to be found first, $t(9) = 3.14, p = .012$, there was a significant effect of memory-to-search-target match condition (i.e., whether the memory item matched the high-salience, low-salience, or neither target) on first hit percentages for high-salience targets, $F(2, 27) = 39.86, p < .001$. That is, attention was guided to the target that matched the item being held in WM (Figure 2). Guidance by WM either enhanced the effects of salience (for the high-salience match) or overrode them (for the low-salience match). Search time analyses also revealed that validity had a significant effect on how quickly participants found the first target, $F(1, 36) = 11.57, p = .002$ (Table 2).

2.4.2 Discussion

In this dual-task paradigm, WM cues were strong enough to not only enhance but also override attentional priority of salience cues. Dual-target search performances for each memory-to-target match condition highlight different attentional guidance effects. In

the neutral condition, without guidance from the contents of WM, participants were more likely to find the high-saliency target first. This demonstrated a baseline effect of visual saliency, in which attention was biased toward the target with greatest local contrast. In the high-match condition, holding a matching high-saliency item in WM enhanced the likelihood of finding the high-saliency target first, replicating the main result of Soto et al. (2006). However, in the low-match condition, holding a matching low-saliency item in WM guided attention toward the less salient target, reversing the predicted pattern of bias by visual saliency. This result supports the idea that visual attention can be guided by a competitive balance between bottom-up visual saliency and the contents of WM.




Figure 2: Data from Experiment 1A demonstrate that WM guidance both enhances and overrides a baseline saliency effect. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors.

2.5 Experiment 1B

In Experiment 1A, holding a less visually salient item in WM enhanced the detection of that particular item, overriding attentional priority of a more visually salient item. The strength of this guidance by the contents of WM is somewhat surprising, so Experiment 1B was intended to replicate the effect with a more complex search task. Specifically, this experiment was identical to Experiment 1A except with the search array size increased to 35 stimuli. Seventeen Duke University students (seven male, ages 18–24 years) participated in the study. Data from five participants were excluded from analysis for having false alarm rates more than 2.5 standard deviations from the mean across all participants.

2.5.1 Results

Memory task. Performance on the memory task was above chance ($M = 76.2\%$, $SD = 8.2\%$), $t(11) = 11.04$, $p < .001$) and was significantly better for valid trials ($M = 79.0\%$, $SD = 7.1\%$) than for neutral trials ($M = 73.3\%$, $SD = 9.6\%$), $t(11) = 3.70$, $p = .003$).

Visual search task. Table 2 summarizes first hit performance of the twelve participants in terms of accuracy and search times for dual-target search trials. False alarms accounted for 2.5% of all clicks, and participants exceeded the search time limit on 4.5% of trials. As in Experiment 1A, the critical measure of performance was first hits. Experiment 1B replicated the two key effects: a baseline effect of visual salience in the neutral condition, $t(11) = 2.69$, $p = .021$, and an overall effect of WM guidance, $F(2,$

33) = 21.63, $p < .001$, in which attention was guided to the target that matched the item being held in WM, regardless of visual salience (Figure 3). Search time analyses also supported this result: participants were significantly faster to find the target that matched the contents of WM, $F(1, 44) = 9.47, p = .004$. Response times were overall longer than in Experiment 1A, which represents the added complexity introduced by increasing the set size from 25 to 35 items.

2.5.2 Discussion

Experiment 1B replicated the findings of Experiment 1A, showing that even in a more complex search task, performance in dual-target search was influenced by an interaction between effects of visual salience and WM guidance. In the neutral condition, participants were more likely to find the high-salience target first, revealing a baseline effect of visual salience. However, when the memory cue matched a search target, participants were more likely to find that particular item first, regardless of salience. This indicates that WM cues are able to override attentional priority of a more salient target and guide attention to a less salient target first.




Figure 3: Data from Experiment 1B replicate the WM guidance effects from Experiment 1A. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors. Note the values may not total to 100% due to rounding.

2.6 Experiment 2

The results of Experiments 1A and 1B demonstrate that WM cues are sufficiently strong to enhance or even override attentional guidance toward targets of high salience. But what happens when the local contrast of the high-salience target is increased even more, effectively strengthening the prioritizing effect of visual salience? To further examine the strength of WM cues in guiding attention, stimulus salience was manipulated in Experiment 2 such that the high-salience target was extremely easy to detect. The stimuli array size and procedure were identical to those in Experiment 1B. Eleven Duke University students (seven male, ages 18–21 years) participated in the study. Data from one participant were excluded for a false alarm rate more than 2.5 standard deviations from the mean across all participants.

2.6.1 Verification of “Pop-Out”-Like Search

Stimulus salience was manipulated such that the high-salience, darker Ts were much easier to detect. As in Experiments 1A and 1B, target Ts were presented at one or both of two visibility levels: high salience (range = 65–75% black) or low salience (range = 30–40% black). However, distractor Ls were presented at varying shades of gray in a muted range (30–52% black) closer to the low-salience range. Thus, high-salience Ts were so distinctive as to seemingly ‘pop-out’ from the search array.

The ‘pop-out’-like character of the high-salience T in this search array was verified by comparing simple search performance across varying set sizes. An independent group of 10 additional participants (seven male, ages 18–30) was recruited for this verification task. Participants were instructed to judge whether a single target T was present or absent in intermixed search arrays of 8, 12, 24, and 36 stimuli. A single target T was present on 50% of the trials, and was either high-salience or low-salience in a field of muted distractors (all salience values match those used in Experiment 2). Search slopes were calculated as the slope of the linear regression line that best fit search times across set sizes. When the target T was high-salience, a nearly flat search slope was generated ($M = 9.5$ ms/item, $SD = 1.2$ ms/item); in contrast, when the target T was low-salience, the search slope was highly positive ($M = 100.1$ ms/item, $SD = 71.2$ ms/item). The search slopes between the two target types were significantly different, $t(9) = 3.72$, $p = .005$, confirming that the high-salience, ‘pop-out’-like target did produce highly efficient search (Figure 4).




Figure 4: Data from Experiment 2 verify the “pop-out” nature of the high-salience target. Search times as a function of set size when searching for low-salience versus high-salience targets among distractors, with slope values provided. Error bars represent standard errors.

2.6.2 Results

Memory task. Performance on the memory task was above chance ($M = 78.7\%$, $SD = 4.9\%$), $t(9) = 18.36$, $p < .001$, and was significantly better for valid trials ($M = 81.3\%$, $SD = 5.2\%$) than for neutral trials ($M = 76.0\%$, $SD = 6.0\%$), $t(9) = 3.25$, $p = .010$.

Visual search task. Table 2 summarizes first hit performance of the ten participants in terms of accuracy and search times for dual-target search trials. False alarms accounted for 0.9% of all clicks, and participants exceeded the search time limit on 0.3% of trials. As in Experiments 1A and 1B, the critical measure of performance was proportion of first hits. A one-way ANOVA revealed a significant effect of memory-to-target match condition, $F(2, 27) = 12.78$, $p < .001$. In the neutral condition, participants

were significantly more likely to find the high-salience target first than the low-salience target first, $t(9) = 18.41, p < .001$. The salience effect was extremely strong, with an advantage of 83.74% for finding the high-salience target first compared to finding the low-salience target first, due to the pop-out-like nature of the high-salience target. Nevertheless, this salience effect was significantly enhanced by WM guidance in the high-match condition, $t(9) = 2.93, p = .017$, with participants finding the high-salience target first almost every trial, $t(9) = 57.07, p < .001$. Furthermore, a one-way ANOVA of search times revealed that the high-salience target was found significantly faster in the high-match condition than in neutral or low-match conditions, $F(2, 27) = 5.46, p = .010$. Crucially, in the low-match condition, although participants were still more likely to find the high-salience target first, $t(9) = 4.41, p = .002$, WM guidance toward the low-salience target significantly attenuated the proportion of finding the high-salience target first, $t(9) = 4.79, p < .001$ (Figure 5).

2.6.3 Discussion

In an already efficient search for a high-salience, pop-out-like target, WM guidance can still enhance search performance, replicating the effect by Soto et al. (2006). Furthermore, although WM guidance does not reverse the pattern of attentional priority of a pop-out-like target, WM cues can still significantly attenuate the effect of visual salience. When the low-salience item was held in WM, the matching low-salience target was more likely to be detected first than when it was not held in memory. This

indicates that WM guidance is strong enough to counteract even a highly robust bottom-up salience effect.




Figure 5: Data from Experiment 2 show that WM guidance modulates an amplified visual salience effect. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors.

2.7 Experiment 3

The aim of Experiment 3 was to test whether the attentional guidance effects observed in Experiments 1A, 1B, and 2 were driven by WM or by visual priming. While previous research has suggested that priming cannot account for attentional guidance such as that observed in the current study (e.g., Downing, 2000; Olivers et al., 2006; Soto et al., 2005), it is nevertheless important to demonstrate that merely presenting the ‘memory’ cue would not be sufficient to guide attention. This experiment was identical to Experiment 1B, except that it did not include the memory test. Specifically, participants were no longer required to hold the initial cue in WM for later comparison to the probe

item. Fourteen Duke University students (four male, ages 18–22 years) participated in the study. Data from four participants were excluded from analysis for reporting that they explicitly used the ‘memory’ cue to guide their search.

2.7.1 Results

Visual search task. Table 2 summarizes first hit performance of the ten participants in terms of accuracy and search times for dual-target search trials. False alarms accounted for 2.7% of all clicks, and participants exceeded the search time limit on 4.7% of trials. As in Experiments 1A, 1B, and 2, the critical measure of performance was first hits. Experiment 3 replicated the baseline effect of visual salience in the neutral condition, $t(9) = 4.97, p < .001$. However, contrary to a visual priming account, there was not an overall effect of WM guidance, $F(2, 27) = 2.37, p = .113$, such that attention was not significantly guided to the target that matched the cued item (Figure 6). In fact, there was a significant effect of visual salience within each match condition, such that the high-salience target was found first more than the low-salience target in every condition (Table 2).

2.7.2 Discussion

The results of Experiment 3 rule out a visual priming account of attentional guidance. Simply presenting the ‘memory’ cue before the visual search array, without any requirement to hold the cue in WM, failed to produce the attentional guidance effects found in Experiments 1A, 1B, and 2. These results support previous studies (e.g., Downing, 2000; Olivers et al., 2006; Soto et al., 2005) that have also found that priming

alone cannot explain attentional guidance effects. All together, these studies are strong indications that the currently observed attentional guidance effects were driven by WM and not by simply priming low-level features in the search array.




Figure 6: Data from Experiment 3 reveal a minimal influence of visual priming. Average rates of which target was correctly found first (“first hits”), as a function of match condition. Error bars represent standard errors.

2.8 General Discussion

The present study demonstrates that the deployment of attention in visual search is modulated by a competitive balance between bottom-up visual salience and WM cues. A previous study found that WM cues can enhance search performance for more salient items (Soto et al., 2006), and the current study specifically investigated the *strength* of WM guidance when contrasted with salience effects. WM guidance was found to not only be able to enhance, but even override, attentional priority of bottom-up visual salience, such that WM biases enhanced the detection of a less salient target. Experiments

1A and 1B confirmed that holding an item in WM significantly biased participants to find that specific item first, regardless of its salience. Experiment 2 further explored the relationship between WM and salience effects on attentional guidance—even in a pop-out-like search with a highly salient target, WM cues still enhanced the detection of a less salient target. Experiment 3 showed that these effects were not found when participants were asked to only observe the cue prior to search and not required to hold the item in WM. Together, these results support a distinct and influential role for WM in the deployment of visual attention, independent from bottom-up visual salience. Thus, these patterns of search performance support models of visual attention driven by reciprocal interactions between stimulus properties, volitional goals, and memory (Awh et al., 2012; Hutchinson & Turk-Browne, 2012).

Here we introduce a novel dual-task paradigm that provides a number of elements not present in similar tasks. First, the current task employs a greater number of stimuli, with as many as 35 items compared to a maximum of eight items in Soto et al. (2005). Second, the ‘T’ and ‘L’ stimuli used are also less discriminable than colored shapes or lines (Wolfe & DiMase, 2003), which makes the memory task and the visual search task more complicated. Third, the current search task features multiple targets, which increases the attentional, memory, and decision-making demands (Cain & Mitroff, 2013). Importantly, these changes to the typical dual-task paradigm did not affect the base effects of WM guidance on search, while allowing for novel questions to be addressed. Specifically, by adopting a search task with multiple targets, these experiments were able

to simultaneously present targets of unequal salience. This allowed for an evaluation of the relative strength of WM guidance, revealing that WM has a powerful influence over the deployment of visual attention. Furthermore, the current results demonstrate that the contents of WM can guide attention even in an effortful and relatively time-consuming search task. This paradigm could be applied in future studies to explore the potential distinction between a fast capture of attention by bottom-up salience and a longer-lived guidance of attention by WM.

The multiple-target nature of the current design offers a potentially powerful methodological tool. Multiple-target searches with targets of unequal salience are common in real-world searches, such as when radiologists search for medical abnormalities in radiological scans or when airport baggage screeners search for safety concerns. This paradigm could be used to explore the effects of WM guidance on real-world objects, as well as with regard to other definitions of salience, such as emotional salience. For example, in baggage X-rays, a water bottle might be more visually salient but less emotionally salient than a handgun. Future research using real-world objects or affective stimuli could simultaneously inform on the interaction between memory and attention in both cognitive psychology and real-world applications.

2.8.1 Working Memory Guidance or Strategy?

An ongoing debate in the field of visual attention revolves around the automaticity of WM guidance, or whether WM guidance is simply the result of top-down, volitional strategy (see Olivers et al., 2011; Soto et al., 2008; Woodman et al.,

2013). Thus, several studies of memory-based guidance have included a critical condition in which the WM cue reappeared only as an invalid distractor, such that attentional biasing by WM content was detrimental to search performance (e.g., Kiyonaga et al., 2012; Olivers et al., 2006; Soto et al., 2005). In these studies, WM effects were still found, even at a cost to performance, supporting the theory that WM guidance of attention is automatic (see Soto et al., 2008). However, in contrast, other studies have found no effect of memory-based attentional guidance (e.g., Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Peters et al., 2009), or even a reversal of the effect (e.g., Woodman & Luck, 2007), demonstrating a strategic avoidance of memory-matching distractors and supporting the flexible and voluntary nature of WM guidance (see Woodman et al., 2013).

Akin to the above hypotheses, the current data could fit one of two interpretations—the observed search performance differences could be explained by either a claim that attention is guided by WM cues or by a claim that participants used an overt strategy to complete the search task. That is, participants are quicker to find the target that matches the memory cue because they were guided by the contents of their WM or because they were explicitly, strategically searching for the memory cue. It is not

possible to definitively rule out either alternative¹, but additional data from the current study and from previous studies support an attentional guidance by WM interpretation.

From the current study, it is possible to inform which interpretation (WM guidance vs. explicit strategy) best explains the current data by considering how each theory predicts performance change after a target is found in search. A memory-based guidance hypothesis predicts that by holding a particular item in WM and then finding that matching target first, the memory representation would be strengthened by additional spatial location information and through a refreshing of the memory representation (Woodman & Luck, 2007). A strengthening of the memory representation would lead to it effectively using a greater proportion of memory resources (Bays & Husain, 2008). This more robust representation would thus interfere with finding the second, non-matching target (Cain & Mitroff, 2013). In terms of the current study, a WM guidance account would predict that participants should be *slower* to find a non-matching target after finding the matching target, compared to a neutral baseline. In contrast, an explicit strategy hypothesis posits that the participant would maintain an overt goal of finding the exact match of the memory cue. In neutral conditions in which the cue does not appear in the search array, participants would be hindered by this overt goal, since they would

¹ The dual-target nature of the current paradigm provides a novel manner to explore interactions between memory and attention, but also eliminates the possibility of including a condition of strategic avoidance, in which the memory cue matches only invalid locations (e.g., Arita et al., 2012; Woodman & Luck, 2007).

continue searching for the memory cue even after finding other targets. In high-match or low-match conditions, participants would find the matching cue and then know to discard that goal target (i.e., a template for rejection; Arita et al., 2012), allowing them to be relatively faster to find the second, non-matching target. In terms of the current study, an explicit strategy account would predict that participants should be *faster* to find a non-matching target after the matching target, compared to a neutral baseline. The data suggest a trend in which participants are slower to find a non-matching target after finding the matching target, compared to a neutral baseline, $t(63) = 1.94, p = .057$. These data are not conclusive, but support a memory-based guidance interpretation more than an explicit strategy hypothesis. Furthermore, as expanded on in Chapter 3, research has shown that while top-down goals can modulate attentional guidance by WM, strategy alone does not account for all WM effects (e.g., Carlisle & Woodman, 2011a; Kiyonaga et al., 2012). Together, these findings suggest that memory-based attentional guidance in the current study were not simply due to overt, top-down strategy.

2.8.2 Conclusion

The primary goal of the present study was to explore the interaction between attentional guidance by bottom-up visual salience and by WM representations. By demonstrating that the contents of WM can powerfully guide attention, even in the face of conflicting visual salience effects, this study illustrates the influential role of WM—distinct from either top-down or bottom-up labeling—in guiding visual attention.

Furthermore, this study also offers a novel paradigm that can be used to study specific interactions between the various influences on visual attention.

3. Top-Down Control Over Working Memory Guidance Operates Via Working Memory Accessibility

While Chapter 2 contrasted the influence of WM versus bottom-up visual salience on attentional deployment, this chapter focuses on how top-down strategic control interacts with WM guidance. In a dual-task paradigm combining WM and visual search, we manipulated the probabilities with which WM contents matched target or distractor locations. The influence of WM was modulated by top-down strategy, such that memory-based guidance was strategically suppressed or enhanced depending on whether WM contents were reliably hurtful or helpful for the search task. These data demonstrate that attentional selection can be driven by interactions between WM and strategic goals, again supporting the distinction between WM guidance and top-down control.

Critically, this study investigated how top-down control over memory-based guidance affected the WM representation itself, either in terms of representational quality (i.e., how precise is it?) or accessibility of WM content (i.e., how easily is it remembered or forgotten?). Using probabilistic model fitting and Bayesian model comparison techniques, we found that top-down control over memory-based guidance impacts the probability of remembering, but not the precision of, items in WM. These findings suggest that whether WM impacts attention may depend on distinct functional states in WM, which are in turn characterized by how easily an item is remembered or forgotten. These data were published as a research article in the journal *Visual Cognition* (Dowd, Kiyonaga, Beck, & Egner, 2015a).

3.1 Introduction

As reviewed in Chapter 1, the contents of WM guide visual attention toward memory-matching objects, as evidenced by dual-task WM and visual search paradigms in which search performance reveals benefits (i.e., faster search times) for valid trials and costs (i.e., slower search times) for invalid trials (see Olivers et al., 2011; Soto et al., 2008; Woodman et al., 2013). Intriguingly, recent studies have found that the WM-attention linkage can be strategically suppressed or enhanced, depending on whether WM contents are reliably hurtful or helpful for attentional task performance (Carlisle & Woodman, 2011a; Kiyonaga et al., 2012). By manipulating the probabilities of trial conditions, these studies demonstrated that participants could amplify the benefits of valid trials when the memory item reliably predicted the location of a target (i.e., even faster responses) and dampen the costs of invalid trials when the memory item was reliably associated with distractors (i.e., less slowed responses).

Top-down strategic control over attentional guidance by WM has been hypothesized to operate via some modulation of the WM representation itself (Kiyonaga et al., 2012); however, the exact way in which the maintenance of the memory item is altered to produce these strategic effects is presently unknown. Recent models of visual WM suggest that a limited pool of cognitive resources (or a limited number of resource slots) can be distributed to items in memory, such that WM representations might vary in quality, or precision (Bays & Husain, 2008; Zhang & Luck, 2008). By testing memory with continuous report tasks (Wilken & Ma, 2004), in which participants recall a target

from a continuous spectrum of responses, memory responses can be fit with a probabilistic mixture model (see Bays, Catalao, & Husain, 2009; Zhang & Luck, 2008) that dissociates performance into two main components: (1) the representational quality, or precision, with which a target is recalled, given that it is remembered, and (2) the accessibility of that representation, or the probability of remembering that target at all (as opposed to randomly guessing or incorrectly reporting another item from the memory array).

Based on this influential distinction between precision and accessibility, the current study investigated how strategic enhancement or suppression of the interaction between WM and attention affects the WM representation, examining whether strategic control of memory-based guidance impacts the quality of WM, the accessibility of WM, or both. Specifically, strategic control over how WM influences attention could alter the precision of WM representations, which may in turn determine performance. For example, memory-based capture could depend on the specificity of the match between WM contents and stimuli in the visual field; strategically enhancing or dampening representational quality could affect WM-attention interactions, such that a more precise WM representation might mean better performance in a matching task, whereas a less precise representation might lead to more behavioral variability in the task. Alternatively, strategic control over WM could change the accessibility of WM. For example, an “enhanced” (prioritized) WM item might be more easily accessed by WM retrieval processes and thus more likely to influence attention, whereas a suppressed WM item

might be stored in a manner that limits its retrieval from WM, and thus its impact on attention (e.g., Olivers et al., 2011). These processes could also work conjointly.

A number of recent studies speak to the question of how knowledge about the task-relevance of a memory item affects the quality of that memory representation. One study used a “directed-forgetting” paradigm, in which certain items in WM were retrospectively cued to be forgotten (i.e., task-irrelevant; Williams, Hong, Kang, Carlisle, & Woodman, 2013). The remaining items in WM (i.e., those that were not cued to be forgotten) had higher precision relative to conditions in which no items were directed to be forgotten, suggesting that the quality of memory representations could be flexibly and strategically controlled (see also Pertzov, Bays, Joseph, & Husain, 2013). In contrast, other studies found no differences in precision between retrospectively cued and uncued items, but rather only advantages in the probability of remembering the cued item (Hollingworth & Hwang, 2013; Murray, Nobre, Clark, Cravo, & Stokes, 2013; Souza, Rerko, Lin, & Oberauer, 2014).

Two studies have also found that the precision of a memory representation is not correlated to attentional capture by that matching item (Hollingworth & Hwang, 2013; van Moorselaar et al., 2014b). However, these previous studies did not manipulate the relevance of the memory item in relation to the subsequent attentional task. Hollingworth and Hwang (2013, Experiment 3) used a retrospective cue to indicate memory relevance before search, and the subsequent search task could contain either a singleton distractor that matched the uncued (memory-irrelevant) item or a distractor that was unrelated to

any feature in WM, but never the cued (memory-relevant) item. While relatively precise representations of the uncued item were *not* sufficient to drive attentional capture by uncued-matching distractors, the study design prevented the testing of whether precision for the cued item was related to attentional capture by cued-matching distractors. Van Moorselaar et al. (2014b, Experiment 3) also found no differences in precision for memory items that did or did not capture attention; however, this study presented the retro-cue *after* the search task, such that there was no manipulation of memory-relevance before search. Furthermore, neither study analyzed the probability of remembering the memory item in relation to attentional capture. In the current study, we manipulate the relevance of the memory item for the subsequent search task, in which the memory-matching item guides attention toward or away from the target, allowing us to examine how strategic modulation of that interaction affects the WM representation itself, both in terms of quality and accessibility.

Based on these findings from previous studies, it remains unclear how strategic control over memory-based attentional guidance would affect the WM representation itself. The disagreement over whether the task-relevance of a WM representation affects its quality and/or accessibility may be explained in several ways. For example, the data could simply be too variable across different task paradigms and different samples. There also may not have been a well-grounded, uniform statistical test applied to appropriately analyze the dataset. Specifically, these previous studies shared a common analytic approach: fitting probabilistic models to the data from each participant, then combining

parameter estimates across participants and comparing group means with t -tests or ANOVAs. While this seems like a straightforward way to analyze parameter estimates, this approach fails to account for the reliability of and correlations between parameter estimates of each subject (Kruschke, 2010; Oberfeld & Franke, 2012; Suchow, Brady, Fougne, & Alvarez, 2013). Moreover, an ANOVA is only capable of identifying a change in parameter values of a given model and is not capable of determining if an entirely different model (perhaps one with more parameters) provides a better description of the behavior in different task conditions. This is particularly problematic in behavioral studies as different subjects may employ a wide range of strategies, each of which might have a unique parameterization, making any simple comparison impossible.

To overcome these limitations, the current study employed Bayesian model comparison. While more computationally intensive, this approach allowed us to compare a large number of possible behavioral models to determine strategies employed by subjects and identify the various sources of suboptimal behavior (i.e., diminished quality vs. accessibility of the WM representation). In Bayesian statistics, for a particular model M with parameter vector θ_M , the marginal likelihood of the data from subject k , D^k , is calculated by integrating over the parameters of model M :

$$p(D^k | M) = \int p(D^k | \theta_M, M) p(\theta_M | M) d\theta_M \quad (1)$$

Bayesian model comparison relies on these marginal data likelihoods (i.e., the model evidences) to compute how likely a particular model is compared to other models.

For example, given either model $M = 1$ or $M = 2$, the posterior probability that $M = 1$ is the model used by subject k is:

$$p(M = 1 | D^k) = \frac{p(D^k | M = 1)}{p(D^k | M = 1)p(M = 1) + p(D^k | M = 2)p(M = 2)} \quad (2)$$

where $p(M = i)$ is the prior probability of model i . At the individual subject level, this approach allows us to compute the probability that a given subject adopts any modeled strategy in response to a given change in task conditions. This approach has the added advantage of automatically taking into account parameter uncertainty and model complexity (Myung & Pitt, 1997). Moreover, consideration of a hierarchical model in which the prior probabilities over models, $p_i = p(M = i)$, are themselves sampled from a Dirichlet distribution, provides us with a natural way to combine data across multiple subjects. This is because inferring a distribution over p_i given data from all of the subjects provides an estimate of the probability that a new (or randomly chosen) subject's behavior will be generated by each of the given models.

Here, we adopted a hypothesis-driven variant of factorial model comparison (see van den Berg, Awh, & Ma, 2014), by generating 20 different models, each of which assumes that a change in task conditions will have a particular effect. The 20 models differ only in the specification of prior probability over parameters used in each task condition. For example, one model may assume that the probability of remembering an item in WM is unaffected by task condition, while another model may assume that it

increases in one condition and decreases in another. In this manner, the current study employed model comparison techniques to investigate how strategic control over the interaction between WM and attention impacts the contents of WM—specifically, we examined how enhancing or suppressing memory-based guidance is related to differences in quality and/or accessibility of WM. In two experiments, we used slight variations of a dual-task paradigm that combined WM and visual search. The main measure of interest, WM performance, was measured with a continuous color recall task, in which participants recalled the memory color value from a spectrum of responses. Memory errors were fitted with a probabilistic model to estimate both the precision of the memory representation, as well as the probability that the item was retained in memory at all (see Zhang & Luck, 2008). Using factorial model comparison, we showed that strategic modulation of WM-attention interactions impacts WM accessibility rather than WM quality.

3.2 General Methods

Both experiments in this study used a similar paradigm, detailed here. Any differences from this paradigm are noted for each experiment.

3.2.1 Apparatus

Both experiments were conducted on a Dell Optiplex 960 computer, running Windows XP, and were programmed in Matlab using Psychophysics Toolbox, Version 3.0 (Brainard, 1997). Participants viewed the experimental displays on a LCD monitor

with a refresh rate of 60 Hz and screen resolution of $1,280 \times 1,024$ pixels at an approximate distance of 60 cm.

3.2.2 Procedure

Figure 7 depicts an example trial sequence. Each trial began with the presentation of a central fixation cross (black, stroke width = 0.1° , subtending $.5^\circ \times .5^\circ$) for 1,000 ms. The cross was followed by a blank screen for 500 ms, then by the memory cue item, a single colored circle (stroke width = 0.3° , subtending 2.8° in diameter), which was presented at the center of the screen for 250 ms. Participants were instructed to remember the color of the memory item for a subsequent test. The memory cue was followed by a blank delay for 2,000 ms, and then by the search array for 100 ms. The search array consisted of three colored circles (stroke width = 0.3° , each 2.8° diameter) at the corners of an imaginary triangle, with each corner approximately 5° from central fixation. Each circle contained a line of 1.2° length. Two of the lines (i.e., search distractors) were vertical, while one line (i.e., the search target) was tilted 15° to the left or right. Each target location and orientation occurred equally often and in a randomized order. Participants were instructed to indicate the orientation of the line via keypress to respond “left” or “right” tilted. The search array was followed by a blank screen until a keypress was recorded. After the search response or search timeout, a blank delay of 2,000 ms was presented, followed by a memory probe item that appeared in the center of the screen until response.

Two types of memory tests were used (Figure 7). In a memory recall test, a continuous-report probe appeared as a multicolored disc (2.8° diameter) with a question mark above it. Participants were instructed to adjust the color of the probe to match the initial memory cue color. The probe's color was adjusted using an input dial (PowerMate USB Multimedia controller, Griffin Technology, USA). Turning the dial caused the probe's color to cycle through the space of possible colors. Accuracy was stressed, and there was no time limit. In a memory recognition test, a change-detection probe appeared as a single colored circle (stroke width = 0.3° , 2.8° diameter) with a question mark above it. Participants were instructed to indicate whether the probe was identical to the initial memory cue color by using designated keypresses to respond "same" or "different." Match and non-match probes occurred equally often, in a randomized order; non-match probes were approximately 20° different from the initial memory cue. Both accuracy and speed were stressed for these trials.

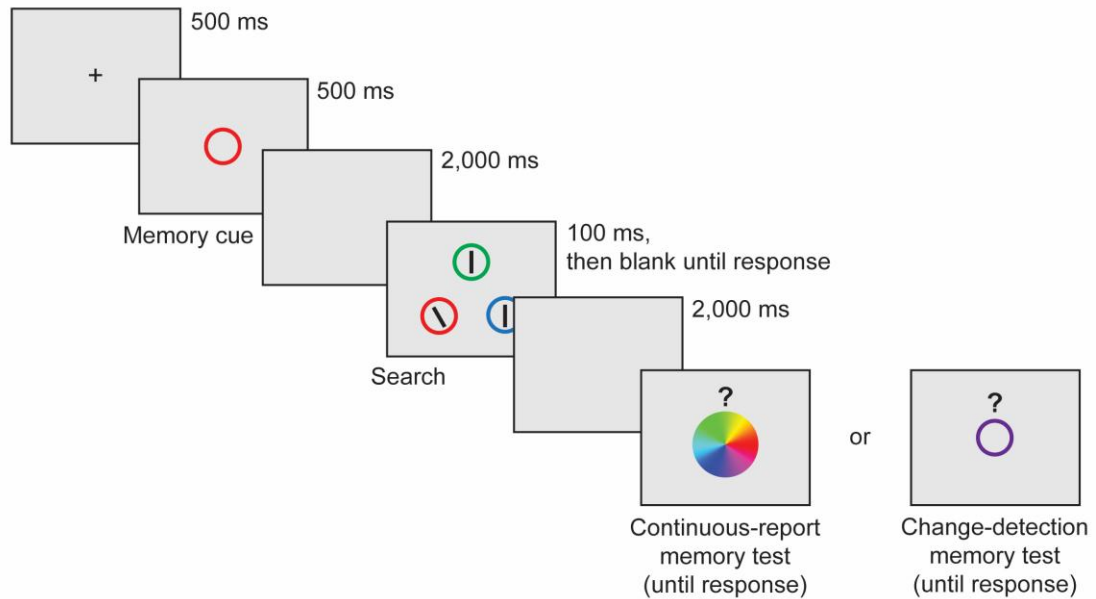


Figure 7: Example trial sequence. Participants were to remember a cue color while performing an intervening visual search task, in which they reported the direction of the tilted line. Then participants were given one memory task: either a continuous-report task (*left*), in which they recalled the memory cue from a continuous spectrum of responses, or a change-detection task (*right*), in which they responded whether the memory probe was same or different from the memory cue.

3.2.3 Design

Trials were classified by validity (i.e., neutral, valid, or invalid), depending on the relationship between the memory cue and the search array. In valid trials, the memory cue matched a circle within the search display, such that the target tilted line was contained within this match. In invalid trials, the memory cue matched a circle within the search display, but the target tilted line was never contained within this match. Memory item colors were pseudorandomly drawn to sample the entire range of circular HSV color

space with constant value and saturation. Search circle colors were set equidistantly from each other and from the memory cue color.

The critical manipulation was that trial conditions were blocked. Blocks could be composed of 100% valid trials (“Valid”), 100% invalid trials (“Invalid”), or 50% valid and 50% invalid trials (“Mixed”). Participants completed a practice session to learn about the different trial conditions and blocks; during the experimental phase, each block began with explicit instructions informing participants of block type and condition percentages.

3.2.4 Data Analysis

Our primary analysis focused on the continuous-report memory task. For each trial, a measure of recall error was calculated as the angular deviation between the memory response and the initial memory cue (range -180° to 180°). Overall memory accuracy was calculated as the root mean squared error (RMSE; see Fan & Turk-Browne, 2013).

To quantify the different components (i.e., quality and accessibility) of memory performance, we employed a mixture modeling technique. For each subject and for each block type, the distribution of recall errors was fit with a standard mixture model (Zhang & Luck, 2008) that attributes memory errors to either noisy recall of the memory cue (a Gaussian, or Von Mises, distribution centered on the cue color value) or to random guessing (a uniform distribution across all colors). This mixture model produces two parameters reflecting precision and accessibility, respectively: *sd* is the standard

deviation of the Gaussian distribution, or the inverse of precision, while g is the probability of random guessing (i.e., the probability of forgetting an item).

As in previous studies, we first compared average precision and average probability of random guessing across block types. However, we extended this analysis by using a hypothesis-driven variant of factorial model comparison (see van den Berg et al., 2014), which allows us to obtain an approximate probability that each subject uses a particular strategy or model. We employed a total of 20 standard mixture models that varied in their application of prior probability distributions (“priors”) on sd and g across block types (Table 3): Model 1, our baseline model, represented a standard mixture model with constant precision and constant guessing across block types; Models 2–10 represented models with constant precision, but variable or ordered guessing; Models 11–17 represented models with variable or ordered precision, but constant guessing; and finally, Models 18–20 represented models with variable precision and variable guessing. For these final models, only a few variants of variable precision and ordered guessing were included, based on the best fitting models within Models 2–17.

As previously mentioned, proper model comparison requires determining exact marginal likelihoods, which is computationally difficult for models with greater numbers of parameters. Instead we obtained approximate marginal likelihoods by computing a related quantity, the deviance information criterion (DIC; Spiegelhalter, Best, Carlin, & van der Linde, 2002), a relative measure of model fit and complexity that can be approximated from samples of the posterior distribution over parameters. These DIC

values were calculated for each model and each subject using the MemToolbox (Suchow et al., 2013). We also employed a Bayesian model selection approach (see Stephan, Penny, Daunizeau, Moran, & Friston, 2009) that treats models as random effects that can vary across subjects. We submitted approximate log model evidences to the `spm_BMS` routine of the SPM8 software suite (<http://www.fil.ion.ucl.ac.uk/spm/software/spm8>). With this method, we approximated the expected posterior probability that a randomly chosen subject will use each model, as well as the exceedance probability P^* of a model being more likely than any other model. While the exceedance probability of a model should not be interpreted as a “Bayesian p -value,” it does express confidence about that model’s favorability (Rigoux, Stephan, Friston, & Daunizeau, 2014).

Table 3: Descriptions of Model Families and Model Prior Probability Distributions for Each Parameter across Block Types

Model no.	Inverse precision (<i>sd</i>)	Probability of guessing (<i>g</i>)
Family 1: Constant precision, constant probability of guessing		
Model 1	Valid = Mixed = Invalid	Valid = Mixed = Invalid
Family 2: Constant precision, variable/ordered probability of guessing		
Model 2	Valid = Mixed = Invalid	{Valid}, {Mixed}, {Invalid}*
Model 3	Valid = Mixed = Invalid	Mixed > Invalid & Mixed > Valid
Model 4	Valid = Mixed = Invalid	Invalid > Mixed & Invalid > Valid
Model 5	Valid = Mixed = Invalid	Valid < Mixed & Valid < Invalid
Model 6	Valid = Mixed = Invalid	Mixed > Invalid > Valid
Model 7	Valid = Mixed = Invalid	Invalid > Mixed > Valid
Model 8	Valid = Mixed = Invalid	Mixed > Invalid & Invalid = Valid
Model 9	Valid = Mixed = Invalid	Invalid > Mixed & Mixed = Valid
Model 10	Valid = Mixed = Invalid	Valid < Invalid & Invalid = Mixed
Family 3: Constant probability of guessing, variable/ordered precision**		
Model 11	{Valid}, {Mixed}, {Invalid}	Valid = Mixed = Invalid
Model 12	Mixed > Invalid & Mixed > Valid	Valid = Mixed = Invalid
Model 13	Invalid > Mixed & Invalid > Valid	Valid = Mixed = Invalid
Model 14	Valid < Mixed & Valid < Invalid	Valid = Mixed = Invalid
Model 15	Mixed > Invalid & Invalid = Valid	Valid = Mixed = Invalid
Model 16	Invalid > Mixed & Mixed = Valid	Valid = Mixed = Invalid
Model 17	Valid < Invalid & Invalid = Mixed	Valid = Mixed = Invalid
Family 4: Variable precision, variable/ordered probability of guessing		
Model 18	{Valid}, {Mixed}, {Invalid}	{Valid}, {Mixed}, {Invalid}
Model 19	{Valid}, {Mixed}, {Invalid}	Mixed > Invalid > Valid
Model 20	{Valid}, {Mixed}, {Invalid}	Invalid > Mixed > Valid

* Curly brackets indicate variable precision or variable probability of guessing, such that the parameter value for each block type is independent.

** Note that *sd* refers to the standard deviation of the Gaussian distribution centered on the target, meaning that it is actually the inverse of precision. Thus, greater *sd* implies less precision, and vice versa.

3.3 Experiment 1

The present experiment was designed to replicate the finding that strategic control can modulate memory-based attentional guidance, as reflected in both search

performance and memory performance (Kiyonaga et al., 2012). Here, the memory task could probe recognition (change-detection) or recall (continuous-report). Including continuous-report responses allowed us to assess how strategic control over memory-based guidance affects the contents of WM.

Twenty volunteers (three men, ages 18–21 years) participated in exchange for course credit or a payment of US\$10 per hour and signed informed consent. Data from two additional volunteers were removed for low accuracy (< 80%) on the search task. Participants completed two blocks of each condition (Valid, Invalid, Mixed), presented in random order. In each type of block, 40% of trials had a change-detection memory probe, while 60% of trials had a continuous-report memory probe. Each block had 42 trials, for a total of 50 continuous-report responses per block type and 252 trials total.

3.3.1 Behavioral Results

Visual search task. Search accuracy was near ceiling ($M = 96.13\%$, $SD = 2.47\%$). We analyzed response times for correct visual search responses according to block and trial conditions (Table 4). A two-way repeated-measures ANOVA across factors of validity (i.e., valid or invalid) and predictability (i.e., 100%-predictable or Mixed) revealed that search was overall faster on valid (vs. invalid trials), $F(1, 19) = 168.16$, $p < .001$, and nearly significantly faster on 100%-predictable (vs. Mixed) blocks, $F(1, 19) = 4.25$, $p = .053$. In other words, valid memory cues led to search benefits, while invalid memory cues led to search costs; furthermore, benefits tended to be enhanced, and costs attenuated, by prior knowledge of whether the memory cues were reliably

helpful or hurtful, respectively. The relative impact of predictability was not modulated by validity, $F(1, 19) = 0.05, p = .827$.

Change-detection memory task. Mean accuracy for the change-detection memory task was 73.82% ($SD = 5.37\%$), confirming that the task was challenging, yet participants were generally able to maintain cues in WM. However, neither accuracy nor response times for the change-detection memory task differed significantly across validity or predictability (Table 4).

Continuous-report memory task. Mean RMSE (i.e., error) on the continuous-report memory task was 24° ($SD = 5.5^\circ$; Table 4). A two-way repeated-measures ANOVA revealed that memory error was significantly greater for invalid (vs. valid) trials, $F(1, 19) = 10.98, p = .004$, but there was no significant effect of predictability and no significant interaction, $ps > .2$. Similarly, a one-way repeated-measures ANOVA across block type (Valid, Invalid, Mixed) showed that memory error was greatest for Invalid blocks and lowest for Valid blocks, $F(2, 38) = 7.15, p = .002$.

3.3.2 Model Fitting and Comparison

Recall errors on the continuous-report memory task were fit with a standard mixture model (see Zhang & Luck, 2008) to obtain average estimates of precision and guessing (Table 5). Because memory load included only a single item, parameter estimates for sd ($M = 17^\circ, SD = 2^\circ$) and g ($M = 4\%, SD = 3\%$) were relatively high; however, these ranges are comparable to those in previous studies with low memory load (e.g., Hollingworth & Hwang, 2013; van Moorselaar et al., 2014b; see also Bays et al.,

2009; Zhang & Luck, 2008). Comparing these parameter estimates across block types showed no significant differences in sd , $F(2, 38) = 0.41, p = .668$, but significant differences in g , $F(2, 38) = 4.54, p = .029$ (Greenhouse-Geisser corrected). In other words, a comparison of group means indicated that while the precision of memory recall did not change across block types, the probability of forgetting that memory item did.

For Bayesian model comparison, recall errors were fit with the 20 different mixture models (Table 3). Figure 8a depicts, for each subject, the DIC value of each model relative to the DIC value of the most likely model for that subject (higher values represent worse fits). We then computed the expected posterior probability that a randomly chosen subject would use each model, as well as the exceedance probability P^* of that model being more likely than any other model. Across subjects, the highest-ranking model specified a standard mixture model with constant sd but ordered g across block types, such that the probability of random guessing was greatest in Invalid blocks, followed by Mixed blocks, and least in Valid blocks (posterior = .26, $P^* = .848$). Critically, all of the models that specified variable or ordered sd were highly disfavored, with the summed posterior probability of all 10 of these models not exceeding .20 (summed posterior = .18, summed $P^* = .002$).

3.3.3 Discussion

Experiment 1 indicated that top-down control, as driven by strategic knowledge about whether memory cues would be reliably helpful or hurtful for the subsequent search task, can modulate WM-attention interactions, replicating the results of Kiyonaga

et al. (2012, Experiment 1). These effects did not extend to memory performance in the change-detection memory task, although change-detection accuracy was numerically higher for valid compared to invalid trials. However, in the more difficult continuous-report task, overall memory accuracy was impacted by validity, with higher accuracy for valid versus invalid trials. These results support the idea that strategic enhancement or suppression of the memory cue during the delay (i.e., across the search task) affected memory performance. Model fitting and comparisons revealed that participants' behavior was best described by a strategy in which the probability of forgetting, as opposed to the precision of WM items, varied across block types. In other words, the relative enhancement or suppression of the WM item was not associated with differences in precision of that item, but rather with differences in the item's accessibility in WM, as reflected in the greater probability of randomly guessing during Invalid blocks.

Table 4: Mean Accuracy and Response Times for Visual Search and Memory Tasks for Each Trial Type

	Visual search task		Change-detection task		Continuous-report task
	Accuracy (%)	RT (ms)	Accuracy (%)	RT (ms)	Accuracy (°, RMSE)
Experiment 1 ($n = 20$)					
100% valid	97.0 (2.3)	590 (137)	75.9 (7.0)	996 (182)	19.5 (4.4)
50% valid	96.3 (5.0)	624 (133)	73.8 (10.0)	987 (234)	21.5 (9.0)
100% invalid	96.0 (3.2)	717 (131)	71.6 (6.9)	962 (204)	26.7 (9.1)
50% invalid	94.5 (4.1)	755 (177)	74.2 (12.6)	997 (218)	25.5 (8.9)
Experiment 2 ($n = 20$)					
100% valid	95.6 (4.2)	590 (118)	-	-	18.2 (3.3)
50% valid	96.4 (2.6)	637 (117)	-	-	18.8 (6.4)
100% invalid	94.0 (4.2)	728 (151)	-	-	24.6 (8.0)
50% invalid	93.3 (5.1)	756 (155)	-	-	23.0 (11.5)

Note. Standard deviations are presented in parentheses.

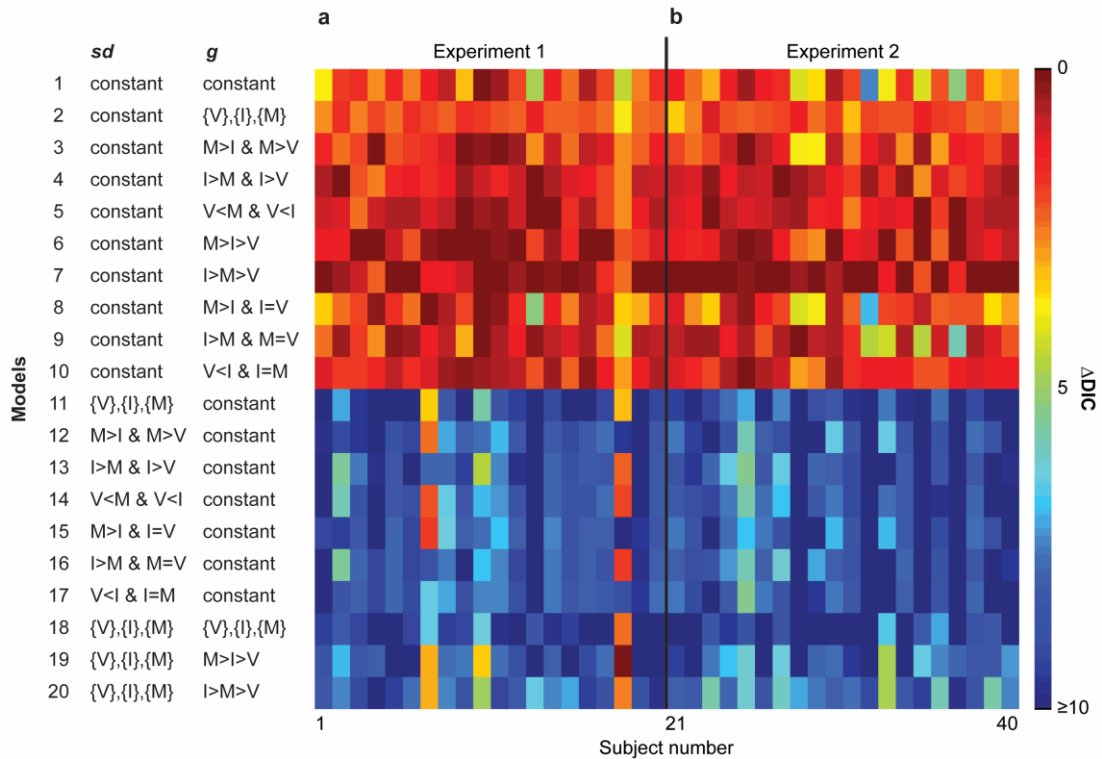


Figure 8: Bayesian model comparison between individual models for each participant. Each column represents a subject, divided by Experiment 1 (A) and Experiment 2 (B), and each row represents a mixture model as identified by model number (see Table 3) and prior probability distributions over two parameters: the standard deviation of the Gaussian component (i.e., precision; *sd*) and the probability of drawing from the uniform component (i.e., guessing; *g*). Prior probabilities were applied across three block types: Valid (V), Invalid (I), and Mixed (M). Curly brackets indicate variable precision or guessing across the three block types. Cell color indicates a model’s deviance information criterion (DIC) relative to that of the subject’s best-fitting model (a higher value means a worse fit). In general, the best-fitting models (redder cells) specify constant precision, whereas the worst-fitting models (bluer cells) specify variable or ordered precision.

3.4 Experiment 2

To strengthen the model comparison results of Experiment 1, Experiment 2 was designed as a replication with enhanced power, by increasing the number of continuous-

report response data points per block type, from 50 to 120 (see Bays et al., 2009). Thus, in Experiment 2, all trials contained recall tests, using continuous-report memory probes.

Twenty volunteers (five men, ages 19–29 years) participated in exchange for course credit or a payment of US\$10 per hour and signed informed consent. The data from three additional participants were removed, one for low search accuracy (< 80%), one for search RTs that exceeded 2.5 standard deviations above the group mean, and one for continuous-report memory accuracy (as measured by RMSE) that exceeded 2.5 standard deviations above the group mean. Participants completed two blocks of each condition (Valid, Invalid, Mixed), in random order. Each block had 60 trials, for a total of 120 continuous-report responses per block type and 360 trials total.

3.4.1 Behavioral Results

Visual search task. Search accuracy was near ceiling ($M = 94.93\%$, $SD = 3.53\%$). A two-way repeated-measures ANOVA across factors of validity and predictability revealed that search was significantly faster on valid (vs. invalid trials), $F(1, 19) = 86.02$, $p < .001$, and significantly faster on 100%-predictable (vs. Mixed) blocks, $F(1, 19) = 6.46$, $p = .020$, with no interaction, $F(1, 19) = 1.40$, $p = .252$ (Table 4). Thus, as in Experiment 1, search benefits and costs were strategically modulated based on prior knowledge of whether the memory cues were reliably helpful or hurtful.

Continuous-report memory task. Mean RMSE (i.e., error) on the continuous-report memory task was 22° ($SD = 6.0^\circ$; Table 4). A two-way repeated-measures ANOVA revealed that memory error was significantly greater for invalid (vs. valid)

trials, $F(1, 19) = 14.27, p = .001$, but there was no significant effect of predictability and no significant interaction, $ps > .5$. Similarly, a one-way repeated-measures ANOVA across block type (Valid, Invalid, Mixed) showed that memory error was greatest for Invalid blocks and lowest for Valid blocks, $F(2, 38) = 7.89, p = .001$.

3.4.2 Model Fitting and Comparison

Mean parameter estimates of sd ($M = 16^\circ, SD = 2^\circ$) and g ($M = 3\%, SD = 3\%$) were comparable to those in Experiment 1 (Table 5). Comparing parameter estimates of precision and probability of random guessing across block types showed no significant differences in sd , $F(2, 38) = 1.57, p = .221$, but significant differences in g , $F(2, 38) = 7.30, p = .002$ (Table 5). As in Experiment 1, the precision of memory recall did not change across block types, but the probability of forgetting did.

Figure 8b depicts for each subject, the DIC value of each model relative to the DIC value of the most likely model for that subject (higher values represent worse fits). Across subjects, the highest-ranking model was the same as in Experiment 1, specifying a standard mixture model with constant sd but ordered g across block types, such that the probability of random guessing was greatest in Invalid blocks, followed by Mixed blocks, and least in Valid blocks (posterior = .40, $P^* = .999$). As in the previous experiment, all of the models that specified variable or ordered sd were highly disfavored (summed posterior = .18, summed $P^* < .001$).

3.4.3 Discussion

Experiment 2 again demonstrated that top-down control can strategically modulate memory-based attentional guidance. Furthermore, this strategic modulation seems to be associated with changes in the accessibility, not the quality, of the WM item. Model fitting was performed on over twice as many continuous-report response data points per block, increasing the reliability of parameter estimates. Results from Bayesian model selection highlighted a specific winning strategy in which the probability of guessing changed across block types, whereas precision remained constant, with even stronger evidence for this particular model in Experiment 2 compared to Experiment 1 ($P^* = .999$ vs. $P^* = .848$, respectively).

Table 5: Mean *sd* and *g* Values for Continuous-Report Memory Task for Each Block Type

	Continuous-report task	
	<i>sd</i> (°)	<i>g</i> (%)
Experiment 1 ($n = 20$)		
Valid	17.0 (2.8)	1.3 (2.1)
Mixed	17.7 (3.4)	3.8 (4.1)
Invalid	17.0 (3.0)	5.6 (7.0)
Experiment 2 ($n = 20$)		
Valid	16.3 (2.2)	1.3 (1.6)
Mixed	15.7 (2.4)	3.7 (5.0)
Invalid	15.4 (2.3)	5.2 (4.7)

Note. Standard deviations are presented in parentheses.

3.5 Comparison of Model Families

For further Bayesian model comparisons, we collapsed the continuous-report memory task data from both Experiments 1 and 2. As reported above, the highest-ranking

and best-fitting model across all 40 participants specified a standard mixture model with constant sd but ordered g across block types (posterior = .50, $P^* > .999$), such that the probability of random guessing was greatest in Invalid blocks, followed by Mixed blocks, and lowest in Valid blocks (Figure 9a).

To directly ask whether strategic modulation of memory-based guidance affects either the quality or the accessibility of WM, we performed Bayesian model comparison of model families. Here “model family” refers to the subset of all models that share a particular level of a particular parameter. From the previous sample of 20 models, four families emerged: constant sd and constant g across block types (Model 1); constant sd , but variable g (Models 2–10); variable sd , but constant g (Models 11–17); and variable sd and variable g (Models 18–20). For each model family, we computed the expected posterior probability that a randomly chosen subject will use that model family, as well as the exceedance probability P^* of that model family being more likely than any other model family. Figure 9b depicts the comparison of model families across all 40 participants. The entire family of models with variable or ordered g across block types was ranked highest (posterior = .92, $P^* = 1$), whereas the entire family of models with variable or ordered sd was ranked lowest (posterior = .02, $P^* = 0$). These results provide compelling evidence that strategic control over the interaction between WM and attention does not impact the precision of the WM item.

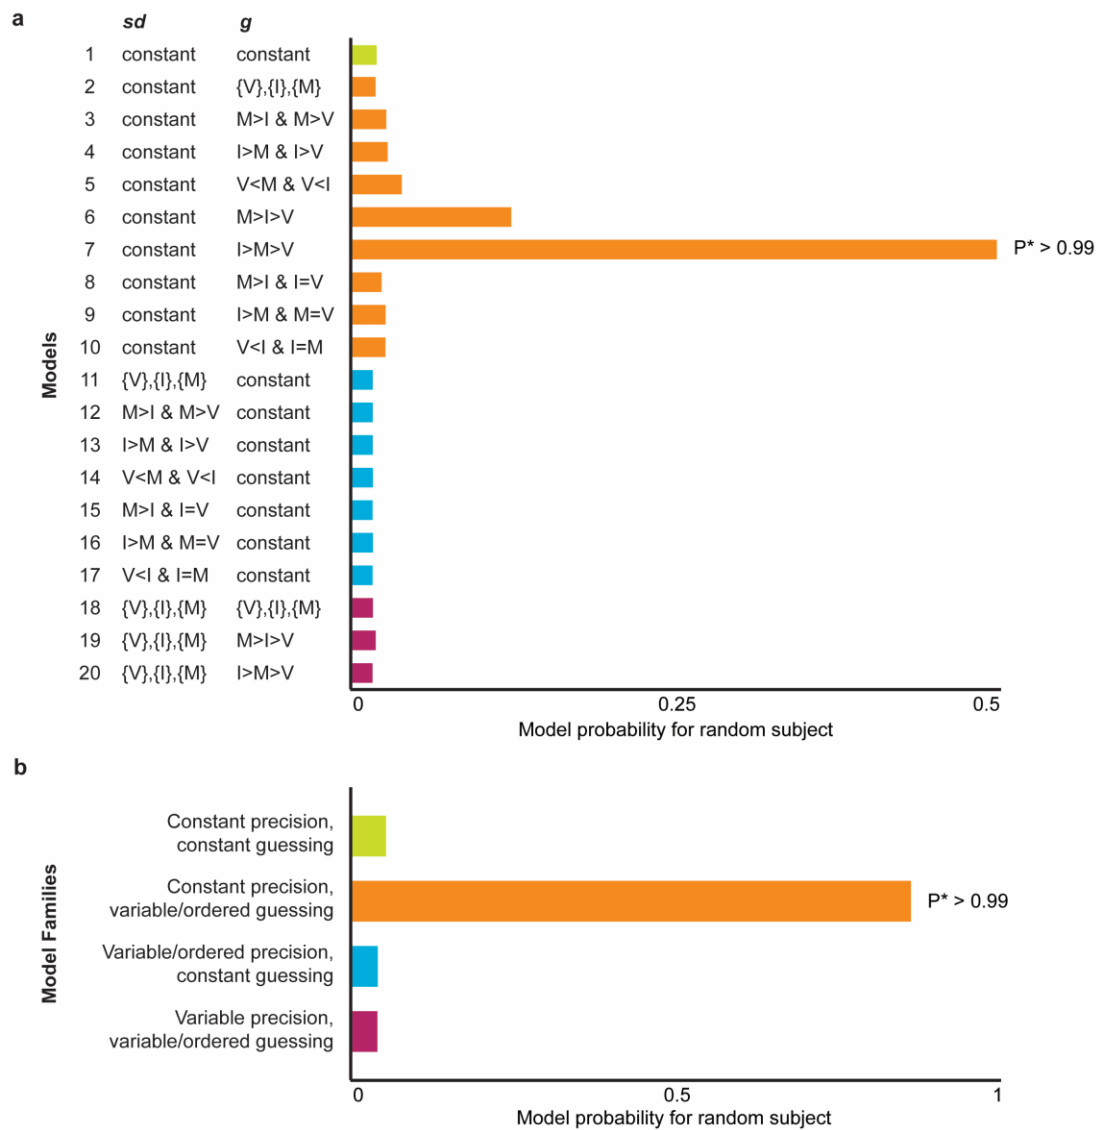


Figure 9: Comparison of individual models (A) and of four major families of models (B). Bars indicate the probability that a given model family generated the data of a randomly chosen subject. P^* values indicate exceedance probability (i.e., posterior probability that a given model is more likely than any other model). The best-fitting model (Model 7) specifies constant precision but ordered guessing (Invalid > Mixed > Valid), and the best-fitting model family specifies constant precision but variable/ordered guessing across block types.

3.6 General Discussion

Across two experiments, we assessed the manner in which a WM representation is modulated when the guidance of attention by WM is strategically enhanced or suppressed. In line with previous studies, valid memory cues benefited search performance, whereas invalid memory cues hindered search performance; prior knowledge about whether WM items would be reliably predictive (or anti-predictive) of the location of a search target allowed for strategic enhancement of valid memory cues and strategic suppression of invalid memory cues. While these data support an interplay between influences of top-down goals and WM on attentional selection, the continued presence of WM biasing effects in blocks of 100% reliably invalid searches also indicate that memory-based attentional guidance is not completely attributable to top-down strategy.

Our model comparison analysis showed that strategic control over the impact of WM on visual attention was reflected in visual WM as a change in the probability of random guessing, but not as a change in precision. Thus, strategic modulation over memory-based attentional guidance appears to be instantiated via changes in the accessibility of that WM representation. Specifically, an enhanced WM item could be more easily accessed by WM retrieval processes and thus more likely to impact attention, whereas a suppressed WM item could be stored in a manner that limits its impact on attention, as well as its ease of subsequent retrieval, thus increasing the likelihood of

forgetting. Importantly, strategic control over WM does not seem to modulate representational quality.

The current study provides strong evidence in favor of a growing literature in which the task-relevance of a memory item affects its accessibility rather than its representational quality (e.g., Hollingworth & Hwang, 2013; Murray et al., 2013; Souza et al., 2014; but see Pertzov et al., 2013; Williams et al., 2013). While these previous studies manipulated memory-relevance via retrospective cues, the current study indicated memory-relevance in relation to the subsequent attentional task. Here, memory-relevance was not based on whether the WM item would be tested, but was instead based on the likelihood that WM content would benefit visual search, thus recruiting strategic control over the interaction between WM and attention. Moreover, instead of simply evaluating changes in model parameter estimates, the current analysis employed a more powerful method, Bayesian model comparison, to determine the likeliest model of behavior for the observed data (Stephan et al., 2009). Thus, by overcoming key limitations of previous studies, we here resolved a central ambiguity in the previous literature, showing that the mechanism by which strategic control modulates the linkage between WM contents and attention seems to operate on the accessibility rather than the quality of memoranda.

Why would top-down strategic control over WM affect accessibility instead of quality? Previous studies have found that WM representations are maintained with intact precision for several seconds before terminating suddenly and completely (Donkin, Nosofsky, Gold, & Shiffrin, 2014; Zhang & Luck, 2009). This “sudden death” of items in

WM is thought to reflect a thresholded forgetting process in which an item gradually decays (not in terms of representational quality, but perhaps in terms of feature-location bindings or distinctiveness) to a certain threshold, then is suddenly dropped from memory (Zhang & Luck, 2009). Along these lines, strategic suppression of memory cues that were reliably unhelpful for subsequent search may have operated via a similar thresholded decay, such that the “suppressed” maintenance of that WM item led to a higher likelihood of being forgotten altogether. An alternative theory suggests that WM accessibility is impacted by the interference and overwriting from simultaneous (Bays et al., 2009) or subsequent items (Gorgoraptis, Catalao, Bays, & Husain, 2011; Makovski, 2012; Souza & Oberauer, 2014). Thus, strategic control over WM could enact an active suppression process that removes task-irrelevant or outdated information from WM (e.g., Ecker, Oberauer, & Lewandowsky, 2014). In the case of the current study, WM items in invalid trials were irrelevant for the search task, but not for the subsequent memory task; active suppression of the WM item during the search task, however, could spill over into the memory task, resulting in reduced accessibility at retrieval.

This idea of active suppression dovetails well with an influential theory of how WM interacts with attention, which proposes multiple, functionally-different states of WM: a more active type of representation that functions as an attentional template and can directly affect perception, and a less active “accessory” representation that does not influence the deployment of attention (Olivers et al., 2011). Another prominent theory of multiple-state WM posits a more distant state of activated long-term memory (Oberauer,

2002), in which WM representations are functionally partitioned away from the direct focus of (internal) attention (i.e., the most active state of WM), and as such might be more prone to retrieval failure. This form of maintenance in less active long-term memory, however, would not necessarily reduce the precision of the WM representation; recent studies have shown that visual long-term memory can store and retrieve thousands of objects and scenes with very high fidelity (e.g., Brady, Konkle, Gill, Oliva, & Alvarez, 2013). Within either framework, top-down control over WM-attention interactions could be instantiated by strategically assigning different memory representations to different functional states in WM, based on their assumed utility for the forthcoming task. The distinction between different functional states in WM may thus be characterized by memory accessibility. In the context of the present study, reliably invalid memory cues could be actively relegated to an accessory or long-term-like state, which may edge closer to a forgetting threshold, thus increasing the probability of forgetting the item and increasing the probability of guessing during a continuous-report memory task. While the current design does not directly address whether changes in memory quality and/or accessibility are *sufficient* to drive attentional guidance, our findings—combined with previous studies that show no impact of WM representational quality on driving attentional selection (Hollingworth & Hwang, 2013; van Moorselaar et al., 2014b)—suggest a functional division between the WM representations that do or do not impact attention based on WM accessibility.

Note that the current study does not differentiate between quantized or continuous models of visual WM capacity (i.e., how cognitive resources are distributed between representations in visual WM), particularly because only a single item was maintained during each trial. However, the present results can be reconciled with both discrete (e.g., Zhang & Luck, 2008) and flexible (e.g., Bays & Husain, 2008; Ma, Husain, & Bays, 2014) models of visual WM. Both camps allow for individual representations to be maintained with some variability and imperfection, such that individual items might gradually decay along some dimension that impacts accessibility. Thus, both discrete and flexible models of WM could be subject to a thresholded, lower bound of memory retrieval, which could be impacted by top-down control.

Furthermore, we recognize that the current analyses build off of a standard mixture model approach whose assumptions have been recently criticized (e.g., van den Berg et al., 2014; van den Berg, Shin, & Chou, 2012). We employed a standard mixture model of two components that reflect either remembering the target with some degree of noise or completely dropping the item from memory (Zhang & Luck, 2008). However, recent studies support a more sophisticated and better-fitting “variable-precision” approach, derived from flexible models of visual WM, which proposes that an item’s encoding precision is itself variable across trials (van den Berg et al., 2012; 2014); thus, our research question could be reframed as whether top-down control over memory-based attentional guidance modulates the *variability* of WM precision—as opposed to the more common dichotomy of an item being “in” or “out” of WM. Based on the current

findings, strategic suppression of the WM representation could thus reflect more variable precision, rather than completely dropping the item from memory. Nevertheless, this interpretation is still consistent with our present claim that different functional states in WM are instantiated through WM accessibility, even if accessibility is defined as more variable and lower precision.

While we propose that top-down control in the current paradigm is enacted through changes to WM accessibility, an alternative explanation is that participants were strategically attending to the memory-matching item during search, in order to refresh their memory representation of that color for the subsequent probe (see Woodman & Luck, 2007). We do not believe that resampling would have specifically impacted WM accessibility across blocks because the memory-matching color reappeared in every single search display, valid or invalid. Thus, the opportunity for memory refreshing was available in all blocks.

In summary, the current findings show that attentional selection can be driven by interactions between WM and strategic goals, again supporting the distinction between WM guidance and top-down control. Furthermore, using a powerful Bayesian model comparison technique, the present study demonstrates how strategic modulation of memory-based attentional guidance impacts the probability of remembering, but not the precision, of items in WM. Thus, different functional states in WM may be characterized by differing retrieval accessibility rather than representational quality, such that top-down

control over the interaction between WM and attention operates by changing how likely an item is to be remembered or forgotten.

4 Individual Differences in Working Memory Guidance

Chapter 3 proposed that whether the contents of WM influence visual attention depends on the functional state of the WM representation itself, as characterized by retrieval accessibility. The current chapter approaches the question of when WM does or does not guide visual attention from a different perspective, arguing that various cognitive factors contribute to and influence the interaction between WM and attention. As described in Chapter 1, several studies using very similar dual-task WM and attentional paradigms have produced divergent results and interpretations about whether memory-based attentional guidance is automatic or under strategic control. This chapter presents an individual-differences study, in which participants completed a battery of cognitive tasks, including two common attentional guidance paradigms, a visual WM task, and a complex WM span task, as well as attention-related self-report assessments. The results suggest that these two task paradigms—which have previously produced contrasting patterns of performance—probe distinct aspects of memory-based attentional guidance, as influenced by individual differences in WM capacity and attentional control. Thus, the impact of WM on attentional deployment is itself a complex and multifaceted construct. These data were published as a brief report in the journal *Attention, Perception, & Psychophysics* (Dowd, Kiyonaga, Egner, & Mitroff, 2015b).

4.1 Introduction

WM and attention are typically considered distinct cognitive constructs, yet they frequently overlap (see Chapter 1). For example, the active maintenance of items in visual WM can bias attention toward matching items in the environment, even at the expense of the current task goals (Soto et al., 2008). However, this bias is not always observed (e.g., Woodman & Luck, 2007), generating debate about whether the coupling between WM and attention is involuntary or under strategic control. While differences in memory-based attentional guidance could be attributed to a complex set of task parameters, it is also possible that the various task paradigms that engender conflicting interpretations are actually tapping into different cognitive processes.

Studies of memory-based attentional guidance typically involve a visual search task while simultaneously maintaining a visual object in WM. The critical manipulation is whether the search target, distractors, or neither matches the stimuli in WM. In one paradigm (e.g., Soto et al., 2005), participants view a colored shape WM item (e.g., a blue circle) before performing a visual search for a slanted-line target among straight-line distractors (Figure 10a). The WM item can reappear in the search display surrounding a target (valid), distractor (invalid), or fail to reappear at all (neutral). The visual search stimuli are thus distinct from, but embedded within, the class of stimuli that could match WM. The canonical finding is that search is speeded when the memory item validly cues the target and slowed when it invalidly cues a distractor (relative to the neutral condition; see Soto et al., 2008). Even when valid trials are eliminated from this task—de-

incentivizing any strategic orienting toward the WM-match—search is still slowed on invalid trials (e.g., Dowd et al., 2015a; Kiyonaga et al., 2012; Olivers, 2009; Soto et al., 2005), suggesting the link between WM and visual attention is automatic (Soto et al., 2008; see also Dowd & Mitroff, 2013; Hollingworth, Matsukura, & Luck, 2013).

In contrast, other seemingly similar memory-based guidance tasks have produced different results. In one prominent task, participants view a colored square to hold in WM before searching among an array of colored Landolt C-like stimuli for a C whose gap is on the top or bottom (Woodman & Luck, 2007; Figure 10b). The search target or one of the distractors could match the WM color. In some studies using this task, visual attention was not captured by the memory-matching items (e.g., Han & Kim, 2009; Woodman & Luck, 2007). In fact, when there was no possibility of WM content matching the search target (i.e., no valid trials), there was a tendency for participants to be *faster* when a distractor matched WM, suggesting that participants were strategically avoiding the memory-matching distractors (Arita et al., 2012; Woodman & Luck, 2007, Experiment 3). These experiments suggest that there is no automatic visual selection bias toward WM-matching items, and that WM content can be strategically avoided (see Carlisle & Woodman, 2011a).

Several studies have attempted to address why these superficially similar paradigms produce divergent results by focusing on specific task parameters that might underlie differences in memory-based guidance (e.g., Dalvit & Eimer, 2011; Dombrowe et al., 2010; Han & Kim, 2009; Olivers, 2009; Soto & Humphreys, 2008; Tsvetanov et

al., 2012; Woodman et al., 2007; Zhang et al., 2010; 2011). These investigations, however, have not coalesced into a unified picture of whether and how these factors impact memory-guided attention. Moreover, contradictory behavioral indices of memory-guided attention have been joined by contradictory electrophysiological findings (i.e., differential modulation of the N2pc event-related potential component) from these two types of paradigms (Carlisle & Woodman, 2011b; 2013; Kumar et al., 2009). Taken together, these findings challenge the assumption in this field that both types of paradigms probe the same underlying processes of how WM contents guide attention. Thus, one possible explanation for why two very similar paradigms engender different interpretations is that they are not measuring the same cognitive construct in the first place. This possibility has not been addressed, so the current study used an individual-differences approach to examine the relationship between performance on the two tasks. As a secondary focus, we also compared memory-based capture in each task relative to independently measured constructs of attention and WM.

4.2 Methods

4.2.1 Participants

Seventy-four individuals from Duke University (24 men, ages 18–27 years) participated for US\$10/hour or course credit. All signed informed consent in accordance with institutional guidelines. Data were excluded from six additional participants who did not complete all four experimental tasks.

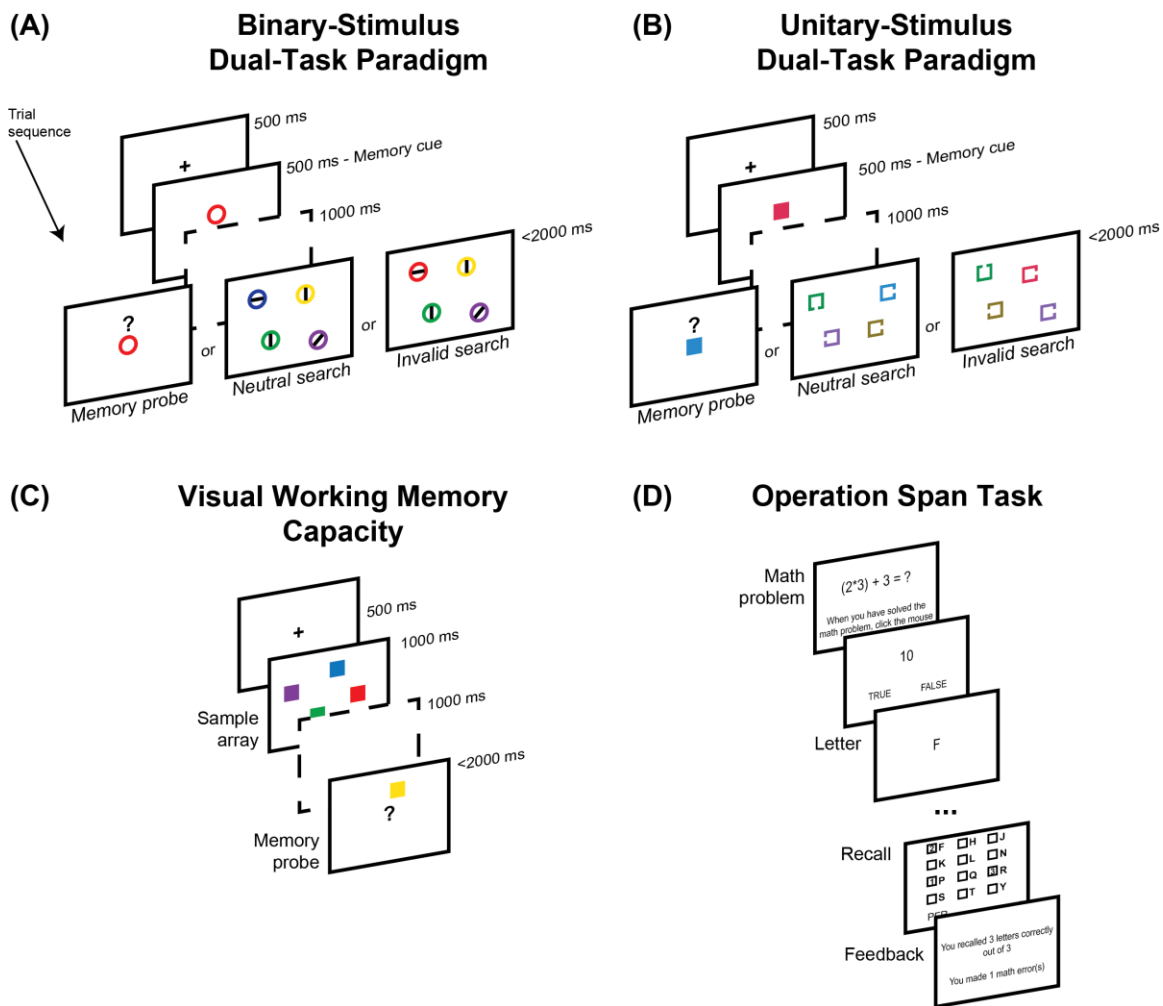


Figure 10: Trial sequences for the four behavioral tasks. (A) Binary-stimulus dual-task paradigm that combines a memory task with an intervening search task; participants searched for the tilted line target among vertical or horizontal line distractors. (B) Unitary-stimulus dual-task paradigm; participants searched for the top or bottom gap target among right or left gap distractors. In both dual-task paradigms, the memory task tested whether the memory probe was same or different from the memory cue. (C) Visual working memory capacity task; participants reported whether memory probe color was same or different from the item at that location in the sample array. (D) Automated operation span task; participants alternated between remembering letters and solving arithmetic problems.

4.2.2 Behavioral Testing

Participants completed four cognitive tasks (Figure 10): two dual-task paradigms that combined WM and visual search (see Soto et al., 2005; Woodman & Luck, 2007), a visual WM capacity task (see Luck & Vogel, 1997), and a complex WM span task (automated operation span task; Unsworth, Heitz, Schrock, & Engle, 2005). Visual WM capacity was included to reflect the strength of a WM representation (e.g., Bays & Husain, 2008), which in turn could affect how memory-matching contents in the visual field capture attention. Complex span was included as a measure of WM maintenance in the face of interference, which resembles the dual-task setup of the attentional guidance paradigms in question here. Tasks were counterbalanced across participants, but the two dual-task paradigms were never completed consecutively.

4.2.2.1 Dual-Task Paradigms

The two dual-task paradigms were based on designs by Soto et al. (2005, 2008) and by Woodman & Luck (2007; see also Carlisle et al., 2011; Carlisle & Woodman, 2011a; Han & Kim, 2009). We refer to these tasks as using “binary stimuli” (when the memory-related circles and search-related lines have separate forms; see Soto et al., 2005) and “unitary stimuli” (when the memory-related color and the search-related gap exist within the same form; see Woodman & Luck, 2007), respectively. We adapted the paradigms to have identical instructions, timing, display sizes, and stimulus colors; the contributions of these elements to memory effects on attention have been addressed by previous studies (e.g., Dombrowe et al., 2010; Olivers, 2009), and thus the current study

focuses on the differences in stimulus arrangement. Importantly, the memory cue *never* matched the location of the search target, removing incentive to strategically look for the memory-matching search item (see Soto et al., 2005, Experiment 4).

Binary-stimulus paradigm. In the binary-stimulus dual-task paradigm (Figure 10a), each trial began with a 500 ms fixation display, followed by a 500 ms memory cue that was a colored circle ($1.3^\circ \times 1.3^\circ$) centered at fixation. Participants were to remember the color of the cue for a potential future test. After a 1,000 ms blank delay, a search array was presented until response (up to 2,000 ms). The search display consisted of four black lines (stroke width = 0.2° , length = 1.0°), each surrounded by a colored circle. Three were horizontal or vertical distractor lines, whereas the target line was tilted left or right at 45° . Each search item was placed with a slight spatial jitter within randomly selected cells of an invisible 5×4 grid subtending $10.3^\circ \times 8^\circ$, with no items overlapping. Each target location and orientation occurred equally often and in randomized order. Participants searched for the tilted line and reported the direction of the tilt via speeded keypress.

To prevent participants from strategically using the search array to refresh their memory of the memory item during the search task (see Woodman & Luck, 2007), memory probes were given on 20% of trials *instead of* a search display. The memory probe was a colored circle ($1.3^\circ \times 1.3^\circ$) presented at the center of the screen beneath a question mark for up to 2,000 ms. Participants indicated whether the probe color was identical to the memory cue color via keypress. The memory cue, search circles, and

memory probe were drawn randomly from five equidistant and equiluminant colors. In the search array, each color could only appear once.

Half of the search trials were neutral (memory item did not reappear) and half were invalid (memory item reappeared surrounding a distractor line). Memory probe trials were either a match (50%) or non-match (50%) to the WM cue. After 10 practice trials, participants completed 200 experimental trials.

Unitary-stimulus paradigm. The unitary-stimulus dual-task paradigm was identical to the binary-stimulus task in timing and structure, but employed different stimuli (Figure 10b). Memory cues and probes were colored squares ($0.65^\circ \times 0.65^\circ$), and the search stimuli were outlined squares (stroke width = 0.1°) with a 0.25° gap on the top, bottom, left, or right. Three of the squares had a gap on the left or right, while the target square had a gap on the top or bottom. Participants searched for the top or bottom gap and reported the orientation of the gap via speeded keypress.

4.2.2.2 Visual Working Memory Capacity

The visual WM task (Figure 10c) was adapted from Luck & Vogel (1997). After a 500 ms fixation display, a sample array of 2, 4, 6, or 8 colored squares ($1.0^\circ \times 1.0^\circ$), jittered across a grid of $15.4^\circ \times 10.7^\circ$, was presented for 1,000 ms. After a 1,000 ms blank delay, a test array was presented, consisting of a single colored square in the same location as one of the sample array items. Participants reported whether the test item was identical to the sample item that had previously been shown in that location. Colors were

pulled pseudorandomly from a circular spectrum of equiluminant colors. If the test item was different from the sample item, the difference was 180° in color space.

4.2.2.3 Operation Span

The operation span task was administered with an automated version packaged by Unsworth et al. (2005). Participants were to remember information (sequences of unrelated letters, set sizes: 3–7) while concurrently engaging in online processing (solving arithmetic problems). On each trial, participants alternated between solving math problems and memorizing single letters until a recall phase wherein they reported the letters in order (Figure 10d). Each participant completed three sets of each set size (i.e., 15 sets). Full details of task structure and timing can be found in Unsworth et al. (2005).

4.2.3 Individual-Differences Assessments

Behavioral task data were collected along with a large battery of self-report assessments (see Appelbaum, Cain, Schroeder, Darling, & Mitroff, 2012). The current study examined two attention-related surveys: the Attention Deficit-Hyperactivity Disorder (ADHD) questionnaire (Jasper & Goldberg, 1993) and the Autism-Spectrum Quotient (ASQ; Baron-Cohen, Wheelwright, Skinner, Martin, & Clubley, 2001) with a focus on the two attention-related sub-categories of attention-switching and attention to detail. Both ADHD and autism have been linked to broader deficits in executive functions, including attentional control and WM (see Gargaro, Rinehart, Bradshaw,

Tonge, & Sheppard, 2011; Pennington & Ozonoff, 1996), which could have important implications for memory-guided attention. No other questionnaire data were analyzed.¹

4.3 Results

In the dual-task paradigms, our primary measure of interest was response time (RT) in the visual search task. Specifically, our key comparison was the time to find the target in neutral versus invalid conditions. To represent a single measure of attentional capture by WM, a normalized RT difference score was calculated for each individual: $(\text{invalid RT} - \text{neutral RT}) / (\text{neutral RT})$, reflecting the absolute RT difference as a proportion of neutral RT. Positive values indicate high attentional capture, while negative values indicate reverse attentional capture.

4.3.1 Binary-Stimulus Paradigm

Performance on memory catch trials was nearly perfect ($M = 99.1\%$, $SD = 1.7\%$). Search trials were excluded for response errors (0.7% of all search trials) and for exceeding the time limit (0.4% of trials). RTs for correct search responses were slower for invalid trials ($M = 1,045$ ms, $SD = 138$ ms) than for neutral trials ($M = 996$ ms, $SD = 131$ ms), $t(73) = 10.57$, $p < .001$ (Figure 11a). There was high variability across participants, with normalized RT difference scores ($M = 0.05$) ranging from +0.16 to -0.03 (for similar reports of individual variability, see Soto, Rotshtein, & Kanai, 2014).

¹ In line with previous research from our lab (e.g., Donohue, Darling, & Mitroff, 2012), although we had a variety of self-report measures available, we only analyzed those that were identified *a priori* as relevant.

4.3.2 Unitary-Stimulus Paradigm

Performance on memory catch trials was nearly perfect ($M = 98.6\%$, $SD = 2.7\%$). Search trials were excluded for response errors (0.4%) and for exceeding the time limit (1.0%). Search RTs for correct responses were slower for invalid trials ($M = 1,172$ ms, $SD = 135$ ms) than neutral trials ($M = 1,128$ ms, $SD = 136$ ms), $t(73) = 7.08$, $p < .001$ (Figure 11a). There was high variability in normalized RT difference scores ($M = 0.04$) ranging from +0.15 to -0.09. This particular paradigm has revealed variable capture effects (Carlisle & Woodman, 2011b; Han & Kim, 2009), and the data here demonstrated significant positive capture, in contrast to the reverse attentional capture effect from Woodman & Luck (2007, Experiment 3).² Importantly though, the presence or direction of a mean capture effect in this protocol is not a key concern of the purpose of the present study, which is to assess the relationship of *individual differences* in performance across tasks.

² It is possible that the specific timing and display differences between the current unitary-task paradigm and Woodman & Luck's task may have engendered discrepant results. However, while both the binary-stimulus and unitary-stimulus tasks here produced positive capture, additional analyses did suggest differences in the magnitude of attentional capture. A Bayesian parameter estimation (Kruschke, 2013) using Markov chain Monte Carlo estimation was used to determine the likely distribution of attentional capture effect sizes in both paradigms. Attentional capture in the binary-stimulus paradigm had a mean effect size of $d = 1.29$ [95% HDI 0.96, 1.62], whereas attentional capture in the unitary-stimulus paradigm was weaker, $d = 0.89$ [95% HDI 0.59, 1.19]. In addition, search RTs for the unitary-stimulus paradigm were overall slower than RTs for the binary-stimulus paradigm, $F(1, 73) = 138.4$, $p < .001$, likely due to the more perceptually demanding search task in the latter paradigm (Han & Kim, 2009).

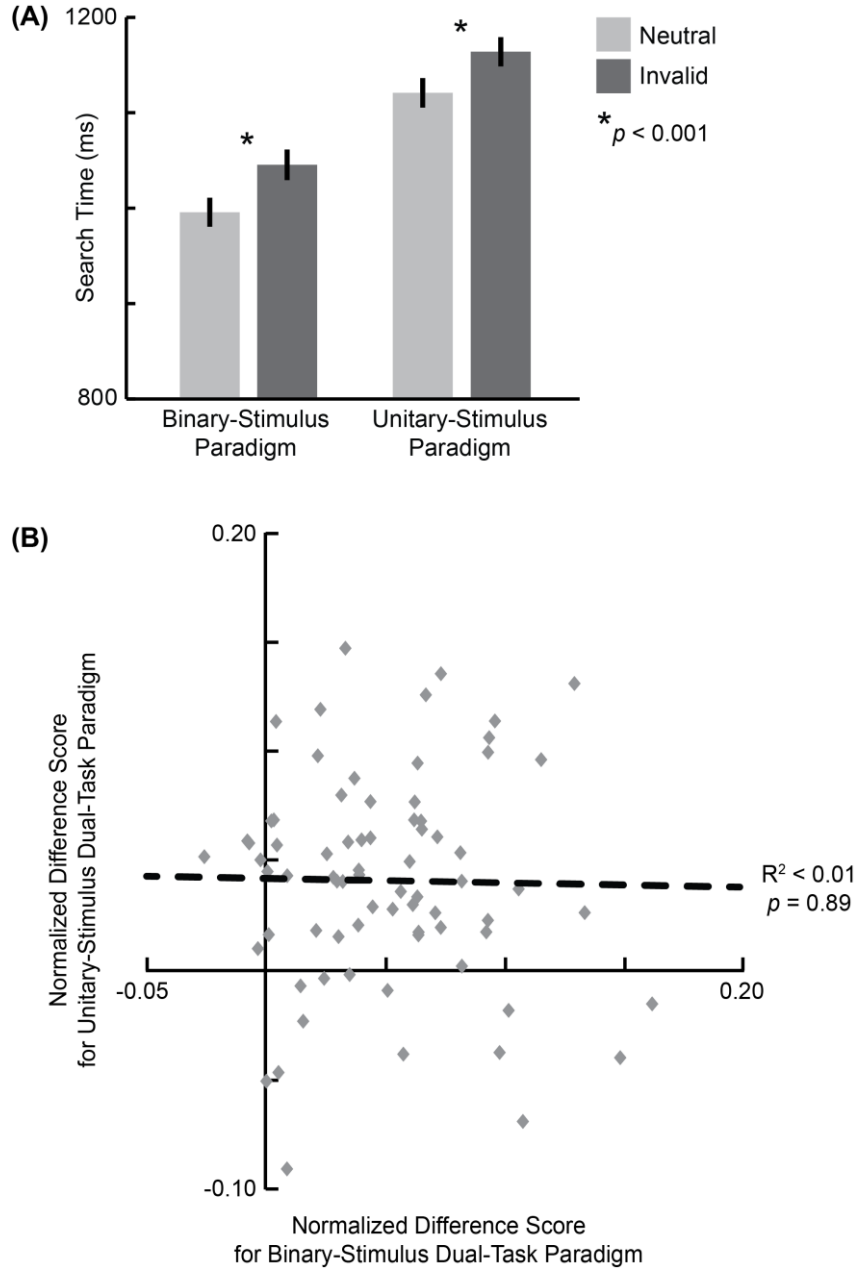


Figure 11: Memory-based capture results. (A) In both the binary-stimulus and unitary-stimulus paradigms, search times were slower for invalid trials compared to neutral trials. Error bars represent 95% confidence intervals. (B) Individual differences in normalized difference scores (i.e., the amount of attentional capture by memory-matching distractors) did not correlate between the binary-stimulus and unitary-stimulus paradigms.

4.3.3 Visual Working Memory Capacity

Visual WM capacity was calculated using Cowan's K (Cowan, 2001): $K = S * (H - F)$, where K is memory capacity, S is set size, H is observed hit rate, and F is false alarm rate. Final visual WM capacity was calculated by averaging K values for set sizes 4, 6, and 8 (Fukuda & Vogel, 2009). As in previous studies (e.g., Vogel & Machizawa, 2004), mean WM capacity for our population was 3.20 ($SD = 1.04$), ranging from 1.05 to 5.39.

4.3.4 Operation Span

Operation span score was calculated by summing all perfectly recalled letter sets. For example, if an individual correctly recalled 3 letters in a set size of 3, 4 letters in a set size of 4, and then the first 2 letters in a set size of 5, the operation span score would be 7 (3 + 4 + 0). A perfect score would be 75. Full details of operation span calculations can be found in (Unsworth et al., 2005). Across participants, mean operation span score was 49.2 ($SD = 13.4$), ranging from 20 to 75.

4.3.5 Regression Analyses

Our analysis strategy was to use regression models to assess whether performance on the two dual-task paradigms was related. In a simple linear regression model, normalized difference scores on the binary-stimulus task did not significantly predict normalized difference scores on the unitary-stimulus task, $\text{adj. } R^2 < .001$, $F(1, 72) = 0.02$, $p = .885$, 95% CI [-0.20, 0.17] (Figure 11b). A Bayesian linear regression (Kruschke, 2011) revealed a posterior distribution of standardized slopes centered on $M = -0.09$

[95% HDI $-0.33, 0.16$], supporting the conclusion that performance on binary-stimulus and unitary-stimulus tasks were not significantly related. To show that this lack of correlation was not due to unreliability of either task, we used attenuation correction to adjust the correlation according to the reliability of the two measures being correlated (Schmidt & Hunter, 1996). With difference score reliabilities of $\rho_{\text{binary}} = 0.45$ and $\rho_{\text{unitary}} = 0.41$, the corrected correlation between attentional capture effects in binary and unitary paradigms was still not significant, $r(74) = -.04, p = .740$.³ To confirm that this lack of correlation was not due to validity effect differences between the tasks, search response times were subjected to a 2×2 repeated-measures ANOVA across validity and task; the interaction was non-significant, $F(1, 73) = 0.51, p = .479$.

The primary analysis of interest was the above correlation between attentional capture effects in the binary-stimulus and unitary-stimulus paradigms. Having found that they are not related to one another, we conducted exploratory regression analyses to see whether performance was differentially predicted by WM capacity or self-report measures of attention. A stepwise linear regression model was applied to both dual-task paradigms to explore what factors best predicted attentional capture for each. Five factors were included in an initial model for each paradigm: visual WM capacity, complex span,

³ Although these reliability estimates seem low, it is important to note that difference scores are mathematically biased toward having relatively lower reliability than raw or composite scores (Webb, Shavelson, & Haertel, 2006). For both paradigms, Cronbach's alpha (a measure of internal consistency) was $> .95$ for both neutral and invalid trials.

self-report scores of ADHD, self-report scores on the attention-switching subcategory of ASQ, and self-report scores on the attention-to-detail subcategory of ASQ.

In predicting performance for the binary-stimulus task, four factors (operation span, ADHD score, and two ASQ sub-category scores) did not significantly contribute to the model (full model, $\text{adj. } R^2 = .01$). Only visual WM capacity significantly predicted attentional capture, $\text{adj. } R^2 = .040$, $F(1, 72) = 3.99$, $p = .049$ (Figure 12a). Outliers were assessed based upon Cook's $D \geq 1$ (Cook & Weisberg, 1982), and no data were trimmed. In predicting performance for the unitary-stimulus task, two factors (visual WM capacity, operation span) did not significantly contribute to the model (full model, $\text{adj. } R^2 = .10$). Self-report scores of ADHD symptoms and the two ASQ sub-categories significantly predicted attentional capture, $\text{adj. } R^2 = .119$, $F(3, 70) = 4.26$, $p = .008$ (Figure 12b). Again, no data points were trimmed for Cook's $D \geq 1$. No factor included in the model had a variance inflation factor above 1.12, indicating that collinearity was not an issue.

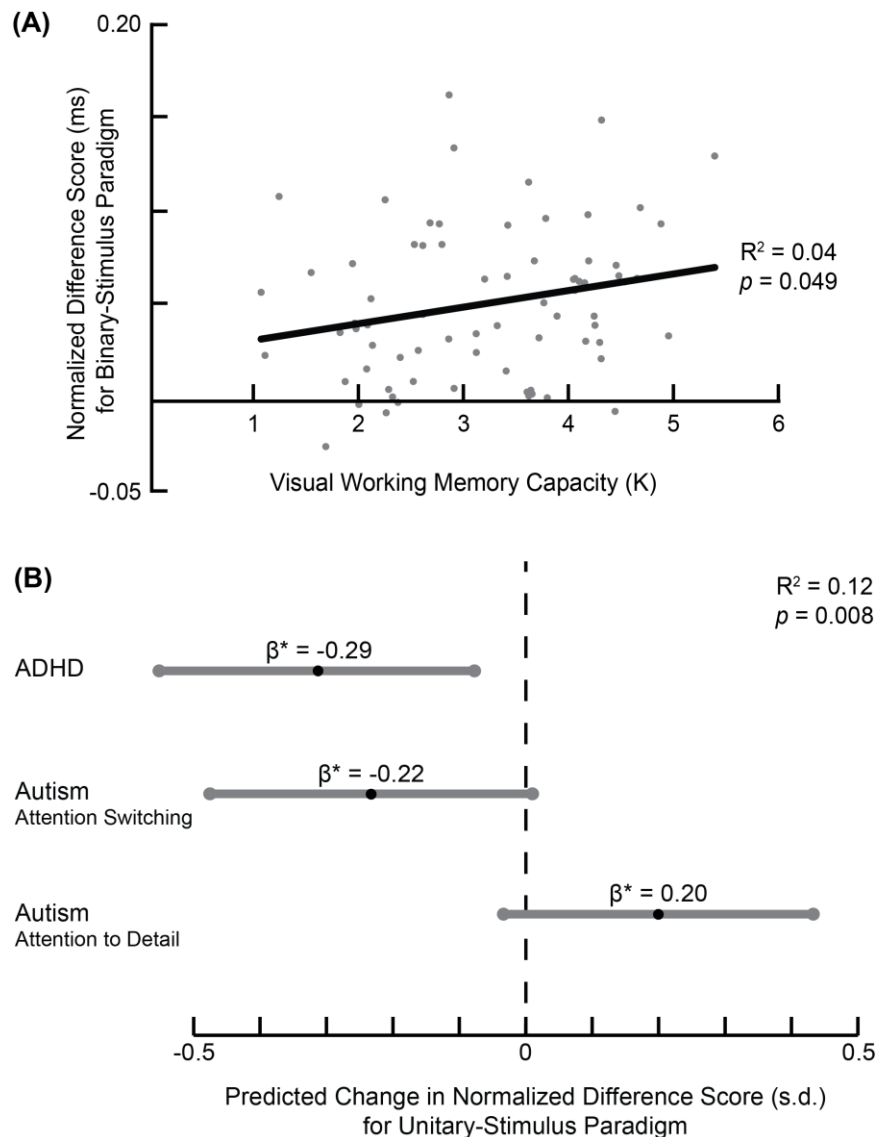


Figure 12: Fitted linear regression models predicting memory-based capture in each dual-task paradigm by selected, significant variables of working memory and attention. (A) Normalized difference scores (i.e., the amount of attentional capture by memory-matching distractors) in the binary-stimulus paradigm were significantly predicted by individual differences in visual working memory capacity. (B) Normalized difference scores in the unitary-stimulus paradigm were significantly predicted by self-report assessments of attention deficit-hyperactivity disorder and two sub-categories of the autism spectrum quotient. Lines depict 95% confidence intervals around mean standardized regression coefficients (black dots) in predicting normalized difference scores in a multiple regression model.

4.4 Discussion

A great number of studies have explored the extent to which WM content can capture visual attention, but debate remains open over whether attentional guidance by WM is automatic (e.g., Soto et al., 2008) or strategic (e.g., Woodman & Luck, 2007). This debate has been fueled by conflicting findings from two seemingly similar dual-task paradigms, which putatively test the same aspect of attentional guidance. The current study employed an individual-differences approach to examine whether performance in one paradigm related to performance in the other. Attentional capture by WM in the binary-stimulus (see Soto et al., 2005) and unitary-stimulus (see Woodman & Luck, 2007) paradigms did not correlate, suggesting a fundamental difference in how each task taps into WM biasing, even within the same individuals.

Much of the evidence that has led to debate over memory-guided attention has been drawn from one or the other of these dual-task paradigms. On the surface, these paradigms seem very similar, generating the assumption that the two are assessing the same underlying phenomenon. In fact, at the group level, the current study observed capture effects of roughly the same magnitude across the two paradigms. Thus, it is all the more striking that the magnitude of attentional capture did not correlate at an individual level. The current behavioral results, along with previous contradictory electrophysiological findings between the two paradigms (see Carlisle & Woodman, 2013), suggest the possibility that these two paradigms—which differ only in terms of stimulus arrangement—are actually tapping into distinct aspects of how memory and

attention interact. The current results thus offer a new perspective on reconciling the disparate literature about memory-guided attention.

What might these distinct aspects of memory-based guidance entail? Using the other individual-differences measures, subsequent exploratory analyses hinted that attentional capture in the two paradigms was differentially predicted by independently measured constructs of WM and attention—specifically, by visual WM capacity for the binary-stimulus task, and by self-report assessments of attention for the unitary-stimulus task. It is important to emphasize that the amount of variance explained by these factors was relatively small; furthermore, the current individual-differences approach did not allow for empirical manipulation of WM- and attention-related factors. Thus, these results only serve to generate initial hypotheses as to why the two paradigms produce different patterns of behavior.

The exploratory analyses revealed a dissociation between predictive factors for the binary-stimulus versus the unitary-stimulus paradigm. For the binary-stimulus paradigm, greater visual WM capacity predicted greater attentional capture. Previous research has suggested that visual search processing taps into WM resources (Woodman & Chun, 2006). We speculate that visual WM capacity might be more relevant for the binary-stimulus paradigm due to the physical distinction between binary and unitary stimuli. In binary-stimulus search arrays, there are two spatially-distinct forms, the colored circle and the search line, which could theoretically be processed as separate “objects” (Duncan, 1984); in contrast, unitary stimuli consist of one form only. Thus,

although only one “object” is encoded during the memory cue display, twice as many “objects” are processed in binary versus unitary search displays, which may tax WM resources to a greater degree.

Memory-based guidance in the unitary-stimulus paradigm was predicted by self-report assessments of ADHD and autism, which have been linked to difficulties with attentional control and flexibility (Geurts, Verte, Oosterlaan, Roeyers, & Sergeant, 2004; Manly, Anderson, & Nimmo-Smith, 2001). These findings point to greater attentional capture in individuals who might be less able to flexibly deprioritize WM items (i.e., relegate them to accessory status; Olivers et al., 2011) during visual search. We speculate that these factors might be specifically predictive for performance in the unitary-stimulus paradigm due to the spatial concentration of task-relevant features, which may require more attentional control to dissociate the search-relevant feature (gap) from the memory-relevant feature (color). This issue goes hand-in-hand with search difficulty, in the sense that distinguishing between memory-relevant and task-relevant features in unitary stimuli is more perceptually demanding. However, general search difficulty in and of itself has not been found to influence memory-based attentional capture (Olivers, 2009). Overall, these exploratory analyses present a preliminary window into how the spatial differences between binary and unitary stimuli might differentially recruit object-encoding and attentional control processes, which may underlie the differences in attentional capture between the binary-stimulus and unitary-stimulus paradigms.

In much of the memory-guided attention literature, there has been an assumption that two superficially very similar and highly influential tasks are testing the same relationship between WM and attention, thus driving a debate about why results differ. By exploring these differences within the same individuals, the present study offers an alternative explanation: these two tasks produce non-correlating effects because they might be tapping into distinct aspects of memory-based attentional guidance. These findings underscore the importance of stimulus attributes and task design in determining the recruitment of particular attention and memory processes, and reveal that a multitude of cognitive factors can contribute to and influence the interaction between memory and attention.

5 Decoding Working Memory Content from Attentional Biases

Thus far, Chapters 2 through 4 have explored the impact of WM content on visual attention with respect to: (1) the traditional dichotomy of top-down and bottom-up attentional control, and (2) the question of under what conditions WM does or does not guide attention. All of these studies point to a powerful and distinct influence of memory-based attentional guidance, supporting the close reciprocity—and perhaps shared mechanisms and representations (see Chapter 1)—between WM and attention. The current chapter tests the strength and reliability of the relationship between WM and attention by asking whether we can infer what a person is holding in mind, based on their attentional bias. The following “behavioral mind-reading” study uses multivariate pattern classification techniques to show that memory-based attentional guidance is diagnostically robust enough to successfully decode the specific contents of WM, even so far as to generalize across individuals. These findings emphasize the robustness of memory-based guidance and further support a very tight link between WM and attention. These data were presented as a conference talk at the 2015 Annual Workshop on Object Perception, Attention, and Memory and are currently being prepared to submit for publication.

5.1 Introduction

Our thoughts guide our behavior. For example, what we are currently thinking about can influence how we interact with the external environment. This idea is

prominent in the process of visual search—finding targets among distractors—in which a mental representation of the target is thought to be held in WM (i.e., the active maintenance and manipulation of internal information), facilitating the detection of target-matching features in the visual environment (e.g., Desimone & Duncan, 1995). Moreover, this link between WM and visual attention (the selective processing of a subset of information in the visual field) is not limited to intentional uses of WM, as it has been shown repeatedly that maintaining *any* information in WM, even information unrelated to a search target, is sufficient to guide visual attention to memory-matching items in the environment (see Chapter 1).

Specifically, studies of memory-based attentional guidance typically use dual-task paradigms, in which participants remember an item (e.g., a colored shape) while performing an intervening but unrelated visual search task during the delay period between the WM cue and subsequent probe (e.g., Soto et al., 2005). Critically, the memory item can reappear in the search display, either coinciding with the location of a target (i.e., “valid” search) or with the location of a distractor (i.e., “invalid” search); or fail to reappear in the display at all (i.e., “neutral” search). The canonical finding is that search is speeded when the memory item matches the target location and slowed when the memory item matches a distractor location (relative to the neutral condition; see Soto et al., 2008). This attentional biasing by WM content occurs even under conditions where it is reliably detrimental to search performance—for instance, in tasks composed of 100% invalid searches (e.g., Dowd et al., 2015a; Kiyonaga et al., 2012).

One explanation for these validity effects hypothesizes that holding an item in WM involves the recruitment or activation of the sensory representation of that item, which in turn facilitates the processing of memory-matching items in the sensory environment (Desimone & Duncan, 1995). This view implies a tight (and perhaps obligatory; Chun, 2011) link between what we think about and where we attend, which has profound implications for our understanding of cognition and behavior. In fact, if this link were truly reliable, then one should be able to infer what a person is holding in mind based on their attentional bias.

Evidence for memory-based attentional guidance has been found to be robust across various paradigms and populations, but only in the form of mean differences in search response times when aggregated over many trials and subjects. However, mean group effects do not necessarily translate into successful classification of categorical data at the single-trial level (see Franz & Luxburg, 2015), as aggregate effects can be driven by a subset of trials or subjects, and/or the signal-to-noise ratio of a given effect measurement might be too small for reliable single-trial inference. Thus, in the present study, we asked whether the effects of WM maintenance on attentional orienting are robust enough to be used diagnostically for inferring a person's mental content. Specifically, we tested whether we could apply multivariate pattern classification techniques to the pattern of response times observed in an unrelated search task during the WM delay period to "decode" the specific item an individual is holding in WM on a particular trial. We first attempted this feat of "behavioral mind-reading" at the within-

subject level, asking if the WM content on a particular trial could be decoded by classifiers trained on all other trials of that subject; we then sought to test the generalizability of attentional guidance by WM at the between-subject level, where classifiers trained on data from a group of subjects were applied to infer the WM content of a single trial from a new, previously unseen individual.

5.2 Method

5.2.1 Participants

100 Amazon Mechanical Turk workers (mean age = 35.2 years, $SD = 10.2$; 50 men, 50 women) participated in exchange for monetary compensation (US\$3.50–4.00). All participants provided informed consent in accordance with the policies of the Duke University Institutional Review Board. Prior to collecting data, we set a data inclusion threshold of 85% accuracy on both memory and search tasks; this performance threshold was set in order to maximize the number of instances (correct memory responses) and number of features (correct search response times) fed into the different classifiers. Thus, data from an additional 19 participants were excluded for below-threshold performance: 8 for memory accuracy, 3 for search accuracy, and 8 for both memory and search accuracies.

5.2.2 Stimuli and Procedure

As illustrated in Figure 13, the dual-task paradigm consisted of a delayed match-to-sample WM test with an intervening visual search task comprised of 12 successive visual searches. We included 12 separate searches within each delay period in order to

form a multidimensional space of classifier inputs, described further in this section. Each trial began by presenting a central fixation cross for 500 ms, followed by the WM cue item under the text “Remember this color!” for 1,000 ms. Each WM cue was a single colored circle (113×113 pixels) in one of four colors (RGB values: red = 227, 2, 24; blue = 48, 62, 152; green = 95, 180, 46; yellow = 251, 189, 18). The WM cue was followed by a blank delay for 2,000 ms, then by a series of 12 visual searches. Each visual search presented a central fixation dot for 500 ms, followed by a search array for 300 ms, and then a blank screen until a response was recorded (or up to 1,200 ms). Each search array consisted of two colored circles presented along the upper and lower halves of the vertical midline, and each circle surrounded a white line (10×75 pixels). One line (i.e., the search distractor) was vertical, while the other line (i.e., the search target) was tilted 45° to the left or right. The visual search stimuli were thus distinct from, but embedded within, the class of stimuli that could match WM. Participants were instructed to indicate the orientation of the line via keypress to respond “left” or “right” tilted, as quickly and accurately as possible. Target location and orientations occurred equally often in a randomized order. After the 12th search, a four-alternative forced-choice memory probe array was presented under the text “Which color did you remember?” for up to 3,000 ms; participants were instructed to report which color matched the initial WM cue color. The color-response mappings of the memory probe array changed across trials. Feedback was presented following all memory responses (i.e., correct or incorrect) and only after erroneous (i.e., timeouts or wrong keypresses) search responses. Participants

completed 4 practice trials, followed by 80 experimental trials (20 of each WM color, randomly intermixed) across 10 blocks.

Importantly, the series of 12 visual searches during the WM delay contained one of each possible combination of two-color (${}_4C_2 = 6$) and target-distractor (2) arrangements (Figure 14), resulting in 12 unique search arrays that were presented in randomized order. We specified 12 as the minimum number of visual search arrays, such that for each trial, no matter the WM color, there was always a decoupling between the specific validity relationships and that memory color. In other words, the full set of 12 possible search arrays was completely agnostic to the WM color. Thus, although each WM color cue was followed by the exact same 12 search combinations (but randomized in order), the corresponding validity condition of each search changed, depending on the specific WM color (but always resulting in 3 valid, 3 invalid, and 6 neutral searches). Figure 13 depicts an example trial sequence in which the memory color is red, followed by 3 valid (top row; target is red), 3 invalid (second row; distractor is red), and 6 neutral (bottom two rows; neither target nor distractor is red) searches. This changing pattern of validity condition—and subsequent attentional guidance effects—is thus the targeted behavioral index by which to decode the contents of WM.

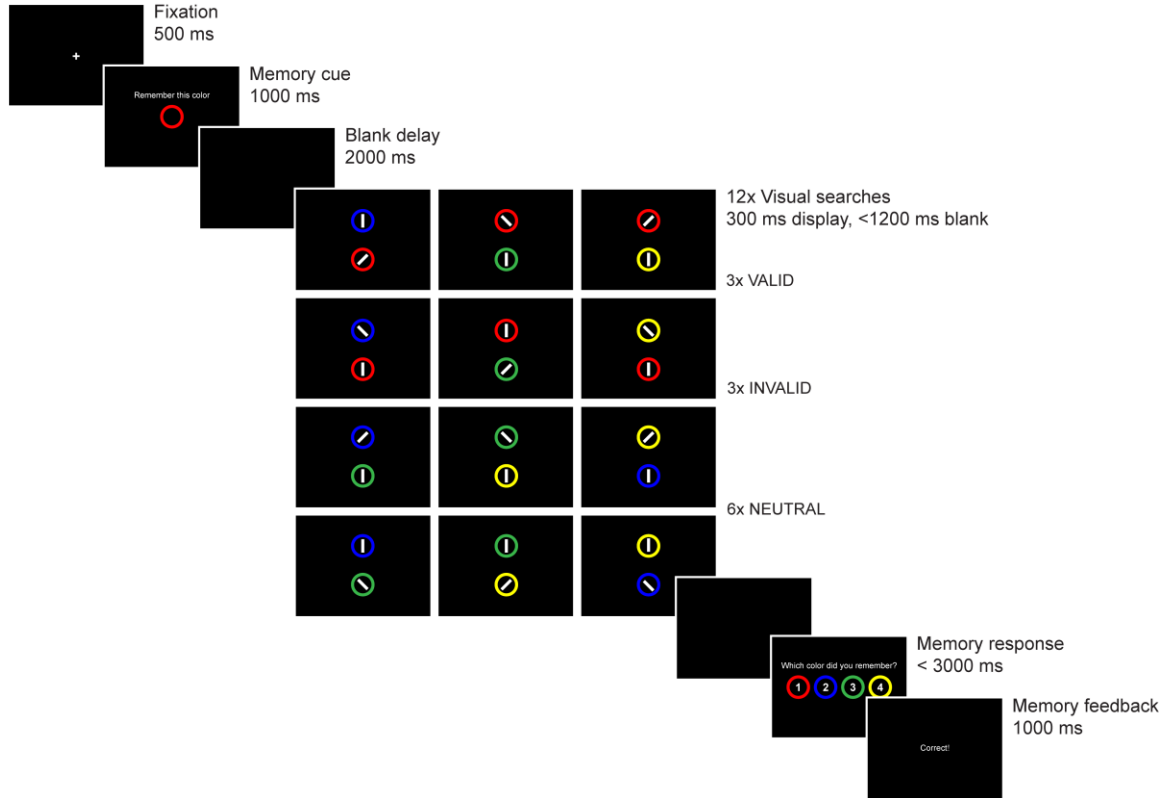


Figure 13: Example trial sequence. Participants were to remember a cue color (red, blue, green, or yellow) while performing a series of 12 intervening visual searches, in which they reported the direction of the tilted line. Then participants were given a four-alternative forced-choice memory probe array and asked to report which color matched the initial memory cue. In this example, the cue color is red; in the series of 12 searches, 3 are valid (top row; target is red), 3 are invalid (second row; distractor is red), and 6 are neutral (bottom two rows; neither target nor distractor is red).

5.2.3 Classification Analysis

Our key analysis focused on using multivariate pattern classification to decode the contents of WM from behavior. The multidimensional feature space consisted of search response times, organized by the 12 search combinations within each trial; each trial thus contributed a vector of 12 search response times, and was labeled by the corresponding

WM color (Figure 14). Trials with incorrect WM responses (4.1%) were excluded from the classifier, and incorrect search responses (3.7% of all visual searches) were replaced with non-informative mean values, which allowed us to discount individual search errors without excluding the entire trial vector. Response times were detrended across blocks and rescaled within each feature column to [0 1], so that each feature was encoded with a similar dynamic range (i.e., relative to a given subject's response time means and variances). To account for potential non-linearity, we also included the squares and square-roots of each feature column, resulting in a 36-dimensional feature space.

We analyzed three separate classifiers: distance analysis, logistic regression, and linear support vector machines. In the first and simplest analysis, trials were classified according to the smallest Euclidean distance between a test vector and the mean training vector for each label (i.e., WM color). The second analysis employed multinomial logistic regression to calculate probabilistic predictions for each label, using the “glmnet” package in R (Hastie & Junyang, 2014). Finally, we also implemented linear support vector machines (SVM), which are commonly used in multivariate decoding analyses of neural data (Haynes, 2015). For multi-class SVM classification, we used the “e1071” package in R to train six binary one-against-one classifiers; for each test instance, each binary classifier output a winning label, and the predicted WM color for that test instance was the most winning label across all binary classifiers (Dimitriadou, Hornik, Leisch, Meyer, & Weingessel, 2006).

All classifiers were evaluated with a cross-validation approach, wherein a subset of trials was used to train the classifier, and the WM color was predicted for an independent set of test trials. For within-subject classification, we adopted a leave-one-trial-out cross-validation scheme. Thus, within each participant's data, the classifier was trained on the vectors corresponding to all but one trial, then tested on its ability to decipher the WM label on that one trial (Figure 14). The training set was iteratively shuffled so that each trial served once as the testing set, and classifier accuracy for a given participant reflected the average classifier performance across all iterations (up to 80 iterations). For between-subject classification, we adopted a similar leave-one-subject-out cross-validation scheme, in which the classifier was trained on the data corresponding to all but one participant, whose data then served as the testing set. This classification was still performed at the level of a single trial within each testing set, not on average values across trials. Due to computational limitations, the between-subject SVM analysis used a leave-10-subjects-out (i.e., 10-fold) cross-validation scheme (while the other two classifiers use a leave-one-subject-out scheme). All between-subject analyses were also iteratively shuffled until each participant had served in the testing set, resulting in a single average classification accuracy across all 100 (or 10 for SVM) iterations.

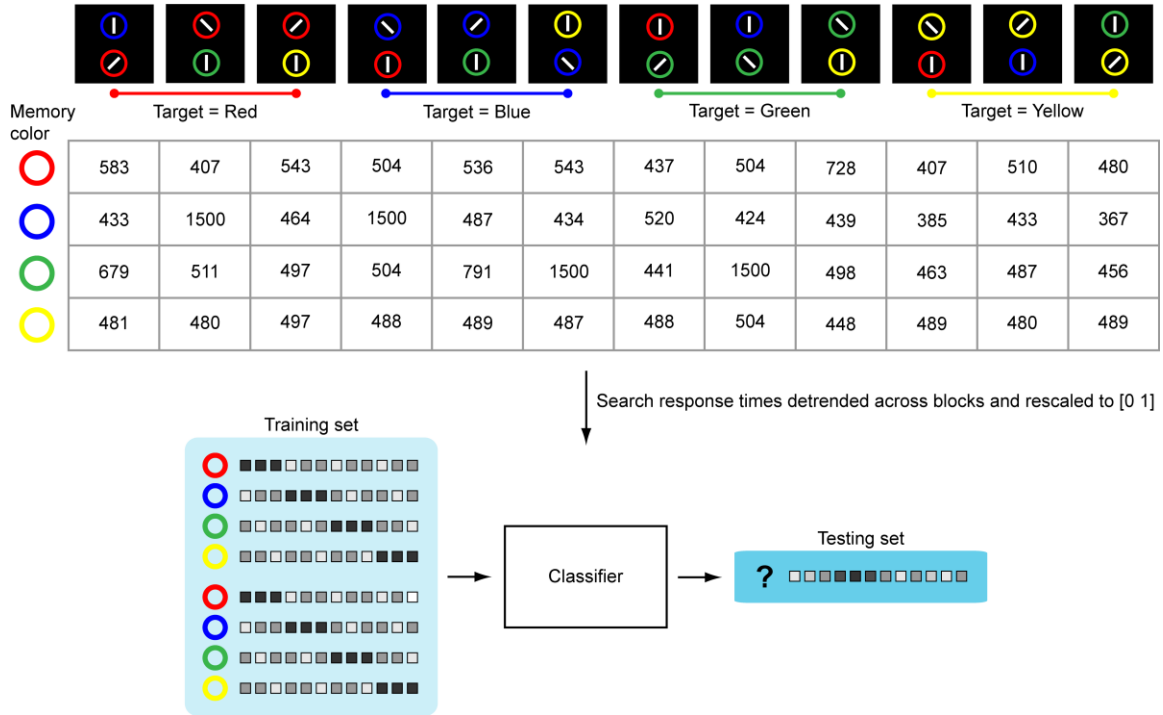


Figure 14: The multidimensional feature space was organized according to the 12 possible two-color (6) and target-distractor (2) arrangements (*top*). Classifier inputs were search response times for each feature. Thus, each trial contributed a vector of 12 search response times, labeled by the corresponding memory color; four example input vectors are shown in the grid. Response times were detrended across blocks and rescaled within each feature column to [0 1]. For within-subject classification analyses, we adopted a leave-one-trial-out cross-validation scheme; within each participant’s data, the classifier was trained on the vectors corresponding to all but one trial (*training set*), then tested on its ability to decipher the memory label on that one trial (*testing set*).

5.3 Results

5.3.1 Mean Performance Data

Both memory accuracy ($M = 95.9\%$, $SD = 4.0\%$) and search accuracy ($M = 96.3\%$, $SD = 3.2\%$) were near ceiling. A repeated-measures ANOVA with a Greenhouse-Geisser correction revealed that response times for correct visual search responses

differed significantly across validity, $F(2, 198) = 250.4, p < .001, \eta^2 = .72$. Search times were overall faster on valid trials ($M = 579$ ms, $SD = 85$ ms), $t(99) = 10.9, p < .001$, and slower on invalid trials ($M = 606$ ms, $SD = 87$ ms), $t(99) = 14.9, p < .001$, compared to neutral trials ($M = 588$ ms, $SD = 85$ ms), replicating canonical validity effects such that attention was biased toward memory-matching contents (i.e., the target location in valid trials and the distractor location in invalid trials; Figure 15a).

While previous dual-task studies have demonstrated group mean validity effects using a single intervening search between the WM cue and probe (e.g., Dowd et al., 2015a; Olivers et al., 2006; Soto et al., 2005), we here employed series of 12 intervening searches, in which the pattern of validity condition was randomly ordered. To test whether memory-based guidance did in fact last across the entire search series, we performed a repeated-measures ANOVA across validity and trial half (i.e., whether a particular search fell in the first or second half of the series), which revealed main effects of validity, $F(2, 198) = 246, p < .001, \eta^2 = .71$, and trial half, $F(1, 99) = 8.2, p = .005, \eta^2 = .08$, such that search times were overall slower in the second half of the search series. A significant interaction effect, $F(2, 198) = 8.4, p < .001, \eta^2 = .08$, indicated that the size of the validity effect was reduced (although still significant and robust, $F(2, 198) = 121.3, p < .001, \eta^2 = .55$) through the second half of the search series (Figure 15b). Thus, maintaining an item in WM impacted attentional guidance across multiple subsequent searches, validating the inclusion of entire patterns of response times as classifier features.

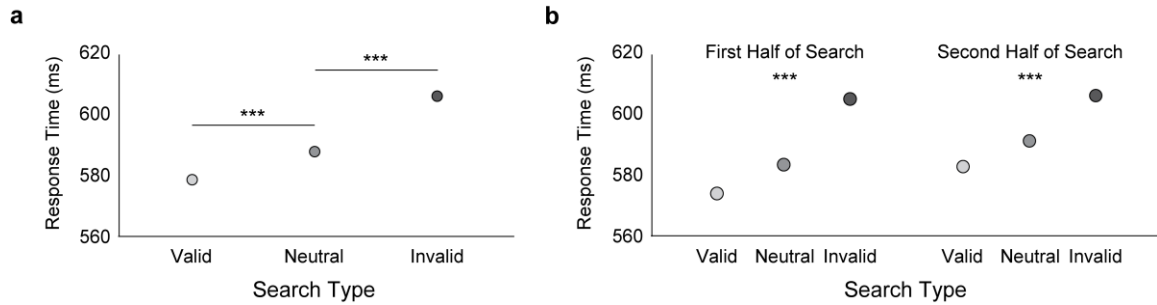


Figure 15: Mean search response times revealed significant validity effects. (A) Response times were overall faster on valid searches and slower on invalid searches, compared to neutral searches. (B) Validity effects were still significant after splitting the series of 12 searches into first and second halves, indicating that a single instance of WM maintenance could impact attentional guidance across multiple subsequent searches. The diameter of each dot represents 95% confidence intervals. * $p < .001$.**

5.3.2 Classification Results

Our main analyses focused on using multivariate pattern classification to decode the contents of WM from behavioral indices of attentional guidance. We first classified the contents of WM within each participant's data, using a leave-one-trial-out cross-validation scheme. Decoding accuracy for the three multivariate classifiers is shown in Figure 16a, where a violin plot illustrates the distribution of each classifier's accuracy over the 100 participants. All three methods accurately classified the trial-specific WM contents within a single participant's set of data, as revealed by one-sample t -tests compared to chance-level performance of 25% (considering four possible memory colors). Overall within-subject decoding accuracy for the distance analysis was 30.7%, $t(99) = 7.1$, $p < .001$, $d = .71$; for linear SVM was 32.1%, $t(99) = 11.8$, $p < .001$, $d = 1.18$; and for logistic regression was 32.0%, $t(99) = 10.9$, $p < .001$, $d = 1.09$. There were no

significant differences across classification methods, as estimated by Wilcoxon signed-ranks tests with a Bonferroni-corrected significance threshold of .017 (Demšar, 2006).

We also asked whether the classification algorithms in fact made use of the degree of memory-based attentional guidance, as we hypothesized, or some more complicated pattern of response time covariance. Thus, we tested whether classifier evidence was related to a coarse measure of the robustness of attentional guidance, as calculated from search response times: (invalid – valid) / neutral. This normalized difference score reflects the absolute response time difference as a proportion of neutral response times, such that positive values indicate greater attentional biasing (Dowd et al., 2015b). Individual differences in classification accuracy for all three approaches were significantly predicted by the magnitude of attentional guidance, with R^2 values ranging from .22 to .33, $ps < .001$ (Figure 16b). In other words, greater average attentional biasing by WM contents was associated with better decoding, suggesting that the trained classifiers indeed capitalized on patterns of search validity effects that emerged across each trial.

We next asked a more profound question of the robustness of attentional guidance by WM, namely, whether a classifier trained on data from a group of individuals could predict single-trial WM content in another, new individual. This between-subject approach thus tests whether patterns of attentional guidance are consistent and generalizable enough to decode WM contents across individuals. Figure 17 shows accuracy distributions of the three between-subject classification methods. Again, all

three classifiers decoded the specific WM contents *across* the group of 100 participants significantly above chance, according to binomial tests (Pereira, Mitchell, & Botvinick, 2009). Overall between-subject decoding accuracy for the distance analysis was 36.3%, $p < .001$, 95% CI [0.352 0.374]; for linear SVM was 36.1%, $p < .001$, 95% CI [0.350 0.372]; and for logistic regression was 36.5%, $p < .001$, 95% CI [0.354 0.376]. Wilcoxon signed-rank tests revealed no significant differences across between-subject classification methods. Note also that the generally higher classifier performance for between-subject compared to within-subject approaches was an artifact of higher volumes of input data for the between-subject analyses.

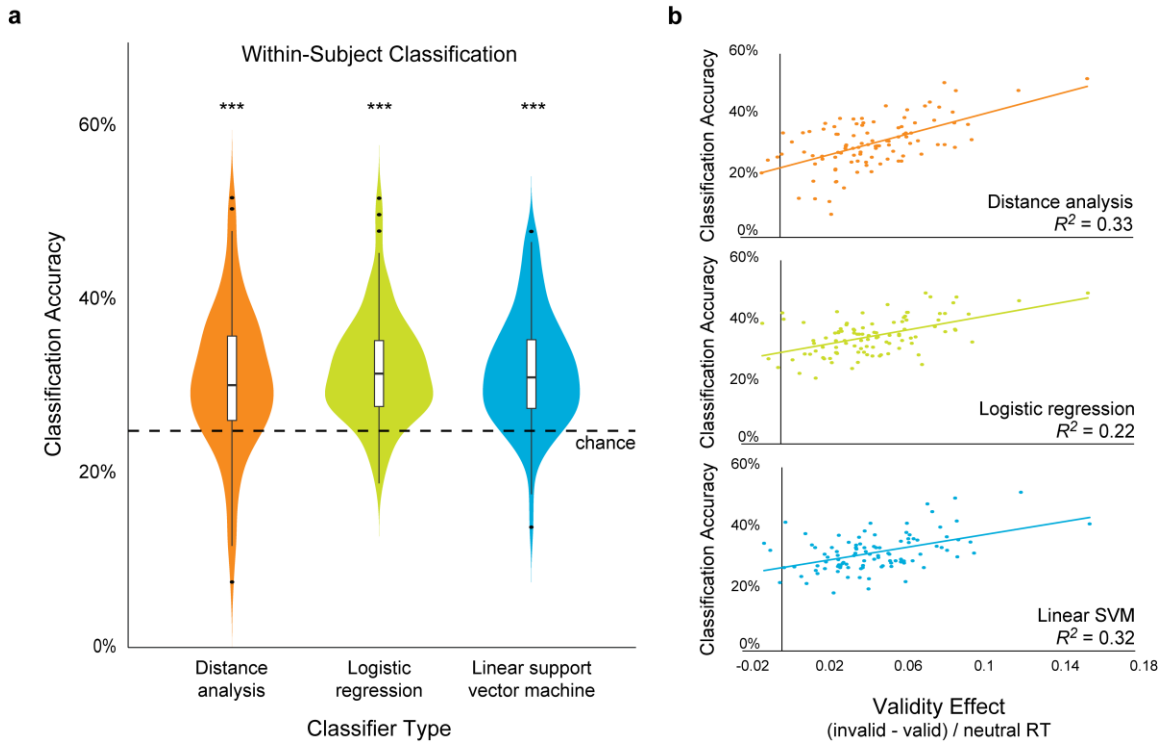


Figure 16: Within-subject classification results. (A) Across three classifiers, within-subject decoding accuracy was significantly above chance (dashed line at 25%). For each classifier, a violin plot illustrates the distribution of that classifier’s accuracy over the 100 participants. The boxes mark the middle two quartiles of classification performance, and the whiskers represent the full range of classification performance (outliers are marked with a dot). (B) Individual differences in classification accuracy for all three approaches were significantly predicted by validity effect (i.e., magnitude of attentional guidance), suggesting that trained classifiers capitalized on patterns of search validity effects that emerged across each trial. * $p < .001$.**

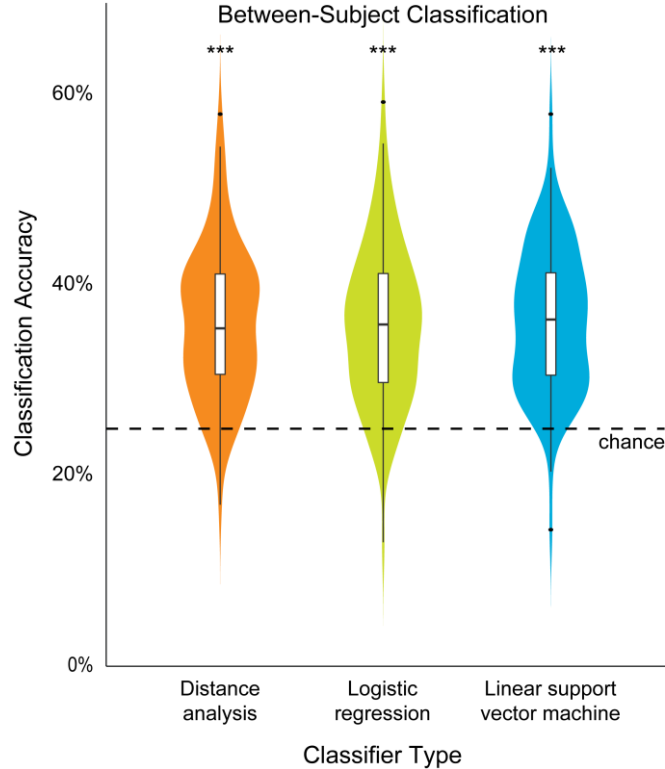


Figure 17: Between-subject decoding accuracy was also significantly above chance for all three classifiers, as evaluated by binomial tests against chance (dashed line at 25%). For each classifier, a violin plot illustrates the distribution of that classifier’s accuracy over the 100 participants. The boxes mark the middle two quartiles of classification performance, and the whiskers represent the full range of classification performance (outliers are marked with a dot). * $p < .001$.**

5.4 Discussion

Recent behavioral and neural evidence supports the idea that WM recruits the same sensory representations that are activated by attending to external stimuli (see Sreenivasan et al., 2014a), such that holding an item in WM automatically facilitates the processing of memory-matching items in the sensory environment (Desimone & Duncan, 1995). Here, we tested the hypothesis that if the link between WM and attentional orienting (i.e., in the form of memory-based attentional guidance) is so highly robust,

then we should theoretically be able to infer what a person is holding in WM, based on their attentional bias. Thus, the current study applied multivariate pattern classification techniques to show that the contents of WM can be reliably inferred from behavioral indices of memory-based attentional guidance.

Here we replicated canonical memory-based attentional biasing effects, in which search response times were speeded to memory-matching target locations and slowed by memory-matching distractor locations. We also showed that these attentional guidance effects could extend across 12 intervening searches, demonstrating robust WM biases that continued across multiple attentional responses and prolonged search sequence durations (cf., Kiyonaga & Egnér, 2014). These simple behavioral outputs (i.e., response times) were then fed into multivariate pattern classification analyses in order to predict the contents of WM. Across three separate within-subject classification methods, patterns of attentional guidance within an individual's search task data could reliably decode which specific color the individual was maintaining in WM for a single, particular trial, highlighting the specificity and reliability of this link between WM and attention. Classifier evidence was furthermore significantly predicted by the magnitude of validity effects, indicating that the classifiers were indeed decoding WM contents from patterns of memory-based attentional guidance.

To further emphasize the robustness of the relationship between WM and attention, the current study also showed that memory-based attentional guidance is both consistent and generalizable enough that a classifier trained on data from a group of

individuals can predict single-trial WM content in another, new individual. Again, across three separate between-subject classification methods, the specific contents of WM—for a completely independent participant—could be reliably inferred from patterns of attentional guidance within a separate group’s search task data. Combined with the within-subject classification results, these between-subject findings suggest that while the impact of WM maintenance on delay-period attentional deployment is subject to individual variation, memory-based attentional biasing nevertheless reflects a common, shared cognitive mechanism across individuals.

In cognitive neuroscience, multivariate pattern-based analyses of brain activity have been used to reveal a person’s mental states—such as what the individual is perceiving, attending to, or remembering (e.g., Harrison & Tong, 2009; Kamitani & Tong, 2005; see Tong & Pratte, 2012). Specifically, the contents of WM have been successfully decoded from within-subject neural activity, with variable ranges of accuracy (e.g., Albers et al., 2013; Emrich, Riggall, LaRocque, & Postle, 2013; Harrison & Tong, 2009; Serences et al., 2009). For example, Serences et al. (2009) classified red versus green stimuli in early visual cortex (V1; 62 voxels) with nearly 65% accuracy (compared to 50% chance), while Albers et al. (2013) classified three grating orientations in V1-V3 (360 voxels) with 54% accuracy (compared to 33% chance), both across different time periods within the WM delay. While it is difficult to directly compare classification accuracies across different feature inputs and task designs, our within-subject results (32% compared to 25% chance) are derived from inexpensive, purely

behavioral inputs. Furthermore, we demonstrate that decoding of WM content generalizes reliably to *completely novel* individuals, a notoriously difficult problem for neuroimaging-based classifiers due to the methodological hurdle of normalizing morphological idiosyncrasies across participants (see Haxby, Connolly, & Guntupalli, 2014; Shinkareva et al., 2008). The current levels of classification accuracy were also achieved with a relatively small multidimensional space of 36 features: 12 search combinations, as well as the squares and square-roots of each combination. Compared to neuroimaging studies, in which classification accuracy across anatomically-defined regions of interest increases as more voxels are included, up to about 500 voxels (although the benefits for greater numbers of features depend on the classifier, number of instances, and other factors; see Etzel, Gazzola, & Keysers, 2009), we employed a very low number of features, which again supports the notion that attentional biasing is a highly informative probe of WM content.

Multivariate classification analyses of neural activity are based on the idea that we can link neural representational space (e.g., patterns of local activity across voxels or neurons) to stimulus representational space (e.g., patterns of stimulus contrast; Haxby et al., 2014). However, the current study instead links stimulus representational space to a multidimensional space as defined by behavioral performance measures. While multivariate pattern analyses have been previously applied to behavioral eye movement statistics to classify the type of viewing task that a participant is engaging in (Borji & Itti, 2014; Henderson, Shinkareva, Wang, Luke, & Olejarczyk, 2013; Kardan, Berman,

Yourganov, Schmidt, & Henderson, 2015; but see Greene, Liu, & Wolfe, 2012), these studies took advantage of systematic differences in a task-relevant behavior (i.e., eye movements) to identify different task strategies—in contrast, the current study focuses on classifying specific mental (WM) representations, as opposed to entire task sets, based on behavioral data from an *unrelated* attentional task. Remember that in the current dual-task design, the appearance of memory-related visual objects (i.e., colored circles) in visual searches was task-irrelevant; they had no bearing on the search task itself (i.e., searching for a target line), and the same 12 combinations of two-color and target-distractor arrangements were *always* presented within every trial. Thus, there was a decoupling between the particular memory color and the subsequent validity relationships, such that there was no strategic incentive to attend to the changing pattern of validity conditions within each trial. The ensuing attentional guidance effects, then, were unlikely to be strategically-driven (see also Soto et al., 2008).

In conclusion, the present study demonstrates for the first time that memory-based attentional guidance effects are so robust that the specific contents of WM can be reliably inferred from patterns of attentional guidance in an unrelated search task. Furthermore, WM contents were also successfully decoded across individuals, suggesting that WM guidance is driven by a common, shared mechanism—specifically one in which the maintenance of any information in WM activates memory-matching sensory features, providing an automatic and robust advantage for further attentional processing in the visual field (Desimone & Duncan, 1995). More generally, the current results provide

support for a very tight link between what we keep in mind and what we attend to, contributing to a fuller theoretical understanding of the interactions between WM and visual attention.

6 Neural Working Memory Representations are Modulated by Attentional Demands

Recent theories suggest that visual WM and visual attention operate upon shared selection mechanisms and sensory representations, as supported by both behavioral and neural evidence (see Chapter 1). While Chapter 5 provided empirical support for the tight and reliable link between these two constructs in terms of memory-based attentional guidance, the current chapter presents a neuroimaging study that tests the neural competition that occurs under simultaneous demands on WM and visual attention. Here, we employed human fMRI to investigate how neural representations of WM content (as assessed by multivariate pattern classification) are affected when attention is occupied by visual search—at varying levels of difficulty—during the delay interval of a WM task. The current results demonstrated behavioral tradeoffs between the two domains, as well as a degradation of neural WM representations when attention had to be diverted to a concurrent search task. Thus, the present study offers novel neural evidence for the hypothesis that WM and visual search compete for shared and capacity-limited attentional resources. This chapter further discusses how the data fit into models of visual WM with multiple representational states or schema. These data were presented as a conference talk at the 2015 Annual Meeting of the Vision Sciences Society and have been submitted for publication.

6.1 Introduction

We are constantly called upon to maintain information temporarily in mind, but this WM must also operate in the face of immediate and variable demands for our attention in the environment (e.g., rehearsing a shopping list while navigating heavy traffic). While attention has typically been described as the selective processing of information that is currently available to the senses—and WM conversely acts on information unavailable to the senses—a large body of evidence indicates that demands on WM and attention reciprocally influence one another (Awh et al., 2006; Awh & Jonides, 2001; Gazzaley & Nobre, 2012). This has encouraged the reconceptualization of WM as internally-oriented attention that endogenously activates perceptual representations in much the same way as attention to external stimuli would (Chun, 2011; Chun & Johnson, 2011). Accordingly, WM-related sustained increases in mean neural population activity, as indexed by univariate fMRI signal (in dorsolateral prefrontal cortex, for instance), were once assumed to represent the information being held in WM; however, those responses are now considered to reflect attentional control over sensory regions that represent the information content itself (see Sreenivasan et al., 2014a; but see also Ester et al., 2015). In other words, even without sensory input, attention is recruited to activate sensory representations for the purpose of WM.

Recent multivariate neural evidence also supports this “sensory recruitment” model of WM, whereby short-term representations are maintained via distributed patterns of activity within the same sensory cortical regions engaged by perceptual attention

toward that content (e.g., area MT for memory of moving dot arrays; Riggall & Postle, 2012). The orientation of a Gabor grating maintained in WM, for instance, can be successfully decoded or reconstructed based on early visual cortex activity patterns derived from actually perceiving oriented stimuli (Ester et al., 2013; 2015; Harrison & Tong, 2009). If WM content is maintained in sensory cortices, via attention-dependent activation, a critical question is: What happens to such internally-attended information in the face of incoming sensory input and concurrent attentional demands? Behavioral studies have shown that WM often suffers when attention is otherwise occupied during the WM delay (Fougnie & Marois, 2009), and the extent of that impairment scales with the time-consumption of the intervening task (Barrouillet, Portrat, & Camos, 2011). Here, we employed fMRI to determine how this competition between internal and external attentional demands impacts the sensory representations of WM content, which is presently unknown.

Many behavioral studies suggest that concurrent attentional demands can alter the “activation status” of a WM representation, relegating it to a distinct format outside of an internal focus of attention (e.g., Stokes, 2015), into which it can be reinstated when it becomes task-relevant again (Gunseli et al., 2015; van Moorselaar et al., 2015). Correspondingly, the active neural trace of a WM representation—as detected by multivariate classification methods—is modulated by internal shifts of attention across a trial; immediately task-relevant representations elicit measurable neural signatures, while

evidence for task-irrelevant memory representations becomes undetectable (LaRocque, Lewis-Peacock, & Drysdale, 2013; Lewis-Peacock et al., 2012).

If activation in visual WM occurs by directing attention internally toward perceptual representations, then directing attention outwardly toward a visual task should similarly modulate WM representational information, and neural activation patterns in regions that represent WM content should become less discriminable. Here, we used multivariate pattern classification of fMRI data to test the prediction that WM decoding is impacted when attention is occupied by visual search—at varying levels of difficulty—during the delay interval of a WM match-to-sample task. If WM and visual search both rely on attention, neural WM representations should be degraded during visual search, and that degradation should be even more pronounced when a more difficult visual search condition diverts attention away from WM maintenance for a longer period of time.

6.2 Methods

6.2.1 Participants

Thirty healthy volunteers gave written informed consent to participate in accordance with the Duke University Institutional Review Board. All participants were fluent in English, reported normal or corrected-to-normal vision, and were compensated US\$20 per hour for their participation. Two participants were excluded for missing data, leaving 28 participants in the final analyses (16 male; mean age: 30; range 18–45).

6.2.2 Design

The experimental protocol was designed to independently vary “internal” (i.e., WM) and “external” (i.e., visual) attentional load in a fully balanced 2 (WM load: 1 item vs. 2) \times 2 (visual search difficulty: easy vs. hard) factorial design. The task comprised a delayed match-to-sample WM test, with a sequence of delay-spanning visual searches (Figure 18a). We employed WM sample stimuli with known cortical sensitivities (i.e., faces and houses), so that we could examine the discriminability of visual cortical WM representations, via classifiers trained on the WM category, in the fusiform face (FFA; Kanwisher, McDermott, & Chun, 1997) and parahippocampal place areas (PPA; Epstein & Kanwisher, 1998). Across different trials, participants had to maintain either one (low WM load) or two (high WM load) faces or houses for a later memory probe (Figure 18b). During the WM delay, participants performed a series of four visual searches for a perfectly vertical target stimulus among horizontal (easy search) or slightly tilted (hard search) distractors (Figure 18c). We borrowed this attentional manipulation approach from the time-based resource-sharing model of WM storage and processing, whereby a harder visual search task should occupy attention—that would otherwise be dedicated to WM maintenance—for a longer period of time (Barrouillet et al., 2011). Our main analyses focus on this search epoch of the trial, as we were interested in characterizing how WM would be impacted by this secondary demand. In order to produce a balanced design, wherein WM category classification would be uncontaminated by overlapping visual input, visual search stimuli were either bodies (which have been demonstrated to

preferentially recruit the extrastriate body area [EBA; Downing, Jiang, Shuman, & Kanwisher, 2001]) or tools (Chao, Haxby, & Martin, 1999). The design thus produced four main conditions: low WM/easy search, low WM/hard search, high WM/easy search, and high WM/hard search.

6.2.3 Dual-Task WM/Visual Search Procedure

The task was programmed and presented in Matlab using the Psychophysics Toolbox (Brainard, 1997). Face stimuli were 144 unique grayscale images of male and female faces, drawn from several databases (Endl, Walla, Lindinger, & Lalouschek, 1998; Kanade, Cohn, & Tian, 2000; Lundqvist, Flykt, & Öhman, 1998; Minear & Park, 2004; Oosterhof & Todorov, 2008; Tottenham et al., 2009), and cropped to include only the “eye and mouth” region. House stimuli were 144 unique grayscale exterior images drawn from local real estate websites. Visual search stimuli were 16 male and female bodies, with heads cropped (Downing et al., 2001), and 16 tools (hammers and wrenches) drawn from freely available online sources. Stimuli were displayed on a back-projection screen against a neutral gray background (RGB values: 128 128 128), and viewed through a mirror mounted to the head coil simulating a viewing distance of approximately 80 cm. Behavioral responses were executed with the left and right hands on MRI-compatible response boxes.

Each trial began with a variable inter-trial interval, followed by the WM sample for 2 s (Figure 18a). Low load WM samples consisted of a single, centrally-presented face or house. High load WM samples consisted of either two faces or two houses

presented side-by-side (Figure 18b). After a variable inter-stimulus interval, a series of four visual search displays appeared for 1.5 s each, separated by 500 ms fixation intervals, producing a search sequence lasting 8 s in total (Figure 18a). Each search array contained four stimuli (either all tool or all body images) at the corners of an imaginary square. The target stimulus was perfectly vertical, while three distractors were tilted to the left or right. The task was to indicate whether the target stimulus was oriented right-side up or upside down. On easy search trials, the distractors were perfectly horizontal (i.e., tilted 90° to the left or right), making them easily discriminable from the vertical target. On hard search trials, on the other hand, distractors were slanted only 15° to the left or right, making them more similar to the vertical target, and thus more difficult to discriminate (Treisman & Gelade, 1980); the amount of perceptual input was nevertheless matched in easy and hard search conditions (Figure 18c). All searches in a given trial were of the same difficulty level. The search sequence was followed by a variable inter-stimulus interval, then a WM probe for 3 s. Participants were asked to rate their confidence, on a 4-point scale, that a single WM probe item was either a match (50% of trials) or non-match to an item from the WM sample set. Underneath the probe image, a visual guide instructed which finger of the left hand should be used to indicate a response of either “Definitely the same,” “Maybe the same,” “Maybe different,” or “Definitely different” (Figure 18a).

WM samples were selected in random order and never repeated across the experiment, except as matching probes. Visual search stimuli, locations, and orientations

were selected at random on every trial. The duration of inter-trial as well as pre- and post-search inter-stimulus intervals were jittered between 2.5–5 s (step-size = 500 ms), selected at random from a pseudo-exponential distribution (Dale, 1999). Participants completed a practice run of 16 trials outside of the scanner, then 9 experimental runs of 16 trials each inside the scanner, for a total of 144 experimental trials. All trial conditions occurred equally often, and in random order, both within and across runs.

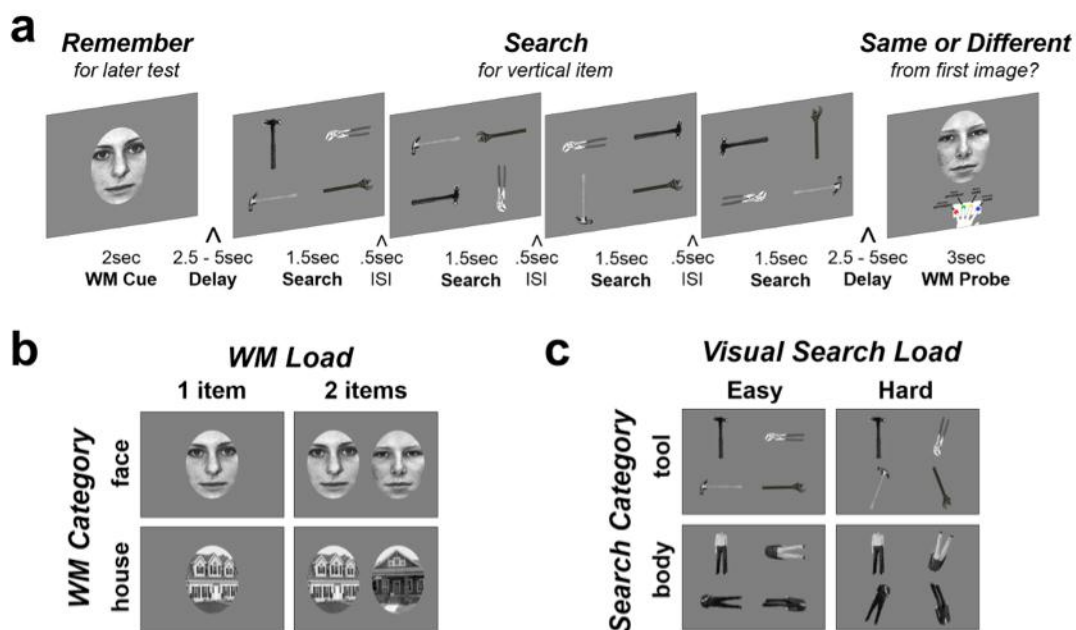


Figure 18: Behavioral task design. (A) During the delay interval of a match-to-sample WM task, participants completed a series of four visual searches. (B) WM load conditions: Participants maintained either 1 (low) or 2 (high) face or house stimuli in WM. (C) Visual search load conditions: Participants searched for the vertical body or tool target among horizontal (easy) or tilted (hard) distractors.

6.2.4 Functional Localizer Procedure

Participants also completed a functional localizer task to define cortical regions of interest (ROIs) that preferred each of the WM and visual search stimulus categories (i.e., faces, houses, bodies, and tools). Each stimulus category was presented in separate blocks; each block entailed a series of 15 images, centrally-presented for 750 ms, and separated by a 250 ms fixation. Participants were asked to make a button response to direct repetitions of a specific stimulus (i.e., 1-back task). The run contained 16 blocks (4 of each condition) separated by 10 s inter-block intervals, occurring in random order.

6.2.5 Image Acquisition

Functional data were recorded on a 3.0 tesla GE MR750 scanner, using a gradient-echo, T2*-weighted multi-phase echoplanar imaging (EPI) sequence. Forty contiguous axial slices were acquired in interleaved order, parallel to the anterior-posterior commissure (AC-PC) plane (voxel size: $3 \times 3 \times 3$ mm; repetition time [TR] = 2 s; echo time [TE] = 28 ms; flip angle = 90° ; FOV: 24 cm). Structural data were obtained with a 3D T1-weighted fast inversion-recovery-prepared spoiled gradient recalled (FSPGR) pulse sequence, recording 154 slices of 1 mm thickness and an in-plane resolution of 1×1 mm.

6.2.6 fMRI Analyses

All analyses were done in Matlab with SPM8 scripts (Wellcome Department of Imaging Neuroscience, London, UK; <http://www.fil.ion.ucl.ac.uk/spm/software/spm8>). The first five volumes of each run were discarded to allow for a steady state of tissue

magnetization. Functional data were then slice-time corrected and spatially realigned to the first volume, coregistered with participants' structural scans, and normalized to the Montreal Neurological Institute (MNI) template brain. Normalized functional images retained their native spatial resolution.

6.2.7 Univariate Analyses

For analyses based on task-related changes in mean signal intensity, the normalized images were spatially smoothed with a Gaussian kernel of 9 mm^3 full width half maximum, before applying a 128 s temporal high-pass filter in order to remove low-frequency noise. A model of the main task was created for each subject via vectors corresponding to the onset of the visual search series (8 s boxcar) for each experimental condition; the model accounted for WM and visual search load conditions, as well as stimulus category for both WM and search task components, resulting in a total of 16 regressors of interest. All univariate analyses collapsed across stimulus category conditions, however, producing four main conditions of interest: low WM/easy search, low WM/hard search, high WM/easy search, high WM/hard search. WM sample and probe periods, error trials (for both visual search and WM probe), head-motion parameters, and grand means of each run were also modeled as separate nuisance regressors. Onset vectors were convolved with a canonical hemodynamic response function (HRF) to produce a design matrix, against which the blood-oxygenation level-dependent signal at each voxel was regressed.

Single-subject contrasts were then calculated to establish the hemodynamic correlates of working memory load (all high WM > all low WM), visual search difficulty (all hard search > all easy search), and their interaction effects (high WM + easy search > low WM + hard search; low WM + hard search > high WM + easy search). Group effects were subsequently assessed by submitting the individual statistical parametric maps to one-sample *t*-tests where subjects were treated as random-effects. To control for false-positives, we applied a whole-brain voxel-wise FDR-correction ($p < .05$, with a cluster extent of 20 voxels). To illustrate the nature of the observed activations, mean beta estimates for each condition were extracted from 6 mm spherical ROIs, centered on peak group activations, using Marsbar software (<http://marsbar.sourceforge.net>).

6.2.8 ROI Definition

Regions of sensitivity for the WM categories were derived from the independent functional localizer task. A model of the localizer was created for each subject via vectors corresponding to the onset of the stimulus block (15 s duration) for each of the four stimulus categories (face, house, body, tool). Single-subject contrasts were then calculated to establish the hemodynamic correlates of face-viewing (all face > all other categories) and house-viewing (all house > all other categories). Group maps were furthermore constrained by anatomical masks of the fusiform and parahippocampal gyri, for the face and house contrasts, respectively, and submitted to FDR-correction ($p < .05$) to identify clusters of maximal responsiveness to the WM categories.

6.2.9 Multivariate Analyses

While standard mass-univariate analyses allow us to localize regions where mean signal intensity is sensitive to internal and external load demands, such variations do not convey precise informational content. Instead, the strength of multivariate decoding can arguably serve as a proxy for the quality of a neural representation (Emrich et al., 2013; Ester et al., 2013). We therefore created a second model for multivariate analyses, using unsmoothed images, with the purpose of gauging how the fidelity of the neural WM representation is impacted when attention is diverted to processing external stimuli. The model included WM category (face vs. house) and visual search load conditions (easy vs. hard) via vectors of onsets (2 s duration) for each event in a trial (i.e., WM sample, WM delay, search trials, pre-probe delay, WM probe). Head-motion parameters and grand means of each run were also modeled as separate nuisance regressors. Searchlight and ROI classification analyses were implemented in The Decoding Toolbox (Hebart, Görgen, & Haynes, 2014), with a L2-norm regularized support vector machine (SVM), using a leave-one-run-out cross-validation procedure. Our design produced 9 experimental runs, wherein each condition occurred equally often. The classifier was thus trained on the patterns corresponding to maintenance of each WM category over 8 runs, then tested on its ability to decipher the remembered category on the 9th run. The training set was then shuffled so that each run served once as the testing set, and classifier accuracy for a given searchlight or ROI would reflect the average classifier performance over those 9 iterations.

6.2.9.1 Searchlight MVPA

We initially attempted to discern whether information about the WM content could be detected (at all) in the face of concurrent sensory input and a secondary task. To that end, we applied whole-brain searchlight multivariate pattern analysis (MVPA) to identify regions conveying locally distributed patterns of WM category information (Haynes et al., 2007; Kriegeskorte, Goebel, & Bandettini, 2006). The searchlight procedure scans through the entire brain, conducting a discrete decoding analysis in small ROIs surrounding every voxel. This produces an accuracy map of the whole brain, for each individual, where the center of each searchlight is ascribed the percentage correct classification achieved within that searchlight. In our case, the searchlight was a spherical cluster with a radius of 3 voxels, thus containing up to 123 voxels. For searchlight analyses, we were interested in classifying whether the remembered category was a face or a house, across the duration of the search sequence. Thus, the inputs to the searchlight were beta estimates for the entire search sequence. The classifier was trained to discriminate the WM category based on multivariate patterns of input from all the voxels in a given searchlight, and that procedure was repeated for searchlights surrounding every gray matter voxel in the brain. The resultant accuracy maps were then submitted to *t*-tests at the group level against chance performance (50%) and thresholded with an FDR-correction of $p < .05$.

6.2.9.2 Time-Resolved ROI-Based MVPA

While searchlight classification decisions are made based on a compact set of voxels, we wanted to assess the quality of WM representations across more widely distributed regions that are engaged for the perception of the WM categories. Specifically—based on prior observations that attentional demands can interfere with WM maintenance—our goal was to examine how the neural activity patterns in regions that represent the WM content would be impacted by the difficulty of a secondary task. Given prior evidence that activity patterns coding WM information vary over the course of a WM delay (Lewis-Peacock, Drysdale, & Postle, 2015; Sreenivasan, Vytlačil, & D'Esposito, 2014b) and that attention demands and time-related decay might interact in their impact on WM maintenance (Kiyonaga & Egner, 2014), we also examined how this impact might accumulate or evolve over the course of a trial as attention continued to be otherwise occupied. Thus, we conducted ROI-based MVPA, wherein classifiers were trained (and tested) using beta values from all voxels in functionally defined FFA and PPA, producing a single accuracy value for the classification based on all the information in that ROI. We conducted separate training and testing of the WM category for each task event (i.e., at 2 s intervals across a trial), at both levels of visual attentional load. We therefore obtained, for each participant, two WM category classification accuracies (one each for easy and hard visual search conditions) at each of 8 trial time points. To assess any potential differences in classification between easy and hard external attention

conditions, these accuracy values were submitted to one-sample *t*-tests against chance (50%).

6.3 Results

6.3.1 Behavioral Results

Visual search accuracy (% correct) was reduced when the search distractors were made less discriminable from the target (easy: $M = 95.0\%$, $SD = 9.1\%$; hard: $M = 84.7\%$, $SD = 12.8\%$), $F(1, 27) = 27.4$, $p < .001$, and search correct response time (RT) was also drastically slower in this high attentional load condition (easy: $M = 848$ ms, $SD = 105$ ms; hard: $M = 1,133$ ms, $SD = 110$ ms), $F(1, 27) = 272.2$, $p < .001$. The external attentional load manipulation was, therefore, an effective means of modulating the time-consumption of WM delay-spanning processing: When the search was harder, it took much more time to complete, suggesting that attention would be diverted from WM maintenance during that time (Figure 19a). Visual search was slightly faster when 2 items were maintained in WM, $F(1, 27) = 4.4$, $ps < .05$, while accuracy was unaffected by internal (WM) load ($p = .6$), and neither accuracy nor RT displayed an interaction between internal and external attention factors (all $ps > .1$). WM probe performance was likewise slower (low: $M = 1,458$ ms, $SD = 287$ ms; high: $M = 1,633$ ms, $SD = 248$ ms), $F(1, 27) = 79.1$, $p < .001$, and less accurate (low: $M = 87.1\%$, $SD = 13.3\%$; high: $M = 75.7\%$, $SD = 14.5\%$), $F(1, 27) = 80.1$, $p < .001$, when 2 items were remembered (vs. 1). Thus, the WM manipulation was effective at increasing internal attentional demands (Figure 19b). WM probe responding was somewhat slower after easier visual search

sequences, $F(1, 27) = 10.5, p < .01$, but neither probe accuracy nor RT displayed an interaction between WM and search load factors (all $ps > .3$).

The absence of an interaction between these load factors was surprising because of the abundant prior evidence that WM performance can be impaired by concurrent attention demands. Thus, we investigated whether the reciprocity between internal and external attention demands may have manifested in another way in the present data set—for instance, by a selective prioritization of one task demand over the other. Therefore, we examined correlations between visual search and WM “load effects” (i.e., high/hard load – low/easy load), and found that the two were negatively correlated with one another, $r = -.41, p < .05$, such that a larger effect of visual search load on search RT was associated with a smaller effect of WM load on probe accuracy (Figure 19c). Thus, those participants who were slowed more dramatically by the harder search condition were less likely to suffer under the higher WM load. A higher visual search load effect was also strongly associated with better WM overall ($r = .71, p < .001$). In other words, participants who dealt with one of the task demands well incurred a higher cost from the other task demand. The lack of a behavioral interaction effect may therefore have stemmed from the fact that participants prioritized one task component (e.g., search) over the other (WM), suggesting that they found it difficult to satisfy both task demands simultaneously. We next examined how these simultaneous demands are reflected in neural measures.

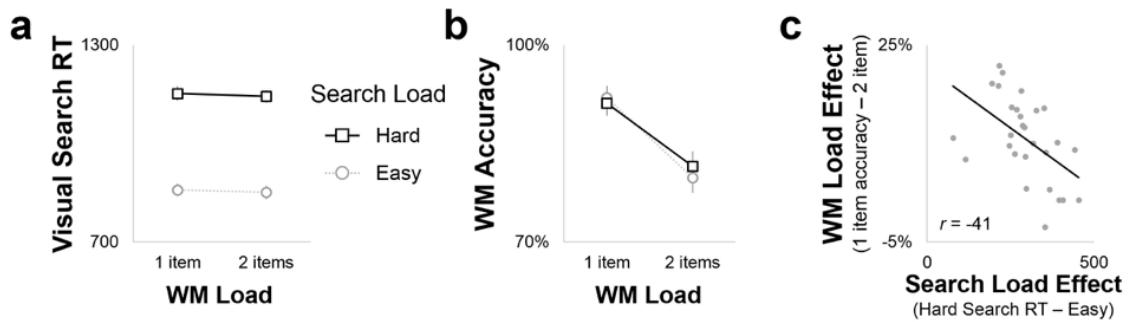


Figure 19: Behavioral results. (A) Mean visual search response time (RT) was faster for easy searches compared to hard searches. (B) WM probe recognition accuracy was higher for low load (1 item) compared to high load (2 item) WM. (C) Load effect scores (i.e., the impact of higher demand on search RT or WM accuracy) on the visual search and WM portions of the dual-task paradigm were significantly correlated.

6.3.2 Univariate fMRI Results

We initially conducted mean signal intensity-based analyses to localize areas that respond to attentional task load, and potentially represent sources of attentional control over the competition between internal and external task demands. First we examined regions that displayed a main effect of the external (visual search) attentional load, during the search sequence (Figure 20a). These encompassed a large bilateral network of frontal, parietal, and occipital regions that are considered part of the “cognitive control network” and are typically engaged when task demands are high (Niendam et al., 2012). While we observed no main effect of WM load (during completion of the search task), there was a robust interaction between visual search and WM demands (Figure 20b). Lateral prefrontal, parietal, and posterior temporal clusters (around the temporo-parietal junction [TPJ]) were sensitive to the combination of load in both the internal and external

domains, suggesting that these regions may be involved in managing the competition between internal and external demands. As illustrated with beta values extracted from an ROI centered on the peak of the interaction (in left middle frontal gyrus), the response to load in one domain (i.e., WM) was dependent on the current load in the other domain (i.e., visual attention). Specifically, the magnitude of the response to a hard search task was dramatically magnified when WM load was high as well (Figure 20b). This neural interaction is thus consistent with the idea that WM and visual attention demands compete with and impact one another. Taken together, both the behavioral data and univariate fMRI results suggest that maintaining visual stimuli in WM was rendered more demanding by concurrent performance of a difficult versus easy search task. These data therefore set the stage for the main goal of our study—that is, the assessment of the fate of WM representations in the face of competition for attention by external stimuli.



Figure 20: Univariate fMRI results. Rendered data are whole-brain FDR-corrected, $p < .05$. (A) A large bilateral network of frontal, parietal, and occipital regions showed elevated activation for hard versus easy visual searches. (B) Lateral prefrontal, parietal, and posterior temporal clusters were sensitive to the interaction between WM and search load levels. Representative beta values are displayed for an ROI surrounding the peak of the interaction contrast in the middle frontal gyrus (MFG).

6.3.3 Multivariate fMRI Results

We initially investigated whether WM category (face vs. house) could be classified during a secondary task in the first place, and then went on to assess variations in this classification as a function of visual search difficulty. We first applied a searchlight procedure to identify regions conveying local patterns of WM category information, across the duration of the search sequence. Indeed, large clusters of searchlights covering the ventral visual stream, and including the fusiform and parahippocampal gyri, classified the WM category significantly above chance (Figure 21a). Thus, even in the face of persistent visual input—not to mention a secondary task being performed on that input—we were able to extract meaningful (i.e., category-diagnostic) WM stimulus information from locally distributed patterns of neural activity

in visual cortex. Moreover, those patterns of decodable activity tended to cluster around the ventral visual regions that typically respond to perception of stimuli from the WM categories.

Having established that WM stimulus information can indeed be extracted in the face of a concurrent attention-demanding task, we next asked (1) how that WM representational content evolves throughout the course of the trial in functionally-defined face- and house-sensitive ROIs, and (2) most importantly, how it is impacted by the time-consumption of an intervening search task. WM category classification in the PPA, overall, displayed a U-shaped pattern across the trial (Figure 21b). Unsurprisingly, regardless of the search difficulty of the current trial, classification of the WM category was highly accurate (88%) during presentation of the WM sample (i.e., when the stimulus was actually being perceived). Classification accuracy dropped off somewhat, but remained well above chance, during the blank delay after the offset of the sample (easy search: $t(27) = 4.4, p < .001$; hard search: $t(27) = 5.2, p < .001$). With progression through the visual search series, however, classification performance began to diverge depending on the attentional demands of that search sequence. When the search was easier, WM category classification remained consistently above chance (S1: $t(27) = 2.1, p = .02$; S2: $t(27) = 1.4, p = .08$; S3: $t(27) = 1.9, p = .04$; S4: $t(27) = 2.4, p = .01$). When the search was more difficult (and therefore diverting attention for a longer period of time), however, WM category classification dropped down to chance levels by the second visual search (S2: $t(27) = 0.0, p = 1.0$), and remained at chance throughout the rest of the

search series (S3: $t(27) = .12, p = .9$; S4: $t(27) = .42, p = .7$). That is, by the second half of the WM delay, WM category information could still be detected when external visual attention demands were low, while classification of that information was no better than chance when those demands were high (Figure 21b). Once the search series ended, classification under the harder search condition recovered to the same level as the easy condition, and both converged to well above chance (75%) when the probe appeared. Here, we thus have evidence that the neural patterns of activity that convey information about the WM content are impeded—or possibly recoded in a different format (Olivers et al., 2011; Stokes, 2015)—when demands on visual attention are high.

Multivariate classification of WM category in the FFA followed the same U-shaped pattern, but the difference between visual search difficulty conditions did not reach significance. The functionally-defined FFA was substantially smaller than the PPA, and indeed, the few prior studies of WM decoding that have used face and house stimuli have found the PPA to provide more reliable classification, in general (Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014; Lewis-Peacock & Norman, 2014).

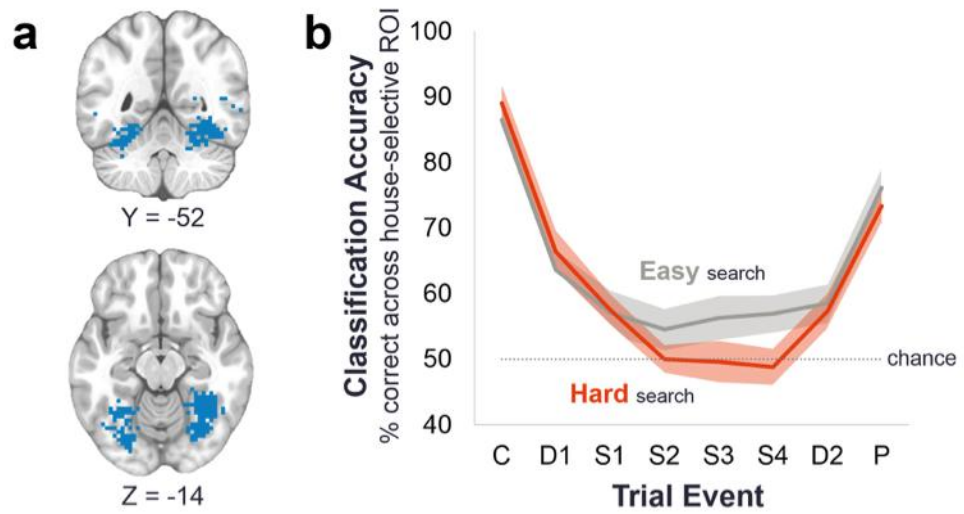


Figure 21: Multivariate fMRI results. (A) Clusters of searchlights across the ventral visual stream classified WM category significantly above chance during the visual search sequence. (B) In functionally-defined parahippocampal place area (PPA), at 2 s timepoints across the trial, WM category classification displayed a U-shaped pattern, suggesting an allocation of attention away from WM during the concurrent search task. WM decoding accuracy remained above chance across easy (gray) visual searches, but dropped below chance later on in the search sequence for hard (orange) visual searches. Ribbons represent ± 1 SEM. C = WM cue; D = delay; S = search; P = probe.

6.4 Discussion

Here, we tested the hypothesis that demands on visual attention should impact neural representations of visual WM content, based on the idea that WM maintenance is accomplished via attention-dependent recruitment of sensory cortices. We manipulated levels of both WM and visual search load in a dual-task paradigm, and found converging behavioral and neuroimaging evidence that these “internal” and “external” attentional demands impact one another. For one, efficient behavioral performance on one part of the task (i.e., WM recognition) was associated with worse performance on the other (i.e.,

visual search), suggesting there is a tradeoff between the two. Secondly, in regions of the fronto-parietal attentional control network, the univariate fMRI response to load in one domain was highly dependent on the load in the other domain, suggesting that the two loads impose demands on a common neural resource. Finally, and most relevant to our goals, the distributed patterns of stimulus information that are thought to convey WM content were degraded under higher visual attentional demands, suggesting that the amount of available attention during the WM maintenance period determines the quality of the neural WM representation.

A searchlight classifier was able to decode the WM category during a delay-spanning series of visual searches, in the local patterns conveyed by clusters of voxels around the ventral visual regions that typically respond to perception of the WM categories (i.e., fusiform and parahippocampal gyri). This result supports the notion that WM maintenance is achieved through activation of sensory representations and marks a substantial advance in the limits of WM decoding. That is, prior studies have successfully decoded WM content from visual cortices over an unfilled delay interval (i.e., no other perceptual input; e.g., Emrich et al., 2013; Ester et al., 2013; Harrison & Tong, 2009; Riggall & Postle, 2012; Serences et al., 2009), and from superior parietal cortex over an interval that included task-irrelevant perceptual input (Bettencourt & Xu, 2015), leaving open the question of what happens to (sensory) WM representations in the face of task-relevant incoming sensory signals (as well as under varying attentional demand). The current data now demonstrate that when WM classifiers are also trained on visual cortex

activity during the completion of a secondary visual task, categorical WM information can be decoded.

Time-resolved pattern classification within the PPA, moreover, illustrated the sensitivity of neural WM representations to concurrent visual attentional demand levels. During a blank delay interval, in the absence of the WM stimulus (both before and after the search sequence), WM category decoding remained well-above chance. Regardless of the difficulty condition of the visual search, WM decoding accuracy declined once the visual task began. As search stimuli were intentionally pulled from different categories than the WM samples, so as to recruit distinct neural representational substrates, this reduction in WM discriminability was likely a product of attention being diverted toward the search task rather than any direct overlap in sensory substrates. Most strikingly, this reduction in decoding accuracy during the search sequence was more profound when the search was more time-consuming, even though the amount of perceptual input was matched in the easy and hard search conditions. Furthermore, this degradation of the cortical WM representation became more severe as attention was diverted from WM maintenance for longer periods of time, as the gap in WM category decoding accuracy between the easy and hard search conditions widened with progression across the search sequence (Figure 21b). The more that attention was required to complete the delay-spanning visual search, the more the discriminability of neural WM representations in ventral visual cortex suffered.

Two intriguing additional findings bear further consideration. For one, while visual attentional demands impacted the neural representation of WM content, they did not lead to deterioration in overall behavioral WM recognition performance. Likewise, once the visual search portion of the trial was completed, neural decoding of the WM category after hard search sequences recovered to levels equal to that after easy sequences, indicating that WM category information was temporarily suppressed, but not eliminated, by the need to direct attention elsewhere. Thus, in the higher demand condition, WM content must be coded or maintained using some other strategy that does not rely on the same patterns of visual representation that are conveyed when attention demands are low. These results are consistent with earlier findings that WM representations can be transferred into a different (and possibly suppressed) activation status to prioritize the immediately relevant task, and then re-prioritized again when they are needed to guide behavior (LaRocque et al., 2013; Lewis-Peacock & Postle, 2012). More generally, this idea of flexible, task-dependent states of WM fits with influential models of multiple-state WM, as established by the allocation of attention (LaRocque et al., 2014; Oberauer, 2002; Olivers et al., 2011). Indeed, interest has grown recently in characterizing a hidden or “silent” WM coding scheme (Stokes, 2015)—possibly implemented via patterned short-term changes in network synaptic weights (Erickson, Maramba, & Lisman, 2010; LaRocque et al., 2014; Stokes, 2015; Stokes et al., 2013; Wolff, Ding, & Myers, 2015)—but it is unclear why and when this maintenance strategy would be used. Our findings raise the intriguing possibility that such a representational

format—that is undetectable with the fMRI methods used here—might be relied upon specifically when WM must be maintained in the absence of sustained attention toward the WM content.

A critical function of WM is to maintain information in the face of competing demands, yet surprisingly little is known about how such attentional demands interact with WM maintenance. The present findings suggest that attention is necessary to maintain visual WM representations in sensory areas, but those distributed activity patterns are not necessarily the sole functional substrate of WM maintenance (Bettencourt & Xu, 2015), since the material can still be remembered when WM decoding falls to chance. Thus, although the quality of multivariate evidence for a WM item is often taken to reflect the precision of the neural WM representation, *per se* (Emrich et al., 2013; Ester et al., 2013), we may need to consider alternate maintenance formats to fully understand how we are best able to juggle our internal goals with persistent concurrent demands for our attention.

7 General Discussion

Although traditionally conceived as separate constructs, the cognitive processes of WM and attention are in fact intertwined and overlapping (Chapter 1). The studies presented in this dissertation further emphasize the reciprocal and competitive interactions between WM and attention, specifically in the visual domain. Chapters 2 through 5 explore this relationship through the phenomenon of memory-based attentional guidance, whereby the active maintenance of items in visual WM reliably biases the deployment of attention to memory-matching items in the visual environment. Specifically, Chapters 2 and 3 present a series of experiments that support a distinction between the (unintentional) effects of WM guidance and the traditional dichotomy of top-down versus bottom-up attentional control. Chapters 3 and 4 also examine the question of under what circumstances the contents of WM impact visual attention, and how the limits of memory-based attentional guidance support models of multiple-state WM. Finally, Chapters 5 and 6 tackle the broader theory of a tight—and even obligatory—link between WM and attention: first by demonstrating the predictive power of WM guidance effects, and second by exploring the competitive tradeoffs in behavioral and neural signals between the two domains. This concluding chapter reviews and synthesizes the present empirical findings and their implications for understanding the nature of WM and attention, while also considering the limitations of current models and open questions for the future.

7.1 Working Memory as a Distinct Influence on Attention, neither “Top-Down” nor “Bottom-Up”

The traditional dichotomy of top-down (i.e., driven by strategic goals) and bottom-up (i.e., driven by physical stimulus attributes) influences on visual attention has been criticized for excluding other impactful factors, such as memory (Awh et al., 2012; Hutchinson & Turk-Browne, 2012). The data from Chapters 2 and 3 support WM as an influential factor on attentional deployment that is distinct from either top-down or bottom-up labeling. In Chapter 2, the effects of WM guidance were found to modulate the effects of bottom-up physical contrast, even so far as to override a baseline salience effect. While a strategic explanation of WM guidance (i.e., explicitly searching for the memory cue) could not be definitively ruled out in that particular paradigm, other studies have demonstrated attentional capture by memory-matching items even when the memory cue was always a distractor (e.g., Olivers et al., 2006; Soto et al., 2005). Conversely, in Chapter 3, top-down control was found to modulate the effects of WM guidance, such that prior knowledge about whether memory items would be reliably predictive (or anti-predictive) of the search target location allowed for strategic enhancement or suppression of WM biasing. Together, these interactions indicate that memory-based attentional guidance is clearly distinct from effects of either bottom-up salience or top-down strategy.

Previous descriptions of memory-based attentional guidance have characterized influences of WM on attention as “top-down” (Soto et al., 2005; 2008), due to the likely

role of WM in maintaining target templates during visual search (Chelazzi et al., 1993; Woodman & Arita, 2011). However, the intentional impact of top-down templates in WM seems intuitively different from the inadvertent biasing from task-irrelevant—and even task-detrimental—items in WM¹ (see also Olivers & Eimer, 2011). The critical distinction is that the active maintenance of items in WM can be separable from, and can interfere with, immediate behavioral goals (e.g., an intervening visual search), although the degree of interference is subject to individual variability (Chapter 4). However, a recent EEG study found that maintaining an item in WM for future recognition and maintaining a target template in WM for visual search actually elicit the same CDA component, an index of WM maintenance (Gunseli, Meeter, & Olivers, 2014a). Thus, an open question is how this distinction between multiple items in WM (i.e., a memory cue from a WM task and a search template from a visual search task) operates mechanistically. One model proposes functionally different states in WM, such that one (and only one) “active” (i.e., more prioritized) WM representation can bias attention (Olivers et al., 2011; van Moorselaar et al., 2014b; but see Beck et al., 2012). Multiple studies have supported this hypothesis, showing that search target templates become

¹ One could argue that because memory items are task-relevant for the purpose of future recognition, they too are “top-down” representations, such that “top-down” labeling simply reflects which goals are being prioritized. However, within the framework of a dual-task paradigm, attentional guidance effects are measured during the visual search task, such that the most immediate goal is search and the “top-down” representation is the target template.

rapidly deprioritized across repetitions (Carlisle et al., 2011; Gunseli et al., 2015), and that while equally-relevant items in WM do *not* interfere with subsequent visual search, simply cueing one WM item to be prioritized elicits search interference effects again (van Moorselaar et al., 2014b; van Moorselaar, Battistoni, Theeuwes, & Olivers, 2014a). In other words, the distinction between an ordinary item in WM and a top-down target template seems to be a function of a rapid and flexible prioritization system within WM, as opposed to qualitatively different WM functions. Another interpretation, then, is that memory-based attentional guidance—inadvertent or intentional—operates via a single shared mechanism, but the “top-down” and volitional nature of a target template alters the balance of prioritization within WM. Thus, while non-volitional WM guidance effects can be separable from and interact with top-down strategy (e.g., Chapter 3), this distinction may be more driven by differences in state (i.e., prioritization or activation) than mechanism (but see Olivers & Eimer, 2011).

A recent modification to the traditional dichotomy of attentional control includes a third category of “selection history” (Awh et al., 2012; Belopolsky, 2015). Selection history refers to the prior processing of a stimulus, classically demonstrated by intertrial priming effects, in which identification of a repeated target becomes more efficient (e.g., Maljkovic & Nakayama, 1994; see Kristjánsson & Campana, 2010). Intertrial priming effects have been described as automatic, implicit memory biases, as the benefits of repetition extend as far as eight trials and cannot be overcome by explicit strategy (e.g., Maljkovic & Nakayama, 1994; see also Maljkovic & Nakayama, 2000). Two accounts of

intertrial priming have been proposed: short-term activation of previously-selected features (Kristjánsson & Campana, 2010), or longer-term retrieval of previous trials as episodic memories (Kruijne & Meeter, 2015). Attentional guidance by WM could thus be reconciled with both accounts of intertrial priming², either as the activation of features that are selected for WM encoding, or as the influence of episodic memory retrieval during WM refreshing. However, one study found separable effects of (but no interactions between) active WM maintenance and selection history (i.e., probability cueing) on attentional capture (Schwark, Dolgov, Sandry, & Volkman, 2013), although that manipulation of selection history did not directly address intertrial priming, per se. Furthermore, as discussed previously, the impact of WM on attention depends on its activation status, such that a new memory cue is more likely to be prioritized *immediately* and thus guide attention (e.g., Gunseli et al., 2015)—in contrast to intertrial priming, where attentional biasing accumulates over multiple repetitions (see Kristjánsson & Campana, 2010). Thus, the question of whether WM guidance effects can be attributed to intertrial priming remains open to future investigation.

² Note that intertrial priming (i.e., repeated attentional processing) is distinct from visual priming (i.e., mere presentation). Chapter 2 presented a control experiment in which attentional guidance effects were *not* present when participants merely viewed a matching cue before visual search, indicating that the processing and encoding of a WM cue for a future test is distinct from passive viewing. For those same reasons, intertrial priming has also been found to be distinct from visual priming (Kristjánsson, Saevarsson, & Driver, 2013).

7.2 Flexible, Multiple-State Models of Working Memory and Attention

The main analysis presented in Chapter 3 addresses the question of how strategic control of memory-based attentional guidance affects the WM representation itself, either in terms of representational quality or accessibility of WM content at retrieval. Using probabilistic model fitting and comparison techniques, we found that strategic control over WM biasing impacted the probability of remembering items in WM—suggesting that whether a WM representation influences attention depends on its accessibility. As described in Chapter 3, this finding supports models of WM that propose multiple, functionally-different states that are characterized by how they interact with attention (Oberauer, 2002; Olivers et al., 2011). Generally, these different representational states reflect two major divisions within WM: (1) a highly-active, capacity-limited “focus of attention” that is immediately accessible for (and can interfere with) current behavior, and (2) a less accessible but still prioritized, larger-capacity buffer of potentially-relevant information (which can in turn be quickly called into a more prioritized state); all of these WM representations are themselves thought to be activated forms of long-term memory representations (see D’Esposito & Postle, 2015). However, the distinction between different functional states may be more conceptual than mechanistic. As addressed previously, it has been shown that WM representations can fluctuate rapidly and flexibly between more and less activated states, in terms of how they interface with visual attention (e.g., Gunseli et al., 2015; van Moorselaar et al., 2014b), and even in terms of

their multivariate neural trace (Lewis-Peacock et al., 2012). Although Chapter 3 speculates that this flexible prioritization may be enacted via changes to WM accessibility, what that represents mechanistically is still unclear. In other words, *how* are WM representations flexibly prioritized and de-prioritized?

One possibility is to think of WM priority as a function of accumulated “resources,” such that activation status reflects the amount of cognitive resources dedicated to that item. While the term “resources” is itself vague and biologically ambiguous (see below), the concept of limited cognitive resources in both attention and memory is common (see Franconeri, Alvarez, & Cavanagh, 2013). For example, one influential conceptualization of visual WM capacity posits a limited and continuous resource that is distributed flexibly among items in WM (Ma et al., 2014). Thus, greater allocation of memory resources may result in higher WM priority. While greater allocation of memory resources has previously been linked to greater representational quality (e.g., Bays et al., 2009; Bays & Husain, 2008), seemingly contradicting the accessibility account put forward by Chapter 3, the two narratives can be reconciled by the “variable-precision” model of WM, in which an item’s encoding precision is itself variable across trials (van den Berg et al., 2012; 2014). Thus, WM activation status may be reflected as accessibility in terms of the variability (i.e., standard deviation) of precision—which is orthogonal to the overall representational quality itself.

Critically, what is the biological basis for these WM “resources”? One hypothesis is that WM resources are essentially limited neural real estate, such that capacity limits

are byproducts of spatial competition between item representations (Franconeri et al., 2013; Liverence & Franconeri, 2015). Similarly, the Neural Theory of Visual Attention (Bundesen, Vangkilde, & Petersen, 2015) interprets resource allocation as the distribution of neural processing via dynamic remapping of receptive fields. An alternative hypothesis is that neural oscillatory rhythms subserve WM, where the degree of oscillatory synchronization in various frequency bands represents different aspects of WM quantity or quality (see D'Esposito & Postle, 2015; Ku, Bodner, & Zhou, 2015). Yet another hypothesis proposes that stimulus features in WM are encoded in the firing rates of a population of neurons, such that WM resources are in fact normalized levels of neural spiking activity (Bays, 2015). This population coding model has been supported by findings that the signal strength in probabilistically spiking neurons is correlated to memory recall errors (Bays, 2014). Furthermore, the population coding model of WM shares many similarities with a normalization model of attention (Reynolds & Heeger, 2009; see also Maunsell, 2015), suggesting a common underlying neural mechanism for resource allocation in both WM and attention.

7.3 Shared Resources and Representations of Working Memory and Attention

Chapters 5 and 6 adopt different approaches to examine the overlap between WM and attention. Using multivariate pattern classification analyses, Chapter 5 showed that the link between WM and visual attention—as indexed by memory-based attentional guidance—is robust and consistent enough to reliably decode the specific contents of

WM from attentional behavior. This link is further supported by the data presented in Chapter 6, in which simultaneous demands on WM and visual search resulted in behavioral tradeoffs across the two task domains. Moreover, neural representations of WM content were modulated by search difficulty, suggesting shared reliance on a common cognitive resource.

The idea of shared cognitive resources between WM and attention has already been acknowledged by current models of WM. For example, in multiple-state models of WM, the distinction between long-term memory storage and activated WM representations is framed as the allocation of “attention” to internal long-term information (D'Esposito & Postle, 2015; Lewis-Peacock & Postle, 2008). Likewise, in flexible-resource models of WM, the unequal distribution of resources is hypothesized to be driven by “attentional” allocation, either externally at encoding or internally during maintenance (see Ma et al., 2014). The long-standing “component processes” view of WM also emphasizes the idea that WM is not its own unique cognitive process, but rather the product of interactions between many processes and many brain regions (see Eriksson, Vogel, Lansner, Bergström, & Nyberg, 2015). Thus, the explicit reconceptualization of WM as internally-directed attention (Chun, 2011; Kiyonaga & Egnér, 2013) is a logical extension of the many parallels and interactions between WM and attention (see Chapter 1).

Central to this reframing of WM is evidence from neuroimaging studies that support shared neural networks for WM and attention. The neural substrates of visual

attention have been identified as a network of frontal and parietal regions that resolve competition in visual cortex (Kastner & Ungerleider, 2000); similarly, evidence from single-unit recordings in non-human primates and EEG and fMRI studies in humans indicate that WM may operate as modulation from these same fronto-parietal areas over more posterior sensory areas (see D'Esposito & Postle, 2015). For instance, one study found that sustained activity in posterior parietal cortex tracked the number of task-relevant stimuli, regardless if the stimuli were currently presented or maintained in WM (Mitchell & Cusack, 2008), suggesting a shared neural component for processing visual load either internally or externally. The fMRI analyses presented in Chapter 6 further support this idea, given the robust interaction effects between *simultaneous* manipulations of internal and external load across frontal and parietal cortical regions, and the modulatory effects of external load on neural representations of WM content in more posterior regions.

Attention and WM have also been argued to share the same neural representations in sensory cortex (i.e., the “sensory recruitment hypothesis”; D'Esposito & Postle, 2015), and multiple studies have used multivariate decoding methods to successfully identify distributed patterns of sensory neural activity that represent both the observation and the maintenance of visual stimuli (e.g., Serences et al., 2009; see Eriksson et al., 2015; Postle, 2015). In one particular fMRI study, participants performed three tasks: maintaining specific spatial locations in WM; covertly attending to specific locations in the periphery; and planning a saccade to a specific location (Jerde, Merriam, Riggall,

Hedges, & Curtis, 2012). Not only were relevant spatial locations successfully discriminated from topographically-specific frontal and parietal activity in all three tasks, but the multivariate classifiers also generalized (i.e., cross-predicted) across tasks, suggesting that the same neural patterns reflect shared representations of spatial location across WM, attention, and intention.

Compelling evidence for shared neural representations between WM and attention has also come from multivariate forward encoding models, which allow for specific hypotheses about underlying neural architectures and how they relate to behavior (see Serences & Saproo, 2012). Ester et al. (2013) identified successful population-level tuning functions that linked fMRI activity in visual cortex to specific orientations of visual stimuli held in WM. Critically, the researchers also found that individual differences in the precision of those tuning functions significantly predicted behavioral memory performance, indicating that neural representations of WM content in sensory cortices are directly linked to behavioral outcomes (Ester et al., 2013). While multivariate decoding analyses have demonstrated generalization of neural classifiers across WM and attention (e.g., Jerde et al., 2012), forward encoding approaches could also be used to directly test how representations of visual stimuli differ during observation versus during WM maintenance.

The sensory recruitment hypothesis is supported by many multivariate neuroimaging analyses that have found reliable discrimination of WM content *only* in sensory cortices (e.g., Albers et al., 2013; Riggall & Postle, 2012; Sreenivasan et al.,

2014b; as well as Chapter 6), questioning the involvement of prefrontal and parietal cortices in representing the *specific contents* of WM and instead ascribing to them the role of top-down task-general modulation (see D'Esposito & Postle, 2015; Sreenivasan et al., 2014a). However, this framework has been challenged by other studies (using a variety of task and analytical designs) that have successfully decoded stimulus-specific WM content from parietal regions (e.g., Bettencourt & Xu, 2015; Christophel et al., 2012; Jerde et al., 2012; Peters et al., 2015). In fact, two recent studies using forward encoding methods were able to reconstruct specific spatial locations and orientations of visual stimuli in WM across occipital, parietal, and frontal regions (Ester et al., 2015; Sprague, Ester, & Serences, 2014), suggesting instead that stimulus-specific WM representations are distributed across sensory and fronto-parietal regions—though the pattern of distribution may depend on more general task-based demands. Thus, if WM is simply internally-directed attention (e.g., Chun, 2011), operating on the same representations and via the same mechanisms, we should expect analogous response patterns in the same cortical networks for attending to either internal or external information. Alternatively, WM could integrate sensory and attentional neural mechanisms into a broader, more distributed program of cognitive representations and functions (e.g., motor planning, decision criteria) that are necessary for flexible, task-relevant behavior (Stokes, 2015).

7.4 Conclusion

In this dissertation, I examine the relationship between WM and attention across multiple studies, primarily focusing on (unintentional) memory-based attentional guidance as a window to conceptual frameworks of attentional control and WM maintenance. The research described here calls into question the utility of traditional frameworks (e.g., the dichotomy of attentional control; Awh et al., 2012) and certain conceptual terms (e.g., Chun et al., 2011) when trying to understand how WM and attention operate. Many open questions remain, such as the dynamics of attentional deployment when guided by WM; how competition for limited processing capacity across internal and external domains affects neural representations; and whether WM is indeed an internal manifestation of attention, or a distinct and broader integration of different cognitive processes. Innovative behavioral paradigms and emerging neuroimaging and modeling techniques will allow for further exploration of these questions, leading to fuller conceptualizations of how we select and maintain information.

Appendix A

Overall Accuracy and Search Times to Find a Target in Single-Target Searches

	Accuracy		Average time in seconds to find target	
	High-salience	Low-salience	High-salience	Low-salience
Experiment 1A (<i>n</i> = 10)				
Neutral	.920 (.08)	.875 (.13)	5.37 (.92)	5.90 (1.53)
High-match	.945 (.06)	--	4.08 (.61)	--
Low-match	--	.915 (.09)	--	4.37 (.82)
Experiment 1B (<i>n</i> = 12)				
Neutral	.833 (.11)	.788 (.14)	6.23 (1.07)	6.92 (1.27)
High-match	.946 (.05)	--	5.32 (1.48)	--
Low-match	--	.842 (.15)	--	5.30 (1.65)
Experiment 2 (<i>n</i> = 10)				
Neutral	.990 (.02)	.750 (.17)	2.47 (.87)	7.11 (1.47)
High-match	1.00 (0)	--	1.79 (.62)	--
Low-match	--	.810 (.17)	--	5.69 (1.04)
Experiment 3 (<i>n</i> = 10)				
Neutral	.785 (.15)	.825 (.12)	5.65 (1.21)	5.85 (1.64)
High-match	.825 (.09)	--	4.98 (.75)	--
Low-match	--	.795 (.12)	--	4.79 (.93)

Note. Standard deviations are presented in parentheses.

References

- Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & de Lange, F. P. (2013). Shared representations for working memory and mental imagery in early visual cortex. *Current Biology*, 23(15), 1427–1431. <http://doi.org/10.1016/j.cub.2013.05.065>
- Anderson, B. A. (2015). The attention habit: How reward learning shapes attentional selection. *Annals of the New York Academy of Sciences*. <http://doi.org/10.1111/nyas.12957>
- Appelbaum, L. G., Cain, M. S., Schroeder, J. E., Darling, E. F., & Mitroff, S. R. (2012). Stroboscopic visual training improves information encoding in short-term memory. *Attention, Perception, & Psychophysics*, 74(8), 1681–1691. <http://doi.org/10.3758/s13414-012-0344-6>
- Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human Perception and Performance*, 38(3), 580–584. <http://doi.org/10.1037/a0027885>
- Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, 5(3), 119–126. [http://doi.org/10.1016/s1364-6613\(00\)01593-x](http://doi.org/10.1016/s1364-6613(00)01593-x)
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437–443. <http://doi.org/10.1016/j.tics.2012.06.010>
- Awh, E., Jonides, J., & Reuter-Lorenz, P. A. (1998). Rehearsal in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 24(3), 780–790. <http://doi.org/10.1037/0096-1523.24.3.780>
- Awh, E., Jonides, J., Smith, E. E., Buxton, R. B., Frank, L. R., Love, T., et al. (1999). Rehearsal in spatial working memory: Evidence from neuroimaging. *Psychological Science*, 10(5), 433–437. <http://doi.org/10.1111/1467-9280.00182>
- Awh, E., Vogel, E. K., & Oh, S. H. (2006). Interactions between attention and working memory. *Neuroscience*, 139(1), 201–208. <http://doi.org/10.1016/j.neuroscience.2005.08.023>
- Baddeley, A. (2003). Working memory: Looking back and looking forward. *Nature Reviews Neuroscience*, 4(10), 829–839. <http://doi.org/10.1038/nrn1201>

- Balani, A. B., Soto, D., & Humphreys, G. W. (2010). Working memory and target-related distractor effects on visual search. *Memory & Cognition*, *38*(8), 1058–1076. <http://doi.org/10.3758/MC.38.8.1058>
- Baron-Cohen, S., Wheelwright, S., Skinner, R., Martin, J., & Clubley, E. (2001). The Autism-Spectrum Quotient (AQ): Evidence from Asperger Syndrome/high-functioning autism, males and females, scientists and mathematicians. *Journal of Autism and Developmental Disorders*, *31*(1), 5–17. <http://doi.org/10.1023/A:1005653411471>
- Barrouillet, P., Portrat, S., & Camos, V. (2011). On the law relating processing to storage in working memory. *Psychological Review*, *118*(2), 175–192. <http://doi.org/10.1037/a0022324>
- Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *Journal of Neuroscience*, *34*(10), 3632–3645. <http://doi.org/10.1523/JNEUROSCI.3204-13.2014>
- Bays, P. M. (2015). Spikes not slots: Noise in neural populations limits working memory. *Trends in Cognitive Sciences*. <http://doi.org/10.1016/j.tics.2015.06.004>
- Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*(5890), 851–854. <http://doi.org/10.1126/science.1158023>
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7–7. <http://doi.org/10.1167/9.10.7>
- Beck, D. M., & Kastner, S. (2009). Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vision Research*, *49*(10), 1154–1165. <http://doi.org/10.1016/j.visres.2008.07.012>
- Beck, V. M., & Hollingworth, A. (2015). Evidence for negative feature guidance in visual search is explained by spatial recoding. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(5), 1190–1196. <http://doi.org/10.1037/xhp0000109>
- Beck, V. M., Hollingworth, A., & Luck, S. J. (2012). Simultaneous control of attention by multiple working memory representations. *Psychological Science*, *23*(8), 887–898. <http://doi.org/10.1177/0956797612439068>

- Belopolsky, A. V. (2015). Common priority map for selection history, reward and emotion in the oculomotor system. *Perception*.
<http://doi.org/10.1177/0301006615596866>
- Belopolsky, A. V., & Theeuwes, J. (2009). No functional role of attention-based rehearsal in maintenance of spatial working memory representations. *Acta Psychologica*, *132*(2), 124–135. <http://doi.org/10.1016/j.actpsy.2009.01.002>
- Belopolsky, A. V., Schreij, D., & Theeuwes, J. (2010). What is top-down about contingent capture? *Attention, Perception, & Psychophysics*, *72*(2), 326–341. <http://doi.org/10.3758/APP.72.2.326>
- Bettencourt, K. C., & Xu, Y. (2015). Decoding the content of visual short-term memory under distraction in occipital and parietal areas. *Nature Neuroscience*, *19*(1), 150–157. <http://doi.org/10.1038/nn.4174>
- Borji, A., & Itti, L. (2014). Defending Yarbus: Eye movements reveal observers' task. *Journal of Vision*, *14*(3), 29–29. <http://doi.org/10.1167/14.3.29>
- Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2013). Visual long-term memory has the same limit on fidelity as visual working memory. *Psychological Science*, *24*(6), 981–990. <http://doi.org/10.1177/0956797612465439>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 443–446.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*(4), 523–547. <http://doi.org/10.1037/0033-295X.97.4.523>
- Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review*, *112*(2), 291–328. <http://doi.org/10.1037/0033-295X.112.2.291>
- Bundesen, C., Vangkilde, S., & Petersen, A. (2015). Recent developments in a computational theory of visual attention (TVA). *Vision Research*, *116*(Part B), 210–218. <http://doi.org/10.1016/j.visres.2014.11.005>
- Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience*, *12*(1), 1–47. <http://doi.org/10.1162/08989290051137585>
- Cain, M. S., & Mitroff, S. R. (2013). Memory for found targets interferes with subsequent performance in multiple-target visual search. *Journal of Experimental*

Psychology: Human Perception and Performance, 39(5), 1398–1408.
<http://doi.org/10.1037/a0030726>

- Cain, M. S., Dunsmoor, J. E., LaBar, K. S., & Mitroff, S. R. (2011). Anticipatory anxiety hinders detection of a second target in dual-target search. *Psychological Science*, 22(7), 866–871. <http://doi.org/10.1177/0956797611412393>
- Carlisle, N. B., & Woodman, G. F. (2011a). Automatic and strategic effects in the guidance of attention by working memory representations. *Acta Psychologica*, 137(2), 217–225. <http://doi.org/10.1016/j.actpsy.2010.06.012>
- Carlisle, N. B., & Woodman, G. F. (2011). When memory is not enough: Electrophysiological evidence for goal-dependent use of working memory representations in guiding visual attention. *Journal of Cognitive Neuroscience*, 23(10), 2650–2664. <http://doi.org/10.1162/jocn.2011.21602>
- Carlisle, N. B., & Woodman, G. F. (2013). Reconciling conflicting electrophysiological findings on the guidance of attention by working memory. *Attention, Perception, & Psychophysics*, 75(7), 1330–1335. <http://doi.org/10.3758/s13414-013-0529-7>
- Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional templates in visual working memory. *Journal of Neuroscience*, 31(25), 9315–9322. <http://doi.org/10.1523/JNEUROSCI.1097-11.2011>
- Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience*, 2(10), 913–919. <http://doi.org/10.1038/13217>
- Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, 363(6427), 345–347. <http://doi.org/10.1038/363345a0>
- Christophel, T. B., Hebart, M. N., & Haynes, J. D. (2012). Decoding the contents of visual short-term memory from human visual and parietal cortex. *Journal of Neuroscience*, 32(38), 12983–12989. <http://doi.org/10.1523/JNEUROSCI.0184-12.2012>
- Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, 49(6), 1407–1409. <http://doi.org/10.1016/j.neuropsychologia.2011.01.029>

- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36(1), 28–71.
<http://doi.org/10.1006/cogp.1998.0681>
- Chun, M. M., & Johnson, M. K. (2011). Memory: Enduring traces of perceptual and reflective attention. *Neuron*, 72(4), 520–535.
<http://doi.org/10.1016/j.neuron.2011.10.026>
- Chun, M. M., & Turk-Browne, N. B. (2008). Associative learning mechanisms in vision. In S. J. Luck & A. Hollingworth (Eds.), *Visual Memory* (pp. 209–245).
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention, *Annual Review of Psychology*, 62(1), 73–101.
<http://doi.org/10.1146/annurev.psych.093008.100427>
- Cook, R. D., & Weisberg, S. (1982). *Residuals and Influence in Regression*. Chapman & Hall/CRC.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 201–215.
<http://doi.org/10.1038/nrn755>
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(01), 87–114.
<http://doi.org/10.1017/S0140525X01003922>
- D'Esposito, M., & Postle, B. R. (2015). The cognitive neuroscience of working memory. *Annual Review of Psychology*, 66(1), 115–142. <http://doi.org/10.1146/annurev-psych-010814-015031>
- Dale, A. M. (1999). Optimal experimental design for event-related fMRI. *Human Brain Mapping*, 8(2-3), 109–114. [http://doi.org/10.1002/\(SICI\)1097-0193\(1999\)8:2/3<109::AID-HBM7>3.0.CO;2-W](http://doi.org/10.1002/(SICI)1097-0193(1999)8:2/3<109::AID-HBM7>3.0.CO;2-W)
- Dalvit, S., & Eimer, M. (2011). Memory-driven attentional capture is modulated by temporal task demands. *Visual Cognition*, 19(2), 145–153.
<http://doi.org/10.1080/13506285.2010.543441>
- de Fockert, J. W., Rees, G., Frith, C. D., & Lavie, N. (2001). The role of working memory in visual selective attention. *Science*, 291(5509), 1803–1806.
<http://doi.org/10.1126/science.1056496>

- Dell'Acqua, R., Sessa, P., Toffanin, P., Luria, R., & Jolicœur, P. (2010). Orienting attention to objects in visual short-term memory. *Neuropsychologia*, *48*(2), 419–428. <http://doi.org/10.1016/j.neuropsychologia.2009.09.033>
- Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, *7*, 1–30.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*(1), 193–222. <http://doi.org/10.1146/annurev.neuro.18.1.193>
- Dimitriadou, E., Hornik, K., Leisch, F., Meyer, D., & Weingessel, A. (2006). *e1071: Misc Functions of the Department of Statistics (e1071)*.
- Dombrowe, I., Olivers, C. N. L., & Donk, M. (2010). The time course of working memory effects on visual attention. *Visual Cognition*, *18*(8), 1089–1112. <http://doi.org/10.1080/13506281003651146>
- Donkin, C., Nosofsky, R., Gold, J., & Shiffrin, R. (2014). Verbal labeling, gradual decay, and sudden death in visual short-term memory. *Psychonomic Bulletin & Review*, *22*(1), 170–178. <http://doi.org/10.3758/s13423-014-0675-5>
- Donohue, S. E., Darling, E. F., & Mitroff, S. R. (2012). Links between multisensory processing and autism. *Experimental Brain Research*, *222*(4), 377–387. <http://doi.org/10.1007/s00221-012-3223-4>
- Dowd, E. W., & Mitroff, S. R. (2013). Attentional guidance by working memory overrides salience cues in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(6), 1786–1796. <http://doi.org/10.1037/a0032548>
- Dowd, E. W., Kiyonaga, A., Beck, J. M., & Egner, T. (2015a). Quality and accessibility of visual working memory during cognitive control of attentional guidance: A Bayesian model comparison approach. *Visual Cognition*, *23*(3), 337–356. <http://doi.org/10.1080/13506285.2014.1003631>
- Dowd, E. W., Kiyonaga, A., Egner, T., & Mitroff, S. R. (2015b). Attentional guidance by working memory differs by paradigm: An individual-differences approach. *Attention, Perception, & Psychophysics*, *77*(3), 704–712. <http://doi.org/10.3758/s13414-015-0847-z>

- Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science, 11*(6), 467–473. <http://doi.org/10.1111/1467-9280.00290>
- Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science, 293*(5539), 2470–2473. <http://doi.org/10.1126/science.1063414>
- Downing, P., & Dodds, C. (2004). Competition in visual working memory for control of search. *Visual Cognition, 11*(6), 689–703. <http://doi.org/10.1080/13506280344000446>
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General, 113*(4), 501–517. <http://doi.org/10.1037/0096-3445.113.4.501>
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96*(3), 433. <http://doi.org/10.1037//0033-295x.96.3.433>
- Ecker, U. K. H., Oberauer, K., & Lewandowsky, S. (2014). Working memory updating involves item-specific removal. *Journal of Memory and Language, 74*(C), 1–15. <http://doi.org/10.1016/j.jml.2014.03.006>
- Eckstein, M. P. (2011). Visual search: A retrospective. *Journal of Vision, 11*(5), 14–14. <http://doi.org/10.1167/11.5.14>
- Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology, 48*, 269–297. <http://doi.org/10.1146/annurev.psych.48.1.269>
- Eimer, M. (2014). The neural basis of attentional control in visual search. *Trends in Cognitive Sciences, 18*(10), 526–535. <http://doi.org/10.1016/j.tics.2014.05.005>
- Emrich, S. M., Riggall, A. C., LaRocque, J. J., & Postle, B. R. (2013). Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. *Journal of Neuroscience, 33*(15), 6516–6523. <http://doi.org/10.1523/JNEUROSCI.5732-12.2013>
- Endl, W., Walla, P., Lindinger, G., & Lalouschek, W. (1998). Early cortical activation indicates preparation for retrieval of memory for faces: an event-related potential study. *Neuroscience, 240*(1), 58–60. [http://doi.org/10.1016/s0304-3940\(97\)00920-8](http://doi.org/10.1016/s0304-3940(97)00920-8)

- Engle, R. W. (2002). Working memory capacity as executive attention. *Current Directions in Psychological Science*, 11(1), 19–23. <http://doi.org/10.1111/1467-8721.00160>
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598–601. <http://doi.org/10.1038/33402>
- Erickson, M. A., Maramba, L. A., & Lisman, J. (2010). A single brief burst induces GluR1-dependent associative short-term potentiation: A potential mechanism for short-term memory. *Journal of Cognitive Neuroscience*, 22(11), 2530–2540. <http://doi.org/10.1162/jocn.2009.21375>
- Eriksen, C. W., & Hoffman, J. E. (1973). The extent of processing of noise elements during selective encoding from visual displays. *Perception & Psychophysics*, 14(1), 155–160. <http://doi.org/10.3758/BF03198630>
- Eriksson, J., Vogel, E. K., Lansner, A., Bergström, F., & Nyberg, L. (2015). Neurocognitive architecture of working memory. *Neuron*, 88(1), 33–46. <http://doi.org/10.1016/j.neuron.2015.09.020>
- 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. http://doi.org/10.1162/jocn_a_00357
- Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron*. <http://doi.org/10.1016/j.neuron.2015.07.013>
- Etzel, J. A., Gazzola, V., & Keysers, C. (2009). An introduction to anatomical ROI-based fMRI classification analysis. *Brain Research*, 1282(C), 114–125. <http://doi.org/10.1016/j.brainres.2009.05.090>
- Fan, J. E., & Turk-Browne, N. B. (2013). Internal attention to features in visual short-term memory guides object learning. *Cognition*, 129(2), 292–308. <http://doi.org/10.1016/j.cognition.2013.06.007>
- Fleck, M. S., Samei, E., & Mitroff, S. R. (2010). Generalized “satisfaction of search”: Adverse influences on dual-target search accuracy. *Journal of Experimental Psychology: Applied*, 16(1), 60–71. <http://doi.org/10.1037/a0018629>
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology*:

Human Perception and Performance, 18(4), 1030–1044.
<http://doi.org/10.1037//0096-1523.18.4.1030>

- Fougnie, D. (2009). The relationship between attention and working memory. In N. B. Johansen (Ed.), *New Research on Short-Term Memory*. New York: Nova Science Publishers.
- Fougnie, D., & Marois, R. (2006). Distinct capacity limits for attention and working memory: Evidence from attentive tracking and visual working memory paradigms. *Psychological Science*, 17(6), 526–534. <http://doi.org/10.1111/j.1467-9280.2006.01739.x>
- Fougnie, D., & Marois, R. (2009). Attentive tracking disrupts feature binding in visual working memory. *Visual Cognition*, 17(1-2), 48–66.
<http://doi.org/10.1080/13506280802281337>
- Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: Competitive content maps for attention and memory. *Trends in Cognitive Sciences*, 17(3), 3–11. <http://doi.org/10.1016/j.tics.2013.01.010>
- Franz, V. H., & Luxburg, von, U. (2015). No evidence for unconscious lie detection: A significant difference does not imply accurate classification. *Psychological Science*, 26(10), 1646–1648. <http://doi.org/10.1177/0956797615597333>
- Fukuda, K., & Vogel, E. K. (2009). Human variation in overriding attentional capture. *Journal of Neuroscience*, 29(27), 8726–8733.
<http://doi.org/10.1523/JNEUROSCI.2145-09.2009>
- Gargaro, B. A., Rinehart, N. J., Bradshaw, J. L., Tonge, B. J., & Sheppard, D. M. (2011). Autism and ADHD: How far have we come in the comorbidity debate? *Neuroscience and Biobehavioral Reviews*, 35(5), 1081–1088.
<http://doi.org/10.1016/j.neubiorev.2010.11.002>
- Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: Bridging selective attention and working memory. *Trends in Cognitive Sciences*, 16(2), 128–134.
<http://doi.org/10.1016/j.tics.2011.11.014>
- Geurts, H. M., Verte, S., Oosterlaan, J., Roeyers, H., & Sergeant, J. A. (2004). How specific are executive functioning deficits in attention deficit hyperactivity disorder and autism? *Journal of Child Psychology and ...*, 45(4), 836–854.
<http://doi.org/10.1111/j.1469-7610.2004.00276.x>

- Gorgoraptis, N., Catalao, R. F. G., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *Journal of Neuroscience*, *31*(23), 8502–8511. <http://doi.org/10.1523/JNEUROSCI.0208-11.2011>
- Greene, C. M., Kennedy, K., & Soto, D. (2015). Dynamic states in working memory modulate guidance of visual attention: Evidence from an n-back paradigm. *Visual Cognition*, *23*(5), 546–560. <http://doi.org/10.1080/13506285.2015.1052033>
- Greene, M. R., Liu, T., & Wolfe, J. M. (2012). Reconsidering Yarbus: A failure to predict observers' task from eye movement patterns. *Vision Research*, *62*(C), 1–8. <http://doi.org/10.1016/j.visres.2012.03.019>
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194. <http://doi.org/10.1162/089892903322598139>
- Gunseli, E., Meeter, M., & Olivers, C. N. L. (2014a). Is a search template an ordinary working memory? Comparing electrophysiological markers of working memory maintenance for visual search and recognition. *Neuropsychologia*, *60*(C), 29–38. <http://doi.org/10.1016/j.neuropsychologia.2014.05.012>
- Gunseli, E., Olivers, C. N. L., & Meeter, M. (2014b). Effects of search difficulty on the selection, maintenance, and learning of attentional templates. *Journal of Cognitive Neuroscience*, *26*(9), 2042–2054. http://doi.org/10.1162/jocn_a_00600
- Gunseli, E., Olivers, C. N. L., & Meeter, M. (2015). Task-irrelevant memories rapidly gain attentional control with learning. *Journal of Experimental Psychology: Human Perception and Performance*. <http://doi.org/10.1037/xhp0000134>
- Han, S. W. (2015). Working memory contents revive the neglected, but suppress the inhibited. *Cognition*, *145*(C), 116–121. <http://doi.org/10.1016/j.cognition.2015.08.012>
- Han, S. W., & Kim, M.-S. (2009). Do the contents of working memory capture attention? Yes, but cognitive control matters. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1292–1302. <http://doi.org/10.1037/a0016452>
- Harel, J., Koch, C., & Perona, P. (2006). Graph-based visual saliency. *Advances in Neural Information Processing Systems*, *19*, 545–552.
- Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–635. <http://doi.org/10.1038/nature07832>

- Hastie, T., & Junyang, Q. (2014). *Glmnet Vignette*.
- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual Review of Neuroscience*, *37*(1), 435–456. <http://doi.org/10.1146/annurev-neuro-062012-170325>
- Haynes, J.-D. (2015). A primer on pattern-based approaches to fMRI: Principles, pitfalls, and perspectives. *Neuron*, *87*(2), 257–270. <http://doi.org/10.1016/j.neuron.2015.05.025>
- Haynes, J.-D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading hidden intentions in the human brain. *Current Biology*, *17*(4), 323–328. <http://doi.org/10.1016/j.cub.2006.11.072>
- Hebart, M. N., Görgen, K., & Haynes, J.-D. (2014). The Decoding Toolbox (TDT): A versatile software package for multivariate analyses of functional imaging data. *Frontiers in Neuroinformatics*, *8*(88), 1–18. <http://doi.org/10.3389/fninf.2014.00088/abstract>
- Hedge, C., Oberauer, K., & Leonards, U. (2015). Selection in spatial working memory is independent of perceptual selective attention, but they interact in a shared spatial priority map. *Attention, Perception, & Psychophysics*, *77*(8), 2653–2668. <http://doi.org/10.3758/s13414-015-0976-4>
- Henderson, J. M., Shinkareva, S. V., Wang, J., Luke, S. G., & Olejarczyk, J. (2013). Predicting cognitive state from eye movements. *PloS ONE*, *8*(5), e64937–6. <http://doi.org/10.1371/journal.pone.0064937>
- Hernández, M., Costa, A., & Humphreys, G. W. (2010). The size of an attentional window affects working memory guidance. *Attention, Perception, & Psychophysics*, *72*(4), 963–972. <http://doi.org/10.3758/APP.72.4.963>
- Hollingworth, A., & Hwang, S. (2013). The relationship between visual working memory and attention: retention of precise colour information in the absence of effects on perceptual selection. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *368*(1628), 20130061–20130061. <http://doi.org/10.1098/rstb.2013.0061>
- Hollingworth, A., Matsukura, M., & Luck, S. J. (2013). Visual working memory modulates rapid eye movements to simple onset targets. *Psychological Science*, *24*(5), 790–796. <http://doi.org/10.1177/0956797612459767>

- Horowitz, T. S., & Wolfe, J. M. (2001). Search for multiple targets: Remember the targets, forget the search. *Perception & Psychophysics*, *63*(2), 272–285. <http://doi.org/10.3758/bf03194468>
- Houtkamp, R., & Roelfsema, P. R. (2006). The effect of items in working memory on the deployment of attention and the eyes during visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(2), 423–442. <http://doi.org/10.1037/0096-1523.32.2.423>
- Hutchinson, J. B., & Turk-Browne, N. B. (2012). Memory-guided attention: Control from multiple memory systems. *Trends in Cognitive Sciences*, *16*(12), 576–579. <http://doi.org/10.1016/j.tics.2012.10.003>
- Irwin, D. E., & Gordon, R. D. (1998). Eye movements, attention and trans-saccadic memory. *Visual Cognition*, *5*(1-2), 127–155. <http://doi.org/10.1080/713756783>
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10-12), 1489–1506. [http://doi.org/10.1016/s0042-6989\(99\)00163-7](http://doi.org/10.1016/s0042-6989(99)00163-7)
- Jasper, L., & Goldberg, I. (1993). *Jasper/Goldberg Adult ADD Questionnaire*.
- 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. <http://doi.org/10.1523/JNEUROSCI.3810-12.2012>
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*(5), 679–685. <http://doi.org/10.1038/nn1444>
- Kanade, T., Cohn, J. F., & Tian, Y. (2000). Comprehensive database for facial expression analysis. *Fourth International Conference on Automatic Face and Gesture Recognition* (pp. 46–53). IEEE. <http://doi.org/10.1109/AFGR.2000.840611>
- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*(11), 4302–4311.
- Kardan, O., Berman, M. G., Yourganov, G., Schmidt, J., & Henderson, J. M. (2015). Classifying mental states from eye movements during scene viewing. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1502–1514. <http://doi.org/10.1037/a0039673>

- Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, *23*, 315–341.
<http://doi.org/10.1146/annurev.neuro.23.1.315>
- Kim, G., Lewis-Peacock, J. A., Norman, K. A., & Turk-Browne, N. B. (2014). Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences*, *111*(24), 8997–9002.
<http://doi.org/10.1073/pnas.1319438111>
- Kiss, M., Van Velzen, J., & Eimer, M. (2008). The N2pc component and its links to attention shifts and spatially selective visual processing. *Psychophysiology*, *45*(2), 240–249. <http://doi.org/10.1111/j.1469-8986.2007.00611.x>
- Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, *20*(2), 228–242. <http://doi.org/10.3758/s13423-012-0359-y>
- Kiyonaga, A., & Egner, T. (2014). Resource-sharing between internal maintenance and external selection modulates attentional capture by working memory content. *Frontiers in Human Neuroscience*, *8*(670), 1–9.
<http://doi.org/10.3389/fnhum.2014.00670/abstract>
- Kiyonaga, A., & Egner, T. (2015). Center-surround inhibition in working memory. *Current Biology*. <http://doi.org/10.1016/j.cub.2015.11.013>
- Kiyonaga, A., Egner, T., & Soto, D. (2012). Cognitive control over working memory biases of selection. *Psychonomic Bulletin & Review*, *19*(4), 639–646.
<http://doi.org/10.3758/s13423-012-0253-7>
- Kiyonaga, A., Korb, F. M., Lucas, J., Soto, D., & Egner, T. (2014). Dissociable causal roles for left and right parietal cortex in controlling attentional biases from the contents of working memory. *NeuroImage*, *100*(C), 200–205.
<http://doi.org/10.1016/j.neuroimage.2014.06.019>
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, *4*, 219–227.
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences*, *103*(10), 3863–3868.
<http://doi.org/10.1073/pnas.0600244103>

- Kristjánsson, Á., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics*, 72(1), 5–18. <http://doi.org/10.3758/APP.72.1.5>
- Kristjánsson, Á., Saevarsson, S., & Driver, J. (2013). The boundary conditions of priming of visual search: From passive viewing through task-relevant working memory load. *Psychonomic Bulletin & Review*, 20(3), 514–521. <http://doi.org/10.3758/s13423-013-0375-6>
- Kruijne, W., & Meeter, M. (2015). The long and the short of priming in visual search. *Attention, Perception, & Psychophysics*, 77(5), 1558–1573. <http://doi.org/10.3758/s13414-015-0860-2>
- Kruschke, J. K. (2010). Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(5), 658–676. <http://doi.org/10.1002/wcs.72>
- Kruschke, J. K. (2011). Bayesian assessment of null values via parameter estimation and model comparison. *Perspectives on Psychological Science*, 6(3), 299–312. <http://doi.org/10.1177/1745691611406925>
- Kruschke, J. K. (2013). Bayesian estimation supersedes the *t* test. *Journal of Experimental Psychology: General*, 142(2), 573–603. <http://doi.org/10.1037/a0029146>
- Ku, Y., Bodner, M., & Zhou, Y.-D. (2015). Prefrontal cortex and sensory cortices during working memory: Quantity and quality. *Neuroscience Bulletin*, 31(2), 175–182. <http://doi.org/10.1007/s12264-014-1503-7>
- Kumar, S., Soto, D., & Humphreys, G. W. (2009). Electrophysiological evidence for attentional guidance by the contents of working memory. *European Journal of Neuroscience*, 30(2), 307–317. <http://doi.org/10.1111/j.1460-9568.2009.06805.x>
- Kuo, B. C., Rao, A., & Lepsien, J. (2009). Searching for targets within the spatial layout of visual short-term memory. *Journal of Neuroscience*, 29(25), 8032–8038. <http://doi.org/10.1523/jneurosci.0952-09.2009>
- Kuo, B.-C., Stokes, M. G., Murray, A. M., & Nobre, A. C. (2014). Attention biases visual activity in visual short-term memory. *Journal of Cognitive Neuroscience*, 26(7), 1377–1389. http://doi.org/10.1162/jocn_a_00577
- LaBar, K. S., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (1999). Neuroanatomic overlap of working memory and spatial attention networks: A functional MRI

- comparison within subjects. *NeuroImage*, 10(6), 695–704.
<http://doi.org/10.1006/nimg.1999.0503>
- Lamy, D. F., & Kristjansson, A. (2013). Is goal-directed attentional guidance just intertrial priming? A review. *Journal of Vision*, 13(3:14), 1–14.
<http://doi.org/10.1167/13.3.14>
- LaRocque, J. J., Lewis-Peacock, J. A., & Drysdale, A. T. (2013). Decoding attended information in short-term memory: An EEG study. *Journal of Cognitive Neuroscience*, 25(1), 127–142. http://doi.org/10.1162/jocn_a_00305
- LaRocque, J. J., Lewis-Peacock, J. A., & Postle, B. R. (2014). Multiple neural states of representation in short-term memory? It's a matter of attention. *Frontiers in Human Neuroscience*, 8. <http://doi.org/10.3389/fnhum.2014.00005>
- Leber, A. B., & Egeth, H. E. (2006). Attention on autopilot: Past experience and attentional set. *Visual Cognition*, 14(4-8), 565–583.
<http://doi.org/10.1080/13506280500193438>
- Leber, A. B., Kawahara, J.-I., & Gabari, Y. (2009). Long-term abstract learning of attentional set. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1385–1397. <http://doi.org/10.1037/a0016470>
- Lewis-Peacock, J. A., & Norman, K. A. (2014). Competition between items in working memory leads to forgetting. *Nature Communications*, 5, 1–10.
<http://doi.org/10.1038/ncomms6768>
- Lewis-Peacock, J. A., & Postle, B. R. (2008). Temporary activation of long-term memory supports working memory. *Journal of Neuroscience*, 28(35), 8765–8771.
<http://doi.org/10.1523/JNEUROSCI.1953-08.2008>
- Lewis-Peacock, J. A., & Postle, B. R. (2012). Decoding the internal focus of attention. *Neuropsychologia*, 50(4), 470–478.
<http://doi.org/10.1016/j.neuropsychologia.2011.11.006>
- Lewis-Peacock, J. A., Drysdale, A. T., & Oberauer, K. (2012). Neural evidence for a distinction between short-term memory and the focus of attention. *Journal of Cognitive Neuroscience*, 24(1), 61–79. http://doi.org/10.1162/jocn_a_00140
- Lewis-Peacock, J. A., Drysdale, A. T., & Postle, B. R. (2015). Neural evidence for the flexible control of mental representations. *Cerebral Cortex*, 25(10), 3303–3313.
<http://doi.org/10.1093/cercor/bhu130>

- Libera, Della, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science, 17*(3), 222–227. <http://doi.org/10.1111/j.1467-9280.2006.01689.x>
- Liverence, B. M., & Franconeri, S. L. (2015). Resource limitations in visual cognition. In R. Scott & S. Kosslyn (Eds.), *Emerging Trends in the Social and Behavioral Sciences* (pp. 1–13).
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390*, 279–281.
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences, 17*(8), 391–400. <http://doi.org/10.1016/j.tics.2013.06.006>
- Lundqvist, D., Flykt, A., & Öhman, A. (1998). The Karolinska directed emotional faces (KDEF).
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience, 17*(3), 347–356. <http://doi.org/10.1038/nn.3655>
- Makovski, T. (2012). Are multiple visual short-term memory storages necessary to explain the retro-cue effect? *Psychonomic Bulletin & Review, 19*(3), 470–476. <http://doi.org/10.3758/s13423-012-0235-9>
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*(2), 369–380. <http://doi.org/10.1037/0278-7393.34.2.369>
- Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition, 22*(6), 657–672. <http://doi.org/10.3758/bf03209251>
- Maljkovic, V., & Nakayama, K. (2000). Priming of popout: III. A short-term implicit memory system beneficial for rapid target selection. *Visual Cognition, 7*(5), 571–595. <http://doi.org/10.1080/135062800407202>
- Manly, T., Anderson, V., & Nimmo-Smith, I. (2001). The differential assessment of children's attention: The Test of Everyday Attention for Children (TEA-Ch), normative sample and ADHD performance. *Journal of Child Psychology and ...*, *42*(8), 1065–1081. <http://doi.org/10.1111/1469-7610.00806>

- Matsukura, M., Luck, S. J., & Vecera, S. P. (2007). Attention effects during visual short-term memory maintenance: Protection or prioritization? *Perception & Psychophysics*, *69*(8), 1422–1434. <http://doi.org/10.3758/BF03192957>
- Maunsell, J. H. R. (2015). Neuronal Mechanisms of Visual Attention. *Annual Review of Vision Science*, *1*(1), 373–391. <http://doi.org/10.1146/annurev-vision-082114-035431>
- McMains, S., & Kastner, S. (2011). Interactions of top-down and bottom-up mechanisms in human visual cortex. *Journal of Neuroscience*, *31*(2), 587–597. <http://doi.org/10.1523/JNEUROSCI.3766-10.2011>
- Minear, M., & Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments, & Computers*, *36*(4), 630–633. <http://doi.org/10.3758/bf03206543>
- Mitchell, D. J., & Cusack, R. (2008). Flexible, capacity-limited activity of posterior parietal cortex in perceptual as well as visual short-term memory tasks. *Cerebral Cortex*, *18*(8), 1788–1798. <http://doi.org/10.1093/cercor/bhm205>
- Moore, E., Laiti, L., & Chelazzi, L. (2003). Associative knowledge controls deployment of visual selective attention. *Nature Neuroscience*, *6*(2), 182–189. <http://doi.org/10.1038/nn996>
- Murray, A. M., Nobre, A. C., Clark, I. A., Cravo, A. M., & Stokes, M. G. (2013). Attention restores discrete items to visual short-term memory. *Psychological Science*, *24*(4), 550–556. <http://doi.org/10.1177/0956797612457782>
- Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin & Review*, *4*(1), 79–95. <http://doi.org/10.3758/bf03210778>
- Niendam, T. A., Laird, A. R., Ray, K. L., Dean, Y. M., Glahn, D. C., & Carter, C. S. (2012). Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. *Cognitive, Affective, & Behavioral Neuroscience*, *12*(2), 241–268. <http://doi.org/10.3758/s13415-011-0083-5>
- Nothdurft, H. C. (1993). The role of features in preattentive vision: Comparison of orientation, motion and color cues. *Vision Research*, *33*(14), 1937–1958. [http://doi.org/10.1016/0042-6989\(93\)90020-w](http://doi.org/10.1016/0042-6989(93)90020-w)
- Nothdurft, H. C. (2002). Attention shifts to salient targets. *Vision Research*, *42*(10), 1287–1306. [http://doi.org/10.1016/s0042-6989\(02\)00016-0](http://doi.org/10.1016/s0042-6989(02)00016-0)

- Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(3), 411–421. <http://doi.org/10.1037//0278-7393.28.3.411>
- Oberfeld, D., & Franke, T. (2012). Evaluating the robustness of repeated measures analyses: The case of small sample sizes and nonnormal data. *Behavior Research Methods*, 45(3), 792–812. <http://doi.org/10.3758/s13428-012-0281-2>
- Oh, S. H., & Kim, M. S. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic Bulletin & Review*, 11(2), 275–281. <http://doi.org/10.3758/bf03196570>
- Olivers, C. N. L. (2009). What drives memory-driven attentional capture? The effects of memory type, display type, and search type. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1275–1291. <http://doi.org/10.1037/a0013896>
- Olivers, C. N. L., & Eimer, M. (2011). On the difference between working memory and attentional set. *Neuropsychologia*, 49(6), 1553–1558. <http://doi.org/10.1016/j.neuropsychologia.2010.11.033>
- Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, 32(5), 1243–1265. <http://doi.org/10.1037/0096-1523.32.5.1243>
- Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, 15(7), 327–334. <http://doi.org/10.1016/j.tics.2011.05.004>
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, 105(32), 11087–11092. <http://doi.org/10.1073/pnas.0805664105>
- Pan, Y., & Soto, D. (2010). The modulation of perceptual selection by working memory is dependent on the focus of spatial attention. *Vision Research*, 50(15), 1437–1444. <http://doi.org/10.1016/j.visres.2009.10.016>
- Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, 42(1), 107–123. [http://doi.org/10.1016/s0042-6989\(01\)00250-4](http://doi.org/10.1016/s0042-6989(01)00250-4)

- Pennington, B. F., & Ozonoff, S. (1996). Executive functions and developmental psychopathology. *Journal of Child Psychology and ...*, *37*(1), 51–87. <http://doi.org/10.1111/j.1469-7610.1996.tb01380.x>
- Pereira, F., Mitchell, T., & Botvinick, M. (2009). Machine learning classifiers and fMRI: A tutorial overview. *NeuroImage*, *45*(S1), S199–S209. <http://doi.org/10.1016/j.neuroimage.2008.11.007>
- Pertsov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(5), 1224–1231. <http://doi.org/10.1037/a0030947>
- Peters, B., Kaiser, J., Rahm, B., & Bledowski, C. (2015). Activity in human visual and parietal cortex reveals object-based attention in working memory. *Journal of Neuroscience*, *35*(8), 3360–3369. <http://doi.org/10.1523/JNEUROSCI.3795-14.2015>
- Peters, J. C., Goebel, R., & Roelfsema, P. R. (2009). Remembered but unused: The accessory items in working memory that do not guide attention. *Journal of Cognitive Neuroscience*, *21*(6), 1081–1091. <http://doi.org/10.1162/jocn.2009.21083>
- Postle, B. R. (2015). The cognitive neuroscience of visual short-term memory. *Current Opinion in Behavioral Sciences*, *1*, 40–46. <http://doi.org/10.1016/j.cobeha.2014.08.004>
- Postle, B. R., Awh, E., Jonides, J., Smith, E. E., & D'Esposito, M. (2004). The where and how of attention-based rehearsal in spatial working memory. *Cognitive Brain Research*, *20*(2), 194–205. <http://doi.org/10.1016/j.cogbrainres.2004.02.008>
- Rerko, L., & Oberauer, K. (2013). Focused, unfocused, and defocused information in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(4), 1075–1096. <http://doi.org/10.1037/a0031172>
- Reynolds, J. H., & Desimone, R. (2003). Interacting roles of attention and visual salience in V4. *Neuron*, *37*(5), 853–863. [http://doi.org/10.1016/S0896-6273\(03\)00097-7](http://doi.org/10.1016/S0896-6273(03)00097-7)
- Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*(2), 168–185. <http://doi.org/10.1016/j.neuron.2009.01.002>
- Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging.

Journal of Neuroscience, 32(38), 12990–12998.
<http://doi.org/10.1523/JNEUROSCI.1892-12.2012>

- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies - Revisited. *NeuroImage*, 84(C), 971–985.
<http://doi.org/10.1016/j.neuroimage.2013.08.065>
- Rutman, A. M., Clapp, W. C., & Chadick, J. Z. (2010). Early top-down control of visual processing predicts working memory performance. *Journal of Cognitive Neuroscience*, 22(6), 1224–1234. <http://doi.org/10.1162/jocn.2009.21257>
- Sahan, M. I., Verguts, T., Boehler, C. N., & Pourtois, G. (2015). Paying attention to working memory: Similarities in the spatial distribution of attention in mental and physical space. *Psychonomic Bulletin & Review*. <http://doi.org/10.3758/s13423-015-0990-5>
- Schmidt, B. K., Vogel, E. K., Woodman, G. F., & Luck, S. J. (2002). Voluntary and automatic attentional control of visual working memory. *Perception & Psychophysics*, 64(5), 754–763. <http://doi.org/10.3758/BF03194742>
- Schmidt, F. L., & Hunter, J. E. (1996). Measurement error in psychological research: Lessons from 26 research scenarios. *Psychological Methods*, 1(2), 199–223.
<http://doi.org/10.1037/1082-989X.1.2.199>
- Schwark, J. D., Dolgov, I., Sandry, J., & Volkman, C. B. (2013). Simultaneous attentional guidance by working-memory and selection history reveals two distinct sources of attention. *Acta Psychologica*, 144(2), 269–278.
<http://doi.org/10.1016/j.actpsy.2013.06.017>
- Serences, J. T., & Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia*, 50(4), 435–446.
<http://doi.org/10.1016/j.neuropsychologia.2011.07.013>
- 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.
<http://doi.org/10.1111/j.1467-9280.2009.02276.x>
- Serences, J. T., Shomstein, S., Leber, A. B., Golay, X., Egeth, H. E., & Yantis, S. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychological Science*, 16(2), 114–122. <http://doi.org/10.1111/j.0956-7976.2005.00791.x>

- Shinkareva, S. V., Mason, R. A., Malave, V. L., Wang, W., Mitchell, T. M., & Just, M. A. (2008). Using fMRI brain activation to identify cognitive states associated with perception of tools and dwellings. *PLoS ONE*, 3(1), e1394–9. <http://doi.org/10.1371/journal.pone.0001394>
- Soto, D., & Humphreys, G. W. (2008). Stressing the mind: The effect of cognitive load and articulatory suppression on attentional guidance from working memory. *Perception & Psychophysics*, 70(5), 924–934. <http://doi.org/10.3758/PP.70.5.924>
- Soto, D., Greene, C. M., Kiyonaga, A., Rosenthal, C. R., & Egner, T. (2012a). A parieto-medial temporal pathway for the strategic control over working memory biases in human visual attention. *Journal of Neuroscience*, 32(49), 17563–17571. <http://doi.org/10.1523/JNEUROSCI.2647-12.2012>
- Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 31(2), 248–261. <http://doi.org/10.1037/0096-1523.31.2.248>
- Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*, 12(9), 342–348. <http://doi.org/10.1016/j.tics.2008.05.007>
- Soto, D., Humphreys, G. W., & Heinke, D. (2006). Working memory can guide pop-out search. *Vision Research*, 46(6-7), 1010–1018. <http://doi.org/10.1016/j.visres.2005.09.008>
- Soto, D., Humphreys, G. W., & Rotshtein, P. (2007). Dissociating the neural mechanisms of memory-based guidance of visual selection. *Proceedings of the National Academy of Sciences*, 104(43), 17186–17191. <http://doi.org/10.1073/pnas.0703706104>
- Soto, D., Llewelyn, D., & Silvanto, J. (2012b). Distinct causal mechanisms of attentional guidance by working memory and repetition priming in early visual cortex. *Journal of Neuroscience*, 32(10), 3447–3452. <http://doi.org/10.1523/JNEUROSCI.6243-11.2012>
- Soto, D., Rotshtein, P., & Kanai, R. (2014). Parietal structure and function explain human variation in working memory biases of visual attention. *NeuroImage*, 89(C), 289–296. <http://doi.org/10.1016/j.neuroimage.2013.11.036>

- Souza, A. S., & Oberauer, K. (2014). Time-based forgetting in visual working memory reflects temporal distinctiveness, not decay. *Psychonomic Bulletin & Review*, 22(1), 156–162. <http://doi.org/10.3758/s13423-014-0652-z>
- Souza, A. S., Rerko, L., Lin, H.-Y., & Oberauer, K. (2014). Focused attention improves working memory: implications for flexible-resource and discrete-capacity models. *Attention, Perception, & Psychophysics*, 76(7), 2080–2102. <http://doi.org/10.3758/s13414-014-0687-2>
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11), 1–29. <http://doi.org/10.1037/h0093759>
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian Measures of Model Complexity and Fit. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 64(4), 583–639. <http://doi.org/10.2307/3088806?ref=search-gateway:e090def79c3cc2570b10e162d6231ed4>
- Sprague, T. C., Ester, E. F., & Serences, J. T. (2014). Reconstructions of information in visual spatial working memory degrade with memory load. *Current Biology*, 24(18), 2174–2180. <http://doi.org/10.1016/j.cub.2014.07.066>
- Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014a). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, 18(2), 82–89. <http://doi.org/10.1016/j.tics.2013.12.001>
- Sreenivasan, K. K., Vytlačil, J., & D'Esposito, M. (2014b). Distributed and dynamic storage of working memory stimulus information in extrastriate cortex. *Journal of Cognitive Neuroscience*, 26(5), 1141–1153. http://doi.org/10.1162/jocn_a_00556
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, 46(4), 1004–1017. <http://doi.org/10.1016/j.neuroimage.2009.03.025>
- Stokes, M. G. (2015). “Activity-silent” working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, 19(7), 394–405. <http://doi.org/10.1016/j.tics.2015.05.004>
- Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, 78(2), 364–375. <http://doi.org/10.1016/j.neuron.2013.01.039>

- Suchow, J. W., Brady, T. F., Fournie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision, 13*(10), 9–9. <http://doi.org/10.1167/13.10.9>
- Summerfield, J. J., Lepsien, J., Gitelman, D. R., Mesulam, M. M., & Nobre, A. C. (2006). Orienting attention based on long-term memory experience. *Neuron, 49*(6), 905–916. <http://doi.org/10.1016/j.neuron.2006.01.021>
- Tan, J., Zhao, Y., Wang, L., Tian, X., Cui, Y., Yang, Q., et al. (2015). The competitive influences of perceptual load and working memory guidance on selective attention. *PloS ONE, 10*(6), e0129533–14. <http://doi.org/10.1371/journal.pone.0129533>
- Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics, 51*(6), 599–606. <http://doi.org/10.3758/BF03211656>
- Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica, 135*(2), 77–99. <http://doi.org/10.1016/j.actpsy.2010.02.006>
- Theeuwes, J., Belopolsky, A., & Olivers, C. N. L. (2009). Interactions between working memory, attention and eye movements. *Acta Psychologica, 132*(2), 106–114. <http://doi.org/10.1016/j.actpsy.2009.01.005>
- Tong, F., & Pratte, M. S. (2012). Decoding patterns of human brain activity. *Annual Review of Psychology, 63*(1), 483–509. <http://doi.org/10.1146/annurev-psych-120710-100412>
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., et al. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research, 168*(3), 242–249. <http://doi.org/10.1016/j.psychres.2008.05.006>
- Treisman, A. (1986). Features and objects in visual processing. *Scientific American, 255*(5), 114–125. <http://doi.org/10.1038/scientificamerican1186-114b>
- Treisman, A. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *The Quarterly Journal of Experimental Psychology Section A, 40*(2), 201–237. <http://doi.org/10.1080/02724988843000104>
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136. [http://doi.org/10.1016/0010-0285\(80\)90005-5](http://doi.org/10.1016/0010-0285(80)90005-5)

- Treue, S. (2003). Visual attention: The where, what, how and why of saliency. *Current Opinion in Neurobiology*, *13*(4), 428–432. [http://doi.org/10.1016/S0959-4388\(03\)00105-3](http://doi.org/10.1016/S0959-4388(03)00105-3)
- Tsvetanov, K. A., Arvanitis, T. N., & Humphreys, G. W. (2012). Dissociating effects of stimulus identity and load on working memory attentional guidance: Lengthening encoding time eliminates the effect of load but not identity. *The Quarterly Journal of Experimental Psychology*, *65*(8), 1475–1483. <http://doi.org/10.1080/17470218.2012.694895>
- Turatto, M., & Galfano, G. (2000). Color, form and luminance capture attention in visual search. *Vision Research*, *40*(13), 1639–1643. [http://doi.org/10.1016/s0042-6989\(00\)00061-4](http://doi.org/10.1016/s0042-6989(00)00061-4)
- Unsworth, N., Heitz, R. P., Schrock, J. C., & Engle, R. W. (2005). An automated version of the operation span task. *Behavior Research Methods*, *37*(3), 498–505. <http://doi.org/10.3758/BF03192720>
- van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, *121*(1), 124–149. <http://doi.org/10.1037/a0035234>
- van den Berg, R., Shin, H., & Chou, W. C. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, *109*(22), 8780–8785. <http://doi.org/10.1073/pnas.1117465109>
- van Moorselaar, D., Battistoni, E., Theeuwes, J., & Olivers, C. N. L. (2014a). Rapid influences of cued visual memories on attentional guidance. *Annals of the New York Academy of Sciences*, *1339*(1), 1–10. <http://doi.org/10.1111/nyas.12574>
- van Moorselaar, D., Olivers, C. N. L., Theeuwes, J., Lamme, V. A. F., & Sligte, I. G. (2015). Forgotten but not gone: Retro-cue costs and benefits in a double-cueing paradigm suggest multiple states in visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(6), 1755–1763. <http://doi.org/10.1037/xlm0000124>
- van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014b). In competition for the attentional template: Can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1450–1464. <http://doi.org/10.1037/a0036229>

- Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic Bulletin & Review*, *19*(5), 871–878. <http://doi.org/10.3758/s13423-012-0280-4>
- Vickery, T. J., King, L. W., & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of Vision*, *5*(1), 8–8. <http://doi.org/10.1167/5.1.8>
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*(6984), 748–751. <http://doi.org/10.1038/nature02447>
- Vuilleumier, P. (2005). How brains beware: Neural mechanisms of emotional attention. *Trends in Cognitive Sciences*, *9*(12), 585–594. <http://doi.org/10.1016/j.tics.2005.10.011>
- Vuilleumier, P. (2015). Affective and motivational control of vision. *Current Opinion in Neurology*, *28*(1), 29–35. <http://doi.org/10.1097/WCO.0000000000000159>
- Webb, N. M., Shavelson, R. J., & Haertel, E. H. (2006). Reliability coefficients and generalizability theory. *Handbook of Statistics*, 81–124. [http://doi.org/10.1016/s0169-7161\(06\)26004-8](http://doi.org/10.1016/s0169-7161(06)26004-8)
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), 1120–1135. <http://doi.org/10.1167/4.12.11>
- Williams, M., Hong, S. W., Kang, M.-S., Carlisle, N. B., & Woodman, G. F. (2013). The benefit of forgetting. *Psychonomic Bulletin & Review*, *20*(2), 348–355. <http://doi.org/10.3758/s13423-012-0354-3>
- Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science*, *9*(1), 33–39.
- Wolfe, J. M. (2007). Guided Search 4.0. *Integrated Models of Cognitive Systems*, 99–119. <http://doi.org/10.1093/acprof:oso/9780195189193.003.0008>
- Wolfe, J. M., & DiMase, J. S. (2003). Do intersections serve as basic features in visual search? *Perception*, *32*(6), 645–656. <http://doi.org/10.1068/p3414>
- Wolff, M. J., Ding, J., & Myers, N. E. (2015). Revealing hidden states in visual working memory using electroencephalography. *Frontiers in Systems ...*, *9*(123). <http://doi.org/10.3389/fnsys.2015.00123>

- Woodman, G. F., & Arita, J. T. (2011). Direct electrophysiological measurement of attentional templates in visual working memory. *Psychological Science*, 22(2), 212–215. <http://doi.org/10.1177/0956797610395395>
- Woodman, G. F., & Chun, M. M. (2006). The role of working memory and long-term memory in visual search. *Visual Cognition*, 14(4-8), 808–830. <http://doi.org/10.1080/13506280500197397>
- Woodman, G. F., & Luck, S. J. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin & Review*, 11(2), 269–274. <http://doi.org/10.3758/bf03196569>
- Woodman, G. F., & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology: Human Perception and Performance*, 33(2), 363–377. <http://doi.org/10.1037/0096-1523.33.2.363>
- Woodman, G. F., Carlisle, N. B., & Reinhart, R. M. G. (2013). Where do we store the memory representations that guide attention? *Journal of Vision*, 13(3), 1–1. <http://doi.org/10.1167/13.3.1>
- Woodman, G. F., Luck, S. J., & Schall, J. D. (2007). The role of working memory representations in the control of attention. *Cerebral Cortex*, 17(suppl 1), i118–i124. <http://doi.org/10.1093/cercor/bhm065>
- Woodman, G. F., Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science*, 12(3), 219–224. <http://doi.org/10.1111/1467-9280.00339>
- Yantis, S. (2000). Goal-directed and stimulus-driven determinants of attentional control. In S. Monsell & J. Driver (Eds.), *Control of Cognitive Processes Attention and Performance XVIII* (pp. 73–103). Cambridge, MA.
- Zelinsky, G. J., & Bisley, J. W. (2015). The what, where, and why of priority maps and their interactions with visual working memory. *Annals of the New York Academy of Sciences*, 1339(1), 154–164. <http://doi.org/10.1111/nyas.12606>
- Zhang, B., Zhang, J. X., Huang, S., Kong, L., & Wang, S. (2011). Effects of load on the guidance of visual attention from working memory. *Vision Research*, 51(23-24), 2356–2361. <http://doi.org/10.1016/j.visres.2011.09.008>
- Zhang, B., Zhang, J. X., Kong, L., Huang, S., Yue, Z., & Wang, S. (2010). Guidance of visual attention from working memory contents depends on stimulus attributes.

Neuroscience Letters, 486(3), 202–206.
<http://doi.org/10.1016/j.neulet.2010.09.052>

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

Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, 20(4), 423–428. <http://doi.org/10.1111/j.1467-9280.2009.02322.x>

Biography

Emma Wu Dowd was born in Dallas, Texas in July 1987. She received a bachelor's degree in linguistics from Harvard University in 2009. She then pursued graduate studies in the United Kingdom on a Marshall Scholarship, receiving a master's degree in psychology from University College London in 2010 and a master's degree in neuroscience from King's College London in 2011. She entered a doctoral program in cognitive neuroscience at Duke University in fall 2011, under the supervision of Drs. Tobias Egner and Stephen R. Mitroff. She received a master's degree in psychology and neuroscience from Duke University in 2014. She begins a postdoctoral fellowship at the Ohio State University in summer 2016 in the laboratory of Dr. Julie D. Golomb.

PEER-REVIEWED PUBLICATIONS

Dowd, E. W., Mitroff, S. R., & LaBar, K. S. (submitted). Fear generalization gradients in visuospatial attention.

Kiyonaga, A.*, **Dowd, E. W.***, & Egner, T. (submitted). Neural representation of working memory content is modulated by visual attentional demand.

* *Shared first-authorship*

Dowd, E. W., Kiyonaga, A., Beck, J. M., & Egner, T. (2015). Quality and accessibility of visual working memory during cognitive control of attentional guidance: A Bayesian model comparison approach. *Visual Cognition*, 23(3), 337-356.

Biggs, A. T., Adamo, S. H., **Dowd, E. W.**, & Mitroff, S. R. (2015). Perceptual and conceptual set biases in multiple-target visual search. *Attention, Perception, & Psychophysics*, 77(3), 844-855.

Dowd, E. W., Kiyonaga, A., Egner, T., & Mitroff, S. R. (2015). Attentional guidance by working memory differs by paradigm: An individual-differences approach. *Attention, Perception, & Psychophysics*, 77(3), 704-712.

- Mitroff, S. R., Biggs, A. T., Adamo, S. H., **Dowd, E. W.**, Winkle, J., & Clark, K. (2015). What can 1,000,000,000 trials tell us about visual search? *Journal of Experimental Psychology: Human Perception and Performance*, *41*(1), 1-5.
- Dowd, E. W.**, & Mitroff, S. R. (2013). Attentional guidance by working memory overrides salience cues in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(6), 1786-1796.
- Bays, P. M., **Wu, E. Y.**, & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, *49*(6), 1622-1631.
- Covington, H. E., Maze, I., Sun, H., Bomze, H. M., DeMaio, K. D., **Wu, E. Y.**, Dietz, D., M., Lobo, M. K., Ghose, S., Mouzon, E., Neve, R., Tamminga, C., & Nestler, E. J. (2011). A role for repressive histone methylation in cocaine-induced vulnerability to stress. *Neuron*, *71*(4), 656-670.
- Covington, H. E., Maze, I., LaPlant, Q. C., Vialou, V. F., Ohnishi, Y. N., Berton, O., Fass, D., M., Renthal, W., Rush, A., J., **Wu, E. Y.**, Ghose, S., Krishnan, V., Russo, S. J., Tamminga, C., Haggarty, S., J., & Nestler, E., J. (2009). Antidepressant actions of histone deacetylase inhibitors. *Journal of Neuroscience*, *29*(37), 11451-11460.

HONORS & FELLOWSHIPS

- 2013-2016** Conference Travel Awards, Duke University
- 2014** Pathfinder Fellow, Information Initiative at Duke, Duke University
- 2014** Summer Institute in Cognitive Neuroscience Fellow, UCSB
- 2014** Summer Research Fellowship, Duke University
- 2009-2011** Marshall Scholarship, Her Majesty's Government