

Where do we store the memory representations that guide attention?

Geoffrey F. Woodman

Vanderbilt University, Vanderbilt Vision Research Center,
Vanderbilt Center for Cognitive and Integrative
Neuroscience, Nashville, TN, USA



Nancy B. Carlisle

Vanderbilt University, Vanderbilt Vision Research Center,
Vanderbilt Center for Cognitive and Integrative
Neuroscience, Nashville, TN, USA

Robert M. G. Reinhart

Vanderbilt University, Vanderbilt Vision Research Center,
Vanderbilt Center for Cognitive and Integrative
Neuroscience, Nashville, TN, USA

During the last decade one of the most contentious and heavily studied topics in the attention literature has been the role that working memory representations play in controlling perceptual selection. The hypothesis has been advanced that to have attention select a certain perceptual input from the environment, we only need to represent that item in working memory. Here we summarize the work indicating that the relationship between what representations are maintained in working memory and what perceptual inputs are selected is not so simple. First, it appears that attentional selection is also determined by high-level task goals that mediate the relationship between working memory storage and attentional selection. Second, much of the recent work from our laboratory has focused on the role of long-term memory in controlling attentional selection. We review recent evidence supporting the proposal that working memory representations are critical during the initial configuration of attentional control settings, but that after those settings are established long-term memory representations play an important role in controlling which perceptual inputs are selected by mechanisms of attention.

our memory (James, 1890; Pillsbury, 1908). Modern theories of attention propose that the internal representations that we use to guide attention to the task-relevant objects we seek in cluttered environments are maintained in working memory (Bundesen, 1990; Bundesen, Habekost, & Kyllingsbaek, 2005; Desimone & Duncan, 1995; Duncan & Humphreys, 1989). Although a growing number of studies have sought to test this proposal using behavioral methods, the findings across experiments are sensitive to a number of factors that could make the pattern difficult to discern. In this paper, we review recent research suggesting that working memory representations can bias attention and often do. However, we also describe how recent electrophysiological studies indicate that the role working memory representations play in controlling attention is predominantly when we begin searching for something new. After we have tuned perceptual attention mechanisms to select a certain target or targets during search, target representations are no longer maintained in working memory and long-term memory representations of the searched-for items appear to control attentional selection.

Introduction

When we walk into our home to retrieve our keys, how do we tune our perceptual machinery to process these targets with the greatest priority? For centuries, psychologists and neuroscientist have proposed that mechanisms of attention are biased to select task-relevant information because of the representations in

Controlling perceptual attention

How do we find what we are looking for? For example, how do we find our suitcase on the packed airport carousel or our child in a clump of young soccer players on the field? This seemingly simple question has

Citation: 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, 1–17, <http://www.journalofvision.org/content/13/3/1>, doi:10.1167/13.3.1.

occupied cognitive scientists for centuries. In the laboratory, we commonly use visual search tasks to understand how we get our brains to select task-relevant information from the complex visual scenes that surround us (Neisser, 1964; Wolfe, 1998, 2003). The modern theoretical perspectives from cognitive psychology and neuroscience break this question down into understanding the bottom-up and top-down factors that guide our selective processing of the most important perceptual inputs.

Bottom-up factors are aspects of the perceptual input that make an item particularly salient relative to its surroundings. For example, one suitcase might be large and fluorescent orange, among smaller black bags. A child might be moving twice as fast as his teammates and competitors (or at half the speed). These heterogeneities in the features of certain potential targets result in attention being pulled to such items regardless of whether they are targets or not (Theeuwes, 1993; Treisman & Gelade, 1980; Wolfe, 1994). In contrast, top-down factors are those that allow us to focus our limited-capacity processing on task-relevant information despite the lack of bottom-up salience. For example, when you look for your black suitcase among many similar bags or the even more extreme task of searching for your small gray bag that is the same color as the conveyor belt among large brightly colored bags with floral prints and stickers from around the world. In this latter example, your target is far less salient than the brightly colored distractors. The top-down influences must be due to some type of internal representation of the thing you are looking for that serves to focus attention on potential targets despite the presence of potent distractors.

Generally, theories of attention propose that bottom-up factors interact with top-down factors to determine how easy it is to find something we seek (Bundesen, 1990; Bundesen et al., 2005; Desimone & Duncan, 1995; Duncan, 1996; Hamker, 2004; Reynolds & Heeger, 2009; Treisman, 1988). What we will discuss in this paper is what we currently know about the source of top-down attentional control. In addition to recent empirical data addressing the dominant hypothesis that the source of top-down attentional control is from working memory, we will also discuss recent studies addressing the hypothesis from theories of learning that working memory representations are particularly important during the acquisition phase of a task.

Working memory templates

The dominant theoretical explanation for how the visual system implements top-down control of visual-spatial attention is by holding a representation of the searched-for item in working memory. This mechanism

of top-down control is shared by leading computational (e.g., Bundesen, 1990; Bundesen et al., 2005) and descriptive models of attention in cognitive neuroscience (e.g., Desimone & Duncan, 1995). The idea is simple and elegant. If you want to look for your suitcase, you actively maintain a representation of your suitcase in visual working memory. By doing this, the working memory representation of that target feeds signals to the neurons in regions of the brain that perform perceptual analysis, targeting the neurons that code for those target features. This top-down source of activation increases the sensitivity of the neurons that detect the bottom-up activation of the target features. This increased activation allows attention to be biased to select target-like inputs to the visual system in the face of nontargets with stronger bottom-up activation.

Recent studies have sought to test the hypothesis that attentional templates in working memory are the source of top-down attentional control. The design of these experiments grows out of an elaboration of the basic logic of the proposal. If working memory representations control the allocation of attention to certain objects in our visual field, then attention should be inherently biased to select any perceptual input matching a representation in visual working memory (Carlisle & Woodman, 2011a, 2011b; Dalvit & Eimer, 2011; Downing, 2000; Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Olivers, 2009; Olivers & Eimer, 2011; Olivers, Meijer, & Theeuwes, 2006; Peters, Goebel, & Roelfsema, 2008; Soto, Heinke, Humphreys, & Blanco, 2005; Soto & Humphreys, 2007, 2008; Soto, Humphreys, & Rotshtein, 2007; Woodman & Luck, 2007). It should be the case that if controlling attention is as simple as holding a representation of the object we seek in working memory, then holding a representation in working memory should be sufficient to guide attention to matching items when they appear in our field of view. Note that this is the simplest version of this hypothesis in which just holding a representation in working memory is all that is needed to bias attention to select similar visual inputs. To test this hypothesis, essentially all of the studies have used a similar dual-task design. Subjects are asked to perform a visual search task for one item while concurrently maintaining another representation in working memory. Figure 1 shows an example used in one of these studies (Woodman & Luck, 2007). If attention is consistently deployed to memory-matching perceptual inputs, then the presence of a memory-matching item in the search array should influence performance of the visual search task.

Variations on the theme change the tune

The seminal study of Downing (2000) had observers respond to a probe stimulus while remembering a

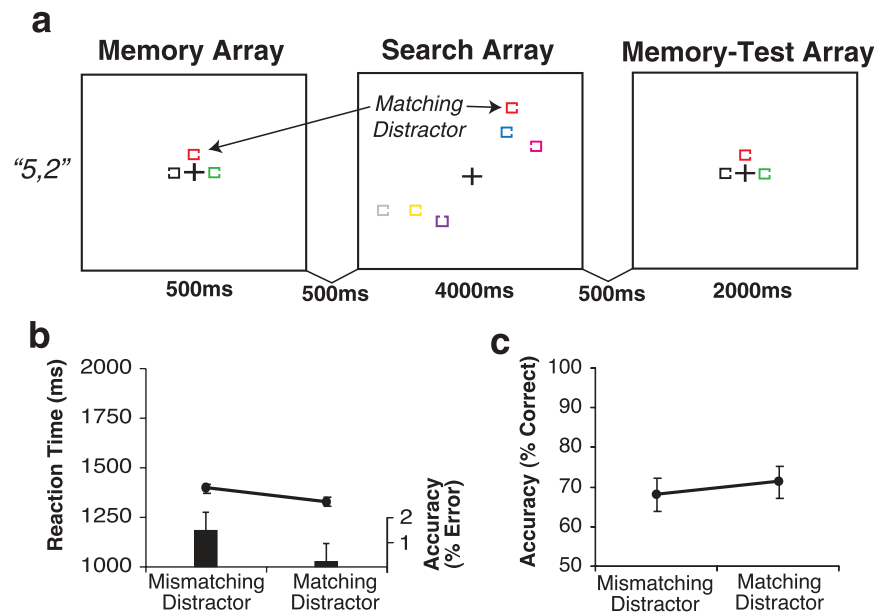


Figure 1. The stimuli and findings of a dual-task experiment examining the relationship between storage of information in visual working memory and the deployment of attention to similar items. (a) The stimuli from experiment 2 of Woodman and Luck (2007). (b) Visual search RT and (c) memory task accuracy (percent correct) for these tasks performed together. Adapted from Woodman and Luck (2007), with permission from the American Psychological Association.

picture of a face for a subsequent memory test. He showed that people were faster to respond to this probe when it was presented at the same location as a memory-matching face stimulus than at the location of a different face. These findings were interpreted as showing that people automatically attend to stimuli in the environment that match those represented in visual working memory.

A number of subsequent studies reported findings consistent with this conclusion (Olivers, 2009; Olivers et al., 2006; Soto et al., 2005; Soto & Humphreys, 2007, 2008; Soto et al., 2007). However, it has been argued that in a number of these studies, people may have adopted a strategy of attending to the memory-matching item (Woodman & Luck, 2007). This is advantageous when the visual search target could match the item in memory, allowing subjects to notice the benefit of voluntarily shifting attention to the memory-matching item first. In addition, it is possible that voluntarily shifting attention to the memory-matching item could facilitate the maintenance of that memory representation in the face of the new perceptual inputs from the visual search array. Shifting perceptual attention to the memory-matching item could mitigate the disturbance of the normal maintenance of the to-be-remembered information by the onset of the visual search array (Woodman & Luck, 2007, 2009). This has been called *perceptual resampling* (Woodman & Luck, 2007) or *refreshing* (e.g., Yi, Turk-Browne, Chun, & Johnson, 2008). These terms describe the idea that working memory maintenance benefits

from re-encoding the to-be-remembered information when it reappears in the environment (but see Balani, Soto, & Humphreys, 2010, for evidence that perceptual resampling cannot explain all such attention effects to memory matching items).

To deal with the alternative explanation that people might be strategically attending to items in the visual field that match a representation in memory, a number of laboratories adopted a paradigm in which the memory-matching item in the visual search array can only hurt search performance if it is attended (e.g., Houtkamp & Roelfsema, 2006; Olivers et al., 2006; Soto et al., 2005; Woodman & Luck, 2007). This is accomplished by always having the memory-matching item appear as a distractor in the visual search array throughout the experiment. Downing and Dodds (2004) were the first to implement such a design and showed the memory-matching distractor did not appear to capture attention because it did not slow visual search. Indeed, their findings suggested that visual search reaction times (RTs) were, if anything, faster when the memory-matching distractor was present in the visual search array. In addition, memory-task performance was better when the memory-matching distractor appeared in the visual search array consisting of three objects. These same patterns of results were found by Woodman and Luck (2007). With either a sufficiently high working memory load or perceptual load of the search array, these benefits of the presence of the memory-matching item on search (see Figure 1b) and the memory task (see Figure 1c) were

reliable and significant. If memory-matching items were capturing attention, then we would expect search RTs to be slower, not faster, when these special distractors were present. This led the latter paper to conclude that the deployment of attention to items matching a representation in working memory is not automatic but is instead under strategic control. This conclusion follows from the fact that subjects knew that the memory matches would never be the targets in the visual search arrays, and this allowed subjects to strategically avoid shifting attention to memory-matching items during visual search. Because these items could never be the target, this speeds RT relative to when no memory-matching item was present and all of the search items were possible targets. If the array remained on the screen after the search response, the memory match could be attended after search was complete to aid in maintaining the representations (Woodman & Luck, 2007).

Findings of the type we just discussed have led to a more sophisticated account in which there appears to be consensus. It does not appear that holding a representation in working memory is sufficient to guarantee attention will be biased to similar perceptual inputs, despite findings being consistent and reliable using a specific set of stimuli and methods (see Olivers, 2009; Olivers, Peters, Houtkamp, & Roelfsema, 2011). In particular, it appears that a number of factors can modulate or eliminate the effect of a memory item on the deployment of attention. These include the probability of the target matching a representation held in working memory (Carlisle & Woodman, 2011a), the timing between the presentation of a to-be-remembered item and the array that might contain a matching item (Dombrowe, Olivers, & Donk, 2010; Han & Kim, 2009), the bottom-up salience of the stimuli and the consistency of the search target (Olivers, 2009), the strength of the representation in working memory (Soto & Humphreys, 2008), as well as the goals of the search task being performed (Carlisle & Woodman, 2011b). This has culminated in the perspective that simply having a representation in visual working memory is not sufficient for attention to be biased toward perceptual inputs that match that representation (e.g., Carlisle & Woodman, 2011a, 2011b; Kiyonaga, Egner, & Soto, 2012; Olivers et al., 2011).

This more moderate hypothesis is also consistent with neuroscientific studies of the effect of working memory representations on the deployment of perceptual attention. Several groups of researchers have used event-related potentials (ERPs) to study the interaction of perceptual attention and working memory representations. Peters et al. (2008) used the P3 component of subjects' ERPs to determine whether inputs matching a representation held in visual working memory were treated like targets in a sequential stream of items. The

parietal P3 is sensitive to the probability of the eliciting category of stimulus and a leading hypothesis proposes that it is related to the updating of memory (Polich, 2012). Peters and colleagues found that memory-matching distractors elicited a P3 that was identical to that of the other distractors that did not match an item in memory, concluding that memory-matching items do not engage limited-capacity processes like targets of the task at hand. In contrast, Kumar, Soto, and Humphreys (2009) sought to directly measure the deployment of covert attention to items in a search array using the N2pc component. The N2pc is a component typically elicited in the time range of the N2, with a posterior and contralateral distribution (hence N2-posterior-contralateral). Kumar and colleagues found that the presence and location of memory-matching distractors modulated the amplitude of the N2pc to the search target, consistent with the idea that covert attention was shifted to these special distractors, at least on some trials. However, using this same index of the covert attention, Carlisle and Woodman (2011b) found that only the current goals of the search task governed the measurements of the N2pc component to memory matches. Specifically, they reported that the memory-matching items elicited a contralateral positivity beginning shortly after the N2pc to the search target, with this positivity indexing a process of distractor suppression (Hickey, Di Lollo, & McDonald, 2009). These different patterns of results demonstrate that whether we attend to items that match the representations in working memory is not simply a matter of maintaining that item. Instead the stimulus and timing parameters used in the experiments appears to be critical (Dombrowe et al., 2010; Han & Kim, 2009; Olivers, 2009) as well as the strategy and high-level goals of the subject in these dual-task experiments.

Adaptive and flexible cognitive control

How do we reconcile the apparently discrepant views that we just discussed regarding the role that working memory plays in providing top-down control over attention? The idea that we have strategic control over how working memory and perceptual attention interact might initially seem like a convenient ad hoc explanation for any pattern of results. However, if we turn a more critical eye to the theories we are testing, then it seems evident that these theories already proposed that we should have strategic control over how working memory representations are used to guide perceptual attention. To illustrate this point we are going to focus on the most highly specified model of attention that we cited above when arguing that theories propose visual working memory representations provide top-down control over attention. This discussion illustrates how

strategic control over the influence of working memory representations is already predicted by at least some of these models.

In previous sections we cited Bundesen's theory of visual attention (TVA; Bundesen, 1990; Bundesen et al., 2005) as an example of a model that explicitly proposes that visual working memory representations are used to control attention. This modeling framework is among the most mathematically sophisticated accounts of attentional selection during visual search and depends upon working memory representations to derive attentional weights. This framework proposes that attentional weights for a given feature of an object represented in working memory can either be set high (i.e., near 1) or low (i.e., near 0). In addition, these attentional weights can be dynamically changed in dual-task paradigms with the executive control of TVA (Logan & Gordon, 2001). For example, if it is initially advantageous to set an attentional weight low for the object features represented in working memory, then attention will be biased away from perceptual inputs with those features. However, if the observer is then asked to switch to a different task (i.e., from visual search to a change-detection task) in which attending to items matching the representation in working memory is necessary, then this is implemented by updating the attentional weights (Logan & Gordon, 2001). Behavioral studies are consistent with this idea that people dynamically change attentional weights in anticipation of performing different tasks during the dual-task procedures described above (Olivers & Eimer, 2011). Moreover, the most recent version of TVA (Bundesen et al., 2005) proposes that even after a representation is held in visual working memory it is necessary for higher-level executive control to configure the system to use this representation to set attentional weights. Thus, this theory does not propose that we are examining a static state of the attentional system in terms of the level of the parameters and representations that set attentional weights for objects. Instead, we are examining an attentional system in which flexibility based on strategic control is the rule, not the exception.

Directly measuring working memory templates in the brain

It is useful to pause and remember that the behavioral experiments sought to indirectly test the hypothesis that working memory representations are the source of the top-down control of attention. That is, all of these studies implicitly or explicitly assumed that if working memory representations normally guide attention, then any representation stored in working memory should influence how attention is deployed. As

already discussed, there is reason to doubt that assumption holds (see also Olivers et al., 2011). It appears that one can have representations in working memory that do not influence how attention is deployed or even guide attention away from similar items in the visual field (Arita, Carlisle, & Woodman, 2012). Thus, the skeptical reader may wonder whether we are even asking the right question. What is the evidence that top-down attentional control comes from working memory representations?

The idea that working memory is the locus of the representations that control attention comes from a series of neurophysiological experiments by Chelazzi and colleagues (Chelazzi, Duncan, Miller, & Desimone, 1998; Chelazzi, Miller, Duncan, & Desimone, 1993). They had monkeys perform a cued visual search task in which the target was shown several seconds before a simple visual search array was presented. They showed that neurons in the inferotemporal cortex (i.e., area IT) that respond vigorously to the cued target object, maintained an elevated firing rate during the interval between the cue and search array presentation. This activity during the delay period was interpreted as evidence that a representation of the target object was actively maintained in working memory in anticipation of the need to locate the target in the search array. Findings such as these were critical in justifying the locus of the representations that are used to guide attention in the limited-capacity store of visual working memory (Bundesen, 1990; Duncan & Humphreys, 1989).

In the time since the demonstrations of Chelazzi and colleagues (Chelazzi et al., 1993; Chelazzi et al., 1998), new empirical findings have clouded the picture. Although the findings from the recordings in IT were consistent with the idea that actively maintained representations in working memory were driving target-selecting neurons even before the search array appeared, recordings from area V4 did not clearly show this pattern (Chelazzi, Miller, Duncan, & Desimone, 2001). This was surprising given the explanation that the working memory representation of the target was feeding back to neurons in perceptual areas, like V4, allowing them to be more sensitive to perceptual input with target features when the visual search array appeared. Given this, we would expect to find that the V4 neurons that code for the target features would also show an elevated firing rate between the cue and the search array, as did the neurons in IT. The observations from the first recordings might not be as ubiquitous as we would hope given that this mechanism is supposed to be the source of attentional control. For example, perhaps the monkeys in the V4 study did not consistently rely upon working memory as the source of top-down control. In addition, from these neurophysiological studies, it is unclear whether humans rely upon

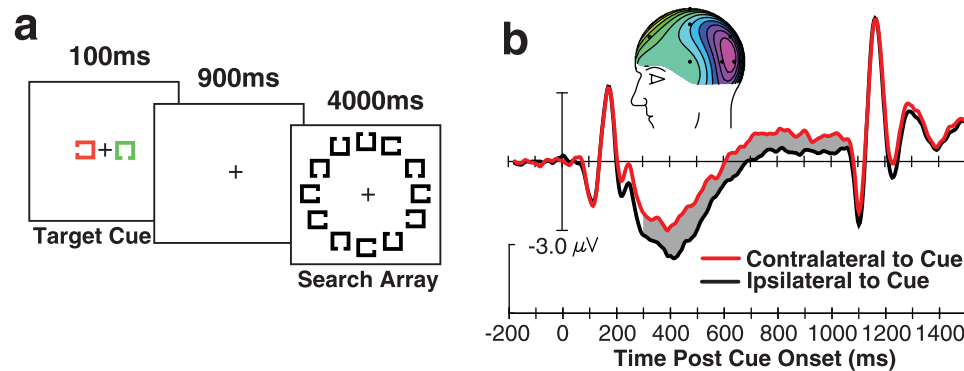


Figure 2. The stimuli and findings for an example experiment measuring ERPs from subjects performing a cued visual search task. (a) Example of the stimulus sequence and (b) the grand-average waveforms from electrodes T5/6, contralateral (red) and ipsilateral (black) to the location of the cue on each trial. The gray region shows the epoch in which the significant CDA was measured and inset shows the voltage distribution. Adapted from Woodman and Arita (2011), with permission from the Association for Psychological Science.

the same cognitive mechanisms to control attention during visual search as do monkeys.

Recently we have sought to determine whether humans maintain representations in visual working memory to provide top-down control over the deployment of attention during visual search. In this work, we used ERPs to directly measure the electrophysiological signature of these attentional templates in working memory during search. This was possible because recent ERP research has shown that when people are maintaining the representations of items in visual working memory initially seen in the left visual field, a sustained negativity is found over the right hemisphere. The opposite pattern is found when to-be-remembered stimuli are presented in the right visual field. As a result, this memory-related ERP component has become known as the contralateral delay activity (CDA; Ikkai, McCollough, & Vogel, 2010; Vogel & Machizawa, 2004; Vogel, McCollough, & Machizawa, 2005). This CDA, discovered in humans performing short-term memory tasks, provides a critical tool with which to address the question of whether we hold attentional templates in visual working memory.

The idea was simple enough, if we presented a lateralized cue to a subject that indicated the target in an upcoming visual search array, would we find that the cue elicited a CDA that continued through until the search array was presented, as Chelazzi and colleagues found with the neurons in IT? Figure 2A shows the stimulus sequence from an example trial in one of these experiments (Woodman & Arita, 2011). After acquiring fixation, a brief cue array indicated the stimulus the subjects were to search for, with the task-relevant cue indicated by color (e.g., green) and the task relevant color alternating across blocks of trials to provide physical stimulus control. Physical stimulus control refers to the fact that presenting items in both hemifield removes the concern that a single item in one hemifield

elicits ERP components that can differ from the earliest ERP components onward, making it impossible to differentiate the lateralized sensory responses from potential memory-related responses (Luck, 2005; Woodman, 2010). Nine hundred milliseconds after the target cue was presented, a search array appeared and the observers had to report with a manual button press whether the target object was present or absent in the search array.

The waveforms shown in Figure 2B illustrate the results we obtained across a number of similar experiments. First, the cue elicited the expected sensory responses immediately following its onset, just as the target cues did in the single-unit recordings from the temporal lobes of monkeys (Chelazzi et al., 1993; Chelazzi et al., 1998). The crucial observation is that following the sensory responses to the cue array we found that contralateral to the task-relevant cue a sustained negativity emerged at posterior electrode sites relative to ipsilateral sites. This CDA was sustained until it was obscured by the presentation of the search array, mirroring a similar abolishment of the template-related activity by the search array presentation in the temporal lobe neurons of Chelazzi and colleagues. These findings show that humans exhibit a pattern of electrophysiological activity just like that initially used to support the proposal that nonhuman primates maintain attentional templates in visual working memory during search.

We performed additional manipulations and analyses to further test the conclusion that we had found a direct measure of the attentional templates in human visual working memory. First, we found that when we cued two possible targets presented in one hemifield that the amplitude of the CDA measured prior to the search array onset was twice as large as when we cued observers to search for a single target (Carlisle, Arita, Pardo, & Woodman, 2011). In addition, we found that

the amplitude of each subject's presearch CDA could be used to predict their search RT (Carlisle et al., 2011) or accuracy before search had even begun (Woodman & Arita, 2011). This is consistent with the idea that the presearch CDA provides a direct measure of the fidelity of target representation in visual working memory that is then used to guide attention in the subsequent search array. If subjects maintained a low quality representation or lost the target representation on a subset of trials, then their CDA would be smaller and their search less efficient. The results we obtained from measurements of the individual observers' presearch CDAs neatly conformed to these predictions.

These findings would appear to indicate that working memory representations are used to guide attention during visual search. However, these ERP findings conflict with dual-task experiments showing that the efficiency of visual search is unchanged when visual working memory is filled with object representations to enable the performance of another concurrent task, compared to conditions in which search is performed alone (Logan, 1978, 1979; Woodman & Luck, 2009; Woodman, Vogel, & Luck, 2001). In these behavioral experiments, people were asked to encode a set of items into visual working memory for a later change-detection judgment, then they searched for the presence of a target object in a visual search array, and finally had their memory tested at the end of the trial on the memoranda presented at the beginning. The slopes of the search functions relating RT to set size were not different between this dual-task condition and one in which people performed search in isolation. This pattern of results was obtained across decades, laboratories, and many experiments, demonstrating the robustness of the findings (Logan, 1978, 1979; Oh & Kim, 2004; Woodman & Luck, 2009; Woodman et al., 2001).

However, all of these experiments involved having observers search for the same target or targets across all trials. When observers were required to search for a different item on every trial, significant interference was found (Woodman, Luck, & Schall, 2007) in the form of steeper search slopes during search with working memory full compared to search performed alone. As described below, these findings were foreshadowed by models of learning that intimated that if attentional templates existed they would be stored in visual working memory when targets change across trials, but in long-term memory when target identity was stable across trials.

The infinite-sized elephant in the room

The inner workings of models of learning and skill acquisition (Anderson, 1982, 2000; Logan, 1988, 2002)

point to an important factor in the processing of complex scenes that our empirical studies of the role of working memory in controlling attention had overlooked. One of the most prominent findings in the visual search literature from the 1960s and 1970s was the observation that there are large learning effects when people search complex scenes for the same target or set of targets trial after trial (Neisser, 1963; Nickerson, 1966; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). This effect has returned to the attention of researchers more recently (e.g., Wolfe, 2012). What ultimately emerged from these empirical observations was the theoretical explanation of these effects in terms of long-term memory representations controlling processing with sufficient practice (Anderson, 2000; Logan, 1988, 1990, 2002; Rickard, 1997). This idea motivated the electrophysiological experiments that we turn to now.

In a recent series of experiments, we sought to use ERPs to directly measure the memory representations controlling attention during brief bursts of learning to search for specific target objects (Carlisle et al., 2011). The paradigm is similar to that described previously in which observers were cued on each trial to look for a specific target object while ERPs were recorded (see Figure 3a). The novel aspect was that we now built in intertrial target repetitions. Across short runs of trials (three to seven trials long), we cued observers to search for the same target object trial after trial before changing the cued target to a different shape. We predicted that if long-term memory quickly acquires instances that can be used to control the deployment of attention, then we should be able to watch the target representation in working memory disappear as long-term memory takes over the control of attention during visual search.

In Figure 3b, we show the CDA results of our first experiment in which we recorded ERPs during these bursts of learning (i.e., experiment 3 of Carlisle et al., 2011). We found that the presearch CDA to the target cue systematically decreased in amplitude across trials of searching for the same target. This was not simply due to a loss of engagement in the task or a reduction in search efficiency. Rather, RTs became faster across these same-target trials. This speeding of RT across trials of searching for the same target would be explained by learning theories as due to an increasing reliance upon long-term memory to control performance (Logan, 1988). The logic is that long-term memory retrieval is fast and automatic. The target in the search context can be used to automatically retrieve the response that is needed to respond correctly. This means that improvements in task performance should be accompanied by a reduced reliance upon visual working memory to control attention in these situations. These findings provided a definitive demonstra-

tion that working memory representations only control the deployment of attention to certain objects for a handful of trials during learning before long-term memory comes to dominate this selection process.

Measuring the footprints of long-term memory along the path of search

In the discussion of our findings so far we have described how we observed an electrophysiological index of the visual working memory representations controlling attention disappear during visual search for the same target trial after trial. This is a necessary observation if the hypothesis is correct that long-term memories take over the guidance of attention from working memory after a fairly modest amount of learning (Anderson, 1982; Logan, 1988; Rickard, 1997). We inferred that long-term memory must be taking over the role of attentional guidance if working memory was giving it up; however, we desired evidence to support this proposal other than the absence of a working memory signal. To this end, we appealed to the ERP literature that has studied long-term memory and the processing of stimuli across repeated presentations (Paller, Voss, & Boehm, 2007; Rugg & Curran, 2007).

Much of the long-term memory ERP work has focused on memory for verbal items or easily nameable materials (words, letters, digits, or familiar objects), and it appears that the ERPs measured when these stimuli are stored in memory are different than those found when stimuli are not readily verbalized (Danker et al., 2008; Paller, Lucas, & Voss, 2012; Voss, Schendan, & Paller, 2010). When the stimuli are not verbal in nature, modulations of an early frontal positivity (the anterior P1) can be used to predict the size of an individual participants' behavioral priming effect well before the response is measured (Voss et al., 2010). Encouraged by this finding, we re-examined our ERPs from frontal sites recorded during experiment 3 of Carlisle and colleagues (Carlisle et al., 2011) and found that we also observed that the effect that Voss and colleagues called the P170 (Voss et al., 2010) was modulated in our task across repeated cues to search for the same object. This component had the same frontal maximum observed in the study of Voss and colleagues. Specifically, Figure 4 shows that the amplitude of the anterior P1 or P170 at Fz became systematically more negative each time subjects were cued to search for the same object. This resulted in an analysis of variance with the factors of target repetition (one, two, three, four, five, six, or seven trials in a row with the same target) and electrode site (Fz, Cz, or Pz) yielding a significant main effect of electrode site, $F(2, 34) = 7.25$, $p < 0.01$, as well as an interaction of target




Figure 3. The stimuli and grand average ERP results, time-locked to the cue presentation (yellow on time axis) until search array onset (shown in blue on time axis) from groups of consecutive trials. Bottom panel shows the CDA amplitude across consecutive trials with the same search target. The gray line shows the power-function fit and the error bars represent ± 1 SEM. Reprinted from Carlisle et al. (2011), with permission from the Society for Neuroscience.

repetition and electrode, $F(12, 204) = 2.16$, $p < 0.05$. Follow-up tests confirmed that this interaction was due to the effect of target repetition being significant at Fz, $F(6, 102) = 5.14$, $p < 0.01$, but not Cz or Pz, $F_s < 1.0$. The onset of this frontal effect seems too early to be explained by eye movements, but to rule this out we examined the ocular electrode channels. We did not




Figure 4. The midline ERPs locked to target cue onset as a function of target repetition. Insets show current density distribution (repetition 7) and P170 amplitude by target repetition (in μV) measured from 150 to 200 ms. Adapted from Carlisle et al. (2011), with permission from the Society for Neuroscience.

find eye movements or blinks in the ocular channels similar to those found at the frontal electrodes, as is often the suspicion with frontal ERP effects (see the bottom, left panel of Figure 4). Thus, it appears that we can measure the systematic changes of the response of the brain to these repeated objects using an ERP component indexing the accumulation of information in long-term memory.

Beyond the theoretical implications, our findings of independent components indexing the role of visual working memory and long-term memory in the performance of visual search have practical implications. The CDA and the anterior P1 (or P170) components have distributions that are sufficiently different from each other such that we can measure these visual working memory and long-term memory effects independently, without cross contamination. These effects also have strikingly different time courses. The frontal effect is elicited very early in the trial (i.e., during the anterior P1), with the CDA elicited much

later. This demonstrates that the P170 and CDA are not due to modulations of the same underlying dipole.

This work establishes a set of ERP tools to independently measure the contributions from visual working memory and long-term memory to task performance. By noninvasively measuring these components of the subjects' waveforms we can track the involvement of different memory representations in controlling cognitive processing on a trial-by-trial basis. Figure 5 shows this graphically by replottting the data from (Carlisle et al., 2011). Here we see that the amplitude of the CDA, indexing the target templates in visual working memory, and the P170, indexing target representations in long-term memory can be independently measured across trials. As the CDA disappears, we can simultaneously observe the P170 effect systematically increase.

We next performed an analysis of this CDA and P170 activity across the course of the experiment to determine how working memory representations versus long-term memories are utilized as the experimental session unfolds. The analyses we have discussed so far focused on the CDA and P170 across trials within a run of trials lasting about a minute at the most. However, it is also possible to examine these ERP components across much longer intervals of time. For example, do these components change across the hours that it takes to collect the ERP data from the observers? In a previous ancillary analysis we found that the amplitude of the CDA elicited by a specific target object cue (i.e., a Landolt-C with a gap 45° to the right of 12 o'clock) did not change across the experiment as people returned to search for that same target object (Carlisle et al., 2011). As shown in Figure 6a, even after searching for that same target during 23 runs of trials (the black symbols and function), another 23 additional




Figure 5. The ERP component amplitudes following the target cue onset as a function of target repetition. The P170 amplitude was measured from 150 to 200 ms postcue and the CDA from 300 to 900 ms postcue. The error bars show ± 1 SEM. Adapted from Carlisle et al. (2011), with permission from the Society for Neuroscience.




Figure 6. (a) Target-locked CDA and (b) P170 amplitudes as a function of target repetitions averaged across the first (black), middle (red), and last 1/3 of runs (green). Error bars represent ± 1 SEM. Adapted from Reinhart and Woodman (in press).

runs of trials (the red symbols and function), the CDA amplitude was not different when the final 23 runs of searching for the same target objects were presented (the green symbols and functions). These data were aggregated across several experiments with the same stimuli and design shown in Figure 3a in which the same target was searched for across runs of trials (Reinhart & Woodman, in press).

The observation that the CDA to the target cues shows the same effect of target repetition across the entire experiment might seem counterintuitive given that people are coming to rely upon storage of these targets in long-term memory to control attention. Why did we not see the CDA get smaller as observers returned to search for the same object later in the 3-hour experiment? Although this might not be obvious without reading about the methods in these studies (Carlisle et al., 2011; Reinhart & Woodman, in press), this is actually consistent with the design of these experiments and the existence of proactive interference across the experiment. When we changed the target to a new object, this meant that the old target had now become a distractor. Classic experiments in cognitive psychology showed that a potential limiting factor in the performance of a task can be interference from the task-relevant information presented on previous trials (Peterson & Gentile, 1965; Wickens, Born, & Allen, 1963). We can see direct evidence for this building proactive interference in the amplitude of the P170

across the course of an experiment. Figure 6b shows that across repeated runs of searching through these same elements the P170 would “reset” somewhat to a lower level at the beginning of the run compared to the end of the previous run, but that this level rose across the experiment as would be expected if representations of the individual items were accumulating in long-term memory.

It has been proposed that one of the primary roles that working memory representations might play is to overcome proactive interference from long-term memory representations so that performance on a current task will not be harmed by previous experience (Kane & Engle, 2002). Our observation that the CDA returns to its full amplitude at the start of each run, even after hours of searching through the same set of objects, is consistent with the idea that these working memory representations need to be reinitiated at the beginning of each run to mitigate the proactive interference from the previous run of trials. Thus, we appear to have found support for the proposal that one of the primary roles of working memory is to overcome the potentially harmful effects of long-term memory using these electrophysiological measurements of observers’ memories.

This shows an example of how these ERP measures of working memory and long-term memory can be used to test fundamental hypotheses about the functions of these different memory systems during the performance of the cued visual search task we have used. Our next goal was to see if we could use these tools to determine the locus of other effects of interest in the attention and visual search literature.

Using ERPs indexing working and long-term memory to localize reward effects

With our initial work establishing the CDA and the anterior P1 or P170 as a useful set of tools to determine the locus of the top-down control of attention, we next sought to use these tools to answer a new question. Specifically, recent studies have shown that reward can be offered during a variety of cognitive tasks to increase the efficiency of attentional selection (Della Libera & Chelazzi, 2006, 2009; Hickey, Chelazzi, & Theeuwes, 2010, 2011; Navalpakkam, Koch, Rangel, & Perona, 2010; Raymond & O’Brien, 2009; Serences, 2008). Our next goal in this line of work was to determine the source of reward benefits during attention-demanding tasks like visual search. Given the existing evidence that the opportunity to earn a reward for task performance can modulate the size of attention effects, we wondered

if the source of these reward benefits might be the concurrent use of multiple memory representations to provide converging control over attention in situations in which just one type of memory representation might suffice.

Our experiments described above showed that within five to seven trials, long-term memory representations of the target had essentially taken over the control of attention from working memory during visual search. However, we hypothesized that it might be possible to bring working memory back online to help supplement the control of attention provided by long-term memory given sufficient reason to do so (Reinhart & Woodman, in press). The reason we gave subjects was the opportunity to earn monetary reward on a given trial.

The paradigm we used was a straightforward modification of those we have already discussed. Figure 7a shows how we added a reward cue prior to our presentation of the target cue, followed by the search array, with the trial ending with feedback about how much reward had been earned (in points that translated into US cents). On many runs of trials, all of the trials were low reward, meaning that subjects received 1 cent for performing the search task correctly. On a subset of trials during the runs of same-target trials, the trial began with a cue indicating that a high reward could be earned (i.e., 5 cents). The critical sequences for our hypothesis were those in which the high-reward cue occurred on the fifth trial of searching for the same target. Based on our previous work, we knew that this would be a time in which subjects would be beginning to rely on long-term memory to control attention. We predicted that if reward triggers the use of multiple memory representations to control selection during the processing of complex scenes, then we should see that the large-reward cue triggers the return of the CDA even after five trials of searching for the same target object.

The behavioral results are shown in Figure 7b. During these runs, we again observed bursts of learning in which RTs decreased across trials of searching for the same object (i.e., the white bars). On the critical sequences, a high-reward cue preceded the fifth trial of searching for the same target, followed by low-reward trials. The RTs from these trials are shown with a green bar on the fifth trial, followed by gray bars. This high-reward cue resulted in a drop in RT that then lingered for the next couple of trials as RTs floated back to the baseline, low-reward RTs. Search accuracy was uniformly at ceiling across all of these trial types (i.e., above 95% correct with no significant differences).

Across the same-target runs of trials, we replicated our basic finding of a systematically decreased CDA amplitude indexing working memory maintenance of the target representation across trials (see Figure 7c). However, when the fifth trial in these sequences was

preceded by a high-reward cue, the CDA returned to its full amplitude. This is as expected if working memory were being recruited to supplement the control of attention provided by long-term memory. To directly assess the nature of the long-term memory representations we concurrently measured the P170. As shown in Figure 7d, the amplitude of the P170 was found to systematically vary with each trial searching for the same target. Moreover, the amplitude of the P170 was insensitive to the reward cue. The blue squares connected with solid lines in Figure 6D show that the P170 did not significantly change following a high-reward cue relative to the runs of trials with all low-reward cues evidenced by the overlapping confidence intervals. Thus, our findings demonstrate that visual working memory is quickly recruited to respond to the opportunity to earn additional reward, whereas long-term memory appears to accumulate information on each trial in an invariant way.

We reasoned that if the recruitment of visual working memory was responsible for the drop in RT following a high-reward cue, then we should be able to use the magnitude of the rebound in the CDA to predict the behavioral benefit of reward on RT that would follow when the search array was presented. Figure 7e shows that for an individual subject, when the rebound of the CDA was larger, the RT benefit triggered by high reward relative to low reward was larger. We found that this relationship between the use of visual working memory and reward was not mediated by the individual's working memory capacity. In a separate task we estimated each subject's visual working memory capacity for storing simple colored squares. We found that regardless of the individuals' working memory capacity, they all used this strategy of recruiting working memory to supplement long-term memory to respond to the high-reward cues and perform the search task as quickly and accurately as possible (Reinhart & Woodman, in press). This indicates that the use of multiple memory representations to control attention is a general mechanism that is used across people that vary in visual working memory capacity.

In summary, our experiments with reward cues show an example of how the concurrent measurement of the working memory and long-term memory representations controlling attention can be used to determine the locus of these experimental effects. These ERP components offer the possibility of tracking the involvement of different memory representations that respond in qualitatively different ways to changing environmental demands. Whereas visual working memory decreases its control across learning, long-term memory increases its control. In addition, when the environment indicates that the stakes are high, visual working memory is quickly recruited while long-term memory

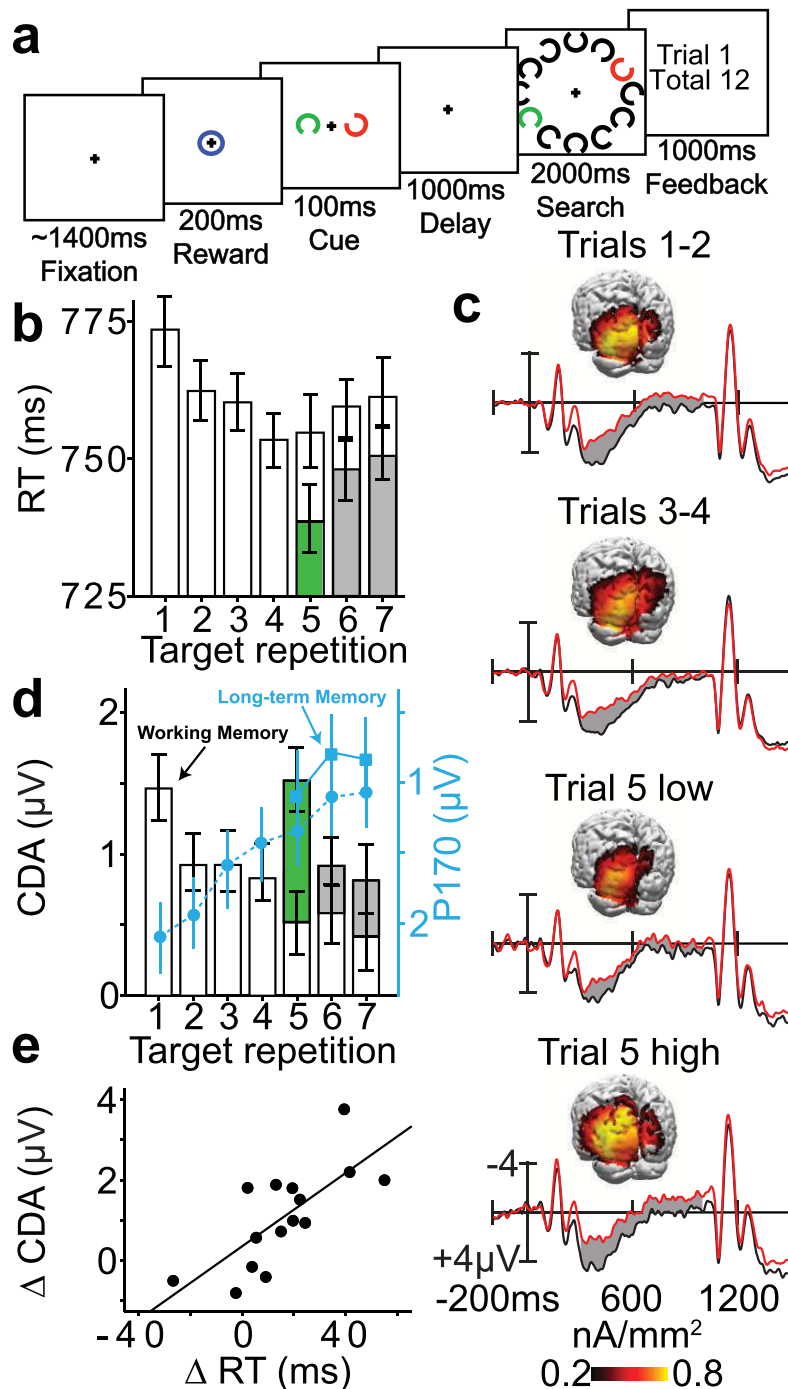


Figure 7. The procedure and findings of an experiment designed to determine the locus of reward-based attention effects. (a) The stimulus sequence in which a blue or yellow circle cued observers that a low or high reward could be obtained for performing the upcoming trial of visual search. (b) The RT in milliseconds on a run of low-reward trials (unfilled bars) and following a high-reward cue before the fifth target repetition shown in green (and the subsequent trials being low reward, gray bars). (c) The event-related potential waveforms with the red traces showing the contralateral waveforms from electrodes T5/6 and the black showing the ipsilateral waveforms. The insets show the current density models on a representative brain. (d) Plots of the mean CDA amplitude, using a format similar to b, and the mean P170 amplitude across low-reward runs (dashed blue line) and following a high-reward cue on the fifth target repetition (solid blue line). (e) The scatterplot showing that the size of the rebound of the CDA following a high-reward cue predicted the size of the search RT benefit that would follow on those trials. Each dot shows an individual subject's change in CDA amplitude and RT on fifth high reward versus fifth low-reward trials. This change was not predicted by the subjects' visual working memory capacity despite considerable variability of capacity across subjects. Adapted from Reinhart and Woodman (in press).

plods along accumulating instances of each trial of task performance.

Conclusions

The nature of the memory representations used to control attention is one of the oldest topics in cognitive science (i.e., James, 1890; Pillsbury, 1908). The interest in the source of this top-down control has only increased across time (e.g., Bundesen & Habekost, 2008; Bundesen et al., 2005; Desimone & Duncan, 1995). Here we reviewed the evidence that observers strategically control how working memory representations interact with the selection of incoming perceptual information during visual search. In addition, we reviewed our evidence for which we now have the electrophysiological tools to noninvasively measure both the working memory and long-term memory representations that determine which perceptual inputs to the visual system are selected by attention. The CDA provides a sensitive measure of the quality of the visual working memory representations that are being held online to guide attention. At the same time, the anterior P1 modulation, also known as the P170, can be measured to determine the nature of the long-term memory representations that are accrued during task performance.

Although our findings demonstrate that our ERP tools exhibit dissociations in line with a number of theories of learning to attend to certain items (e.g., Logan, 1988, 2002), there are a number of questions remaining about the types of representations we are measuring. Some of this work will need to focus on the nature of the ERP components themselves and other lines on the relationship between the findings using this paradigm and other work we reviewed earlier in the paper.

First, the precise nature of the visual working memory representations indexed by the CDA is not yet definitively known. For example, in our cued-search task, subjects were shown a target to search for that could appear in either the left or the right visual field. Because observers were equally accurate at finding the cued target on the left or on the right side, even on the first trial after a target change, the representation of the target in working memory must have effectively influenced neurons that process inputs across the entire visual field. This is despite the lateralized nature of the CDA component we were measuring. However, we do find that RTs tend to be faster when the target in the search array is presented in the same hemifield as the target cue. The degree to which the visual working memory representations are bilateral will be a question for future research. It seems likely that with the CDA

we measure a slight hemispheric imbalance in maintenance-related activity of visual working memory. However, we do not believe that the disappearance of the CDA that we observed during learning is simply due to this imbalance being washed out across trials in which the cues appear in both hemifields. Contrary to this account, we found that we when consistently presented the target cues in the same hemifield within a run of trials that the CDA to the target cues disappeared at the same rate. These findings suggest that the disappearance of the CDA across target repetitions is not due to the effect being washed out; however, we acknowledge the need for additional research to understand the nature of taking working memory representations offline during the learning paradigm we used in this work.

Second, we do not yet know the precise type of long-term memory representations that the anterior P1 effect (the P170) is measuring. The initial work with the P170 demonstrates that it provides an electrophysiological measure of long-term, perceptual priming that is highly related to the behavioral priming effect of subjects when shown new versus previously exposed stimuli (Voss et al., 2010). However, we do not know whether this ERP component is measuring an explicit representation in long-term memory or a process of accessing a long-term memory. A large body of work has focused on these questions with regard to other ERP indices of long-term memory (Paller et al., 2012; Rugg & Curran, 2007). However, given the recent discovery of the P170 we need to understand more about how this sensitive measure of memory maps onto the mechanisms of long-term memory that are foundational in models of memory. Studies of long-term memory that found effects like the P170 seem to suggest that it indexes the familiarity of an item that we encounter (Duarte, Ranganath, Winward, Hayward, & Knight, 2004; Diana et al., 2005; Friedman, 2004; Tsivilis et al., 2001).

Third, the initial behavioral experiments that we discussed in this paper have yielded some puzzling results. Specifically, Olivers (2009) showed that the effect of a distractor in a search array that matched a working memory representation was stronger when the target was the same across trials compared to one in which the search target changed from trial to trial. This could indicate that a working memory representation can more easily influence processing when perceptual attention is relying upon a target representation in long-term memory (consistent with the reward effects on the CDA described above); however, further work is necessary to understand how representations in different memory stores interact to influence the deployment of perceptual attention.

The line of work we discussed here offers the opportunity to understand the nature of top-down

attentional control with unprecedented precision. In many models of visual attention, there is a high degree of specification of how the bottom-up inputs are processed both in computational (e.g., Bundesen, 1990) and neural terms (Desimone & Duncan, 1995). However, the memory representations that control attention have been less specified in these models due to a lack of methods with which to study them (Woodman & Chun, 2006). Recent work has shown that this is a great need because the intersection of attention and memory is increasingly heavily trafficked (Awh, Belopolsky, & Theeuwes, 2012; Awh & Jonides, 2001; Chun, Golomb, & Turk-Browne, 2011; Logan, 2002; Olivers et al., 2011). Our hope is that the tools we have laid out here will provide much needed leverage with which to pry the cover off of the mechanisms that control selective visual processing.

Keywords: visual attention, visual working memory, long-term memory, event-related potentials

Acknowledgments

We would like to thank Jason Arita, Debbie Pardo, and Julianna Ianni who helped with the collection of the data described. Gordon Logan provided many useful discussions of this work. This research was made possible by grants from the National Eye Institute of NIH (RO1-EY019882 and P30-EY008126) and the National Science Foundation (BCS-0957072).

Commercial relationships: none.

Corresponding author: Geoffrey Woodman.

Email: geoffrey.f.woodman@vanderbilt.edu.

Address: Department of Psychology, Vanderbilt University, Nashville, Tennessee.

References

- Anderson, J. R. (1982). Acquisition of a cognitive skill. *Psychological Review*, *89*, 369–406.
- Anderson, J. R. (2000). *Learning and memory*. New York: Wiley.
- Arita, J. T., Carlisle, N., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 580–584.
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*, 437–443.
- Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*, 119–126.
- Balani, A. B., Soto, D., & Humphreys, G. W. (2010). Working memory and target-related distractor effects on visual search. *Memory & Cognition*, *38*, 1058–1076.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*, 523–547.
- Bundesen, C., & Habekost, T. (2008). *Principles of visual attention: Linking mind and brain*. New York: Oxford University Press.
- Bundesen, C., Habekost, T., & Kyllingsbaek, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review*, *112*, 291–328.
- Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional templates in visual working memory. *Journal of Neuroscience*, *31*, 9315–9322.
- Carlisle, N. B., & Woodman, G. F. (2011a). Automatic and strategic effects in the guidance of attention by working memory representations. *Acta Psychologica*, *137*, 217–225.
- Carlisle, N. B., & Woodman, G. F. (2011b). When memory is not enough: Electrophysiological evidence for goal-dependent use of working memory representations in guiding visual attention. *Journal of Cognitive Neuroscience*, *23*, 2650–2664.
- Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of Neurophysiology*, *80*, 2918–2940.
- Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, *363*, 345–347.
- Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (2001). Responses of neurons in macaque area V4 during memory-guided visual search. *Cerebral Cortex*, *11*, 761–772.
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*, 73–101.
- Dalvit, S., & Eimer, M. (2011). Memory-driven attentional capture is modulated by temporal task demands. *Visual Cognition*, *19*, 145–153.
- Danker, J. F., Hwang, G. M., Gauthier, L., Geller, A., Kahana, M. J., & Sekuler, R. (2008). Characterizing the ERP Old-New effect in a short-term memory task. *Psychophysiology*, *45*, 784–793.
- Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary reward. *Psychological Science*, *20*, 222–227.

- Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*, 778–784.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
- Diana, R., Vilberg, K., & Reder, L. (2005). Identifying the ERP correlate of a recognition memory search attempt. *Cognitive Brain Research*, *24*, 674–684.
- Dombrowe, I., Olivers, C. N. L., & Donk, M. (2010). The time course of working memory effects on visual attention. *Visual Cognition*, *18*, 1089–1112.
- Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, *11*, 467–473.
- Downing, P. E., & Dodds, C. M. (2004). Competition in visual working memory for control of search. *Visual Cognition*, *11*, 689–703.
- Duarte, A., Ranganath, C., Winward, L., Hayward, D., & Knight, R. (2004). Dissociable neural correlates for familiarity and recollection during the encoding and retrieval of pictures. *Cognitive Brain Research*, *18*, 255–272.
- Duncan, J. (1996). Cooperating brain systems in selective perception and action. In T. Inui & J. L. McClelland (Eds.), *Attention and performance XVI: Information integration in perception and communication* (pp. 549–578). Cambridge, MA: MIT Press.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458.
- Friedman, D. (2004). ERP studies of recognition memory: differential effects of familiarity, recollection, and episodic priming. *Cognitive Sciences*, *1*, 81–121.
- Hamker, F. H. (2004). A dynamic model of how feature cues guide spatial attention. *Vision Research*, *44*, 501–521.
- 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*, 1292–1302.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*, *30*, 11096–11103.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2011). Reward has a residual impact on target selection in visual search, but not on the suppression of distractors. *Visual Cognition*, *19*, 117–128.
- Hickey, C., Di Lollo, V., & McDonald, J. J. (2009). Electrophysiological indices of target and distractor processing in visual search. *Journal of Cognitive Neuroscience*, *21*, 760–775.
- 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*, 423–442.
- Ikkai, A., McCollough, A. W., & Vogel, E. K. (2010). Contralateral delay activity provides a neural measure of the number of representations in visual working memory. *Journal of Neurophysiology*, *103*, 1963–1968.
- James, W. (1890). *The principles of psychology*. New York: Holt.
- Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin & Review*, *9*(4), 637–671.
- Kiyonaga, A., Egner, T., & Soto, D. (2012). Cognitive control over working memory biases of selection. *Psychonomic Bulletin & Review*, *19*, 639–646.
- Kumar, S., Soto, D., & Humphreys, G. W. (2009). Electrophysiological evidence for attentional guidance by the contents of working memory. *European Journal of Neuroscience*, *30*, 307–317.
- Logan, G. D. (1978). Attention in character classification tasks: Evidence for the automaticity of component stages. *Journal of Experimental Psychology: General*, *107*, 32–63.
- Logan, G. D. (1979). On the use of a concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 189–207.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.
- Logan, G. D. (1990). Repetition priming and automaticity: Common underlying assumptions. *Cognitive Psychology*, *22*, 1–35.
- Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review*, *109*, 376–400.
- Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, *108*, 393–434.
- Luck, S. J. (2005). *An introduction to the event-related potential technique*. Cambridge, MA: MIT Press.
- Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences, USA*, *107*, 5232–5237.
- Neisser, U. (1963). Decision-time without reaction-

- time: Experiments in visual scanning. *American Journal of Psychology*, 76, 376–385.
- Neisser, U. (1964). Visual search. *Scientific American*, 210(6), 94–102.
- Nickerson, R. S. (1966). Response times with a memory-dependent decision task. *Journal of Experimental Psychology*, 72, 761–769.
- Oh, S. H., & Kim, M. S. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic Bulletin and Review*, 11(2), 275–281.
- 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, 1275–1291.
- Olivers, C. N. L., & Eimer, M. (2011). On the difference between working memory and attentional set. *Neuropsychologia*, 49, 1553–1558.
- Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory contents affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, 32, 1243–1265.
- Olivers, C. N. L., Peters, J. C., 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, 327–334.
- Paller, K. A., Lucas, H. D., & Voss, J. L. (2012). Assuming too much from ‘familiar’ brain potentials. *Trends in Cognitive Sciences*, 6, 313–315.
- Paller, K. A., Voss, J. L., & Boehm, S. G. (2007). Validating neural correlates of familiarity. *Trends in Cognitive Sciences*, 11, 243–250.
- Peters, J. C., Goebel, R., & Roelfsema, P. R. (2008). Remembered but unused: The Accessory items in working memory that do not guide attention. *Journal of Cognitive Neuroscience*, 21, 1081–1091.
- Peterson, L. R., & Gentile, A. (1965). Proactive interference as a function of time between tests. *Journal of Experimental Psychology*, 70, 473–478.
- Pillsbury, W. B. (1908). *Attention*. New York: Macmillan.
- Polich, J. (2012). Neuropsychology of P300. In S. J. Luck & E. Kappenman (Eds.), *Oxford handbook of event-related potential components* (pp. 159–188). New York: Oxford University Press.
- Raymond, J. E., & O’Brien, J. L. (2009). Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science*, 20, 981–988.
- Reinhart, R. M. G., & Woodman, G. F. (in press). High stakes trigger the use of multiple memories to enhance the control of attention. *Cerebral Cortex*.
- Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61, 168–185.
- Rickard, T. C. (1997). Bending the power law: A CMPL theory of strategy shifts and the automatization of cognitive skills. *Journal of Experimental Psychology: General*, 126, 288–311.
- Rugg, M. D., & Curran, T. (2007). Event-related potentials and recognition memory. *Trends in Cognitive Sciences*, 11, 251–257.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing. I: Detection, search and attention. *Psychology Review*, 84, 1–66.
- Serences, J. (2008). Value-based modulations in human visual cortex. *Neuron*, 60, 1169–1181.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing. II: Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127–190.
- 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.
- Soto, D., & Humphreys, G. W. (2007). Automatic guidance of visual attention from verbal working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 730–737.
- Soto, D., & Humphreys, G. W. (2008). Stressing the mind: The role of verbal suppression and cognitive load on the time course of attentional guidance from working memory. *Perception and Psychophysics*, 70(5), 924–934.
- 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, USA*, 104, 17186–17191.
- Theeuwes, J. (1993). Visual selective attention: A theoretical analysis. *Acta Psychologica*, 83, 93–154.
- Treisman, A. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, 40, 201–237.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Tsivilis, D., Otten, L., & Rugg, M. D. (2001). Context effects on the neural correlates of recognition

- memory: an electrophysiological study. *Neuron*, *31*, 497–505.
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*, 748–751.
- Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, *438*, 500–503.
- Voss, J. L., Schendan, H. E., & Paller, K. A. (2010). Finding meaning in novel geometric shapes influences electrophysiological correlates of repetition and dissociates perceptual and conceptual priming. *NeuroImage*, *49*, 2879–2889.
- Wickens, D. D., Born, D. G., & Allen, C. K. (1963). Proactive inhibition and item similarity in short-term memory. *Journal of Verbal Learning & Verbal Behavior*, *2*(5-6), 440–445.
- Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, *1*, 202–238.
- Wolfe, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13–73). London: University College London Press.
- Wolfe, J. M. (2003). Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Sciences*, *7*, 70–76.
- Wolfe, J. M. (2012). Saved by a log: How do humans perform hybrid visual and memory search? *Psychological Science*, *23*, 698–703.
- Woodman, G. F. (2010). A brief introduction to the use of event-related potentials (ERPs) in studies of perception and attention. *Attention, Perception & Psychophysics*, *72*(8), 2031–2046.
- Woodman, G. F., & Arita, J. T. (2011). Direct electrophysiological measurement of attentional templates in visual working memory. *Psychological Science*, *22*, 212–215.
- Woodman, G. F., & Chun, M. M. (2006). The role of working memory and long-term memory in visual search. *Visual Cognition*, *14*, 808–830.
- 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*, 363–377.
- Woodman, G. F., & Luck, S. J. (2009). Why is information displaced from visual working memory during visual search? *Visual Cognition*, *18*, 275–295.
- Woodman, G. F., Luck, S. J., & Schall, J. D. (2007). The role of working memory representations in the control of attention. *Cerebral Cortex*, *17*, i118–i124.
- Woodman, G. F., Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science*, *12*, 219–224.
- Yi, D.-J., Turk-Browne, N. B., Chun, M. M., & Johnson, M. K. (2008). When a thought equals a look: Refreshing enhances perceptual memory. *Journal of Cognitive Neuroscience*, *20*(8), 1371–1380.