

Attentional Templates in Visual Working Memory

Nancy B. Carlisle, Jason T. Arita, Deborah Pardo, and Geoffrey F. Woodman

Vanderbilt University, Vanderbilt Vision Research Center, Center for Integrative and Cognitive Neuroscience, Nashville, TN 37240-7817

Most theories of attention propose that we maintain attentional templates in visual working memory to control what information is selected. In the present study, we directly tested this proposal by measuring the contralateral-delay activity (CDA) of human event-related potentials during visual search tasks in which the target is cued on each trial. Here we show that the CDA can be used to measure the maintenance of attentional templates in visual working memory while processing complex visual scenes. In addition, this method allowed us to directly observe the shift from working memory to long-term memory representations controlling attention as learning occurred and experience accrued searching for the same target object. Our findings provide definitive support for several critical proposals made in theories of attention, learning, and automaticity.

Introduction

Theories of attention propose that we hold target representations, known as attentional templates, in working memory to provide top-down biasing signals to the neurons that perform perceptual processing (Desimone and Duncan, 1995; Bundesen et al., 2005). The neurophysiological evidence for this proposal comes from single-unit recordings in the inferotemporal cortex (IT) of monkeys performing search (Chelazzi et al., 1993, 1998). In those experiments, monkeys performed a search task in which the target object was cued at the beginning of each trial. Neurons in IT exhibited elevated firing rates between the cue and the search array, suggesting that the target representation was being maintained in working memory.

Despite the importance of attentional templates in working memory for explaining the top-down control of attention, we lack direct empirical support for this proposal from neuroscientific studies of humans. Moreover, the cued visual search tasks used to support the proposal of working memory attentional templates are unlike those typically used in studies of humans in which the identity of the targets is constant across all trials (Woodman et al., 2007). Given that attentional templates in working memory are the theoretical lynchpins that allow many models to account for top-down attention effects, it is critical that we determine under what conditions these target representations in working memory may exist.

In this study, we directly measured neural activity related to maintaining information in visual working memory by recording event-related potentials (ERPs) from humans performing visual search. Each trial began with the presentation of a cue array, followed by a search array. The task-relevant object in the cue

array, indicated by color (e.g., red for the first block of trials), signaled the shape of the target object that was to be searched for on that trial in the subsequent search array (e.g., a red square with a gap to the right) (see Fig. 1A). We focused on the ERP waveforms following the cue to determine whether the contralateral-delay activity (CDA) was present. The CDA is a component that indexes the active maintenance of object representations in visual working memory (Vogel and Machizawa, 2004; Vogel et al., 2005).

We predicted that if target-template representations are maintained in visual working memory to control perceptual attention during visual search, then we should see that the target cue elicits a CDA that continues until search is performed. Furthermore, we expected that when subjects searched for two targets, the amplitude of the CDA would be twice as large as when a single object was searched for (experiment 1). Next, we tested the prediction that long-term memory representations, and not working memory representations, enable performance when searching for a target whose identity is static across trials (experiment 2) (Logan, 1988; Woodman et al., 2007). We concluded by directly observing the handoff of the attentional template from working memory to long-term memory by measuring the CDA across short runs of trials in which subjects searched for the same target (experiment 3).

Materials and Methods

Participants. Volunteers (18–35 years of age) participated in the experiments in exchange for monetary compensation. Twenty subjects (11 female) were used across the two conditions of experiment 1, 10 subjects (5 female) in experiment 2, and 18 subjects (11 female) in experiment 3 to increase the signal-to-noise ratio for the trial-based analyses. In experiment 1, subjects participated in either the one-target or two-target condition. Informed consent was obtained before any experimental procedures began and all procedures were approved by the Vanderbilt Institutional Review Board. All participants had normal color vision, no history of neurological problems, and normal or corrected-to-normal visual acuity.

Stimuli. The participants viewed the stimuli at a distance of 114 cm on a homogeneous light gray background (54.3 cd/m^2). A black fixation cross ($<0.01 \text{ cd/m}^2$, $0.4^\circ \times 0.4^\circ$ of visual angle) was visible throughout

Received March 1, 2011; revised April 19, 2011; accepted May 5, 2011.

Author contributions: N.B.C., J.T.A., and G.F.W. designed research; N.B.C., J.T.A., D.P., and G.F.W. performed research; N.B.C. and G.F.W. analyzed data; G.F.W. wrote the paper.

This work was supported by National Science Foundation (BCS-0957072) and National Eye Institute (R01-EY019882). We thank Jaime Coffino for help with data collection while conducting experiment 3.

Correspondence should be addressed to Geoffrey F. Woodman, Department of Psychology, Vanderbilt University, PMB 407817, 2301 Vanderbilt Place, Nashville, TN 37240-7817. E-mail: geoffrey.f.woodman@vanderbilt.edu.

DOI:10.1523/JNEUROSCI.1097-11.2011

Copyright © 2011 the authors 0270-6474/11/319315-08\$15.00/0

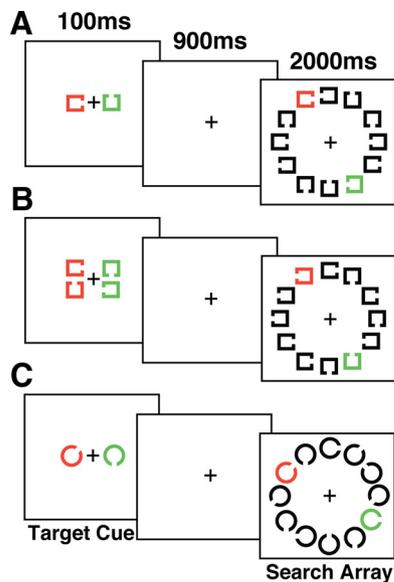


Figure 1. Example of the stimuli used in experiments 1–3. *A*, The stimulus sequences used in the one-target condition of experiments 1 and 2. The stimuli for experiment 2 were identical to the one-target condition of experiment 1 except that the cued target shape appearing on the left or right was the same on each trial for a given participant. *B*, Two possible targets were cued on each trial of in the two-target condition of experiment 1. *C*, The stimuli in experiment 3 were Landolt Cs with eight possible gap locations.

each trial. The two cue stimuli were outlined squares ($0.7^\circ \times 0.7^\circ$, 0.1° line thickness), one of which was green ($x = 0.281$, $y = 0.593$, 45.3 cd/m^2) and the other red ($x = 0.612$, $y = 0.333$, 15.1 cd/m^2). Each cue stimulus was presented 2.2° to the left or right of the center of the monitor. The visual search arrays were composed of outlined Landolt square stimuli ($0.7^\circ \times 0.7^\circ$, composed of line segments 0.1° thick) with a gap (0.5°) on the left, right, top, or bottom. Figure 1 shows that in all experiments, the search array contained one red and one green Landolt square stimulus and 10 black distractor stimuli (<0.01 cd/m^2) arranged similar to the number locations on a clock face (centered 4.4° from the middle of the monitor). In addition, the target shape could only appear in the task-relevant color. In experiment 3, the elements in the cue and search arrays were Landolt C stimuli (0.88° diameter, 0.13° thick, 0.22° gap width) of eight possible orientations (gap on the top, right, bottom, left, or 45° of rotation off of these cardinal directions) instead of the Landolt squares used in experiments 1 and 2. This allowed us to have a set of eight possible search elements that could be cued on each trial (i.e., a larger stimulus set).

Procedure. The beginning of each trial was indicated by the presentation of the fixation point 1200–1600 ms before the onset of the cue stimulus. Next, the cue stimuli were presented for 100 ms, beginning 1000 ms before the onset of the search array (i.e., a 1000 ms cue-to-search-array stimulus onset asynchrony). The search array was then presented for 2000 ms. The final stimulus event during each trial was the offset of the fixation point and search array after the 2000 ms presentation. These offsets served to indicate to the participants the beginning of the intertrial interval (duration of 1200–1600 ms, randomly jittered using a rectangular distribution) during which they were allowed to blink. During the intertrial interval, a reminder appeared on the screen indicating the currently task-relevant cue color (i.e., “attend red” or “attend green”).

In all experiments, a target was presented on half of the trials and matched the shape of the task-relevant cue (or one of the two task-relevant cues in the two-target condition of experiment 1). The cued target shape (i.e., a top, bottom, left, or right gap in experiments 1 and 2, or every 45° beginning at the vertex of the Landolt C in experiment 3), target presence (present or absent), and target location, when present, were randomly selected on each trial. In experiment 3, there was the additional factor of length of the same-target run (3, 5, or 7 trials). The

length of each run within each block of trials was randomly ordered, with equal numbers of each run length.

Participants were instructed to respond to the search array by pressing a button on a handheld gamepad (Logitech Precision). Subjects were instructed to respond as accurately and quickly as possible. In experiments 1 and 2, they used the index finger of their right hand to press a button indicating target present and the middle finger of their right hand to press a different button indicating a target-absent array. In experiment 3, subjects pressed two buttons with their right-hand thumb and were cued in the intertrial interval if there was going to be an upcoming change in target identity because we wanted to ensure that they noticed that a new target was being cued. A failure to perceive the cue on the change trial would make the memory-related ERP effects difficult to interpret.

Each participant performed 18–24 practice trials before the beginning of the first block of experimental trials to ensure they understood the task and could maintain proper fixation. In experiments 1 and 2, participants then performed four blocks of 192 trials while ERPs were recorded. Every 64 trials, the subjects were given a short break to allow them to rest. In experiment 3, two blocks of 360 trials were performed with short breaks approximately every 65 trials, constrained to occur after a run of trials with the same target. Across blocks of trials, observers switched between the red and green items being task relevant to rule out physical stimulus explanations for the lateralized ERP effects (Hillyard and Picton, 1987; Woodman, 2010).

ERP recording and analysis. The electroencephalogram (EEG) was recorded from tin (Sn) electrodes held on the scalp by an elastic cap (Electrocap International). A subset of the International 10/20 System sites were used (Fz, Cz, Pz, F3, F4, C3, C4, P3, P4, PO3, PO4, T3, T4, T5, T6, O1, and O2) as well as the nonstandard sites OL (halfway between O1 and T5) and OR (halfway between O2 and T6). The right mastoid electrode served as the online reference for these active electrode sites. The signals were re-referenced offline to the average of the left and the right mastoids (Nunez, 1981). The electrooculogram (EOG) was recorded by placing electrodes 1 cm lateral to the external canthi to measure horizontal eye movements and by placing an electrode beneath the left eye, referenced to the right mastoid, to measure vertical eye movements and blinks. The EEG and EOG were amplified by an SA Instrumentation amplifier with a gain of 20,000 and a bandpass of 0.01–100 Hz. The amplified signals were digitized at 250 Hz by a PC and averaged offline. Trials accompanied by incorrect behavioral responses or ocular or myogenic artifacts were excluded from the averages.

We used a two-step procedure for ocular artifact rejection, as described previously (Woodman and Luck, 2003). Briefly, trials with large eye movements were rejected before averaging, and averaged horizontal EOG (HEOG) waveforms were used to reject any subjects with significant unrejected eye movements. This procedure led to the rejection of an average of 16.2% of trials per participant (with a single participant maximum of 24.9%) and required us to replace four, four, and five participants in experiments 1, 2, and 3, respectively, due to excessive eye movements (either $>25\%$ of individual trials rejected or any residual systematic eye movement that resulted in HEOG voltage deflections >3.2 μV , corresponding to an ocular deviation of $\pm 0.1^\circ$).

CDA was measured as the difference between electrode sites contralateral versus ipsilateral to the location of the task-relevant cue shape in the cue array for each trial. On the basis of previous CDA experiments using change-detection tasks, we focused on the period 300–1000 ms after the onset of the target cue (Vogel and Machizawa, 2004; Vogel et al., 2005).

ANOVAs were used for all statistical tests, and p values were adjusted using the Greenhouse-Geisser epsilon correction for nonsphericity (Jennings and Wood, 1976). The data were collapsed across cued color because we did not observe any effects due to sensory differences between which color was task relevant. We first entered the data into an ANOVA with the following within-subjects factors: hemisphere (left vs right), contralaterality with respect to the cue (ipsilateral vs contralateral), and electrode site (PO3/4, O1/2, OL/R, vs T5/6). We then performed mixed-model ANOVAs between experiments using the within-subjects factors described above and the between-subjects factor of condition (one target vs two targets in experiment 1) or experiment (data from the one-target condition of experiment 1 vs experiment 2). For experiment 3, the ERP

One-Target Condition

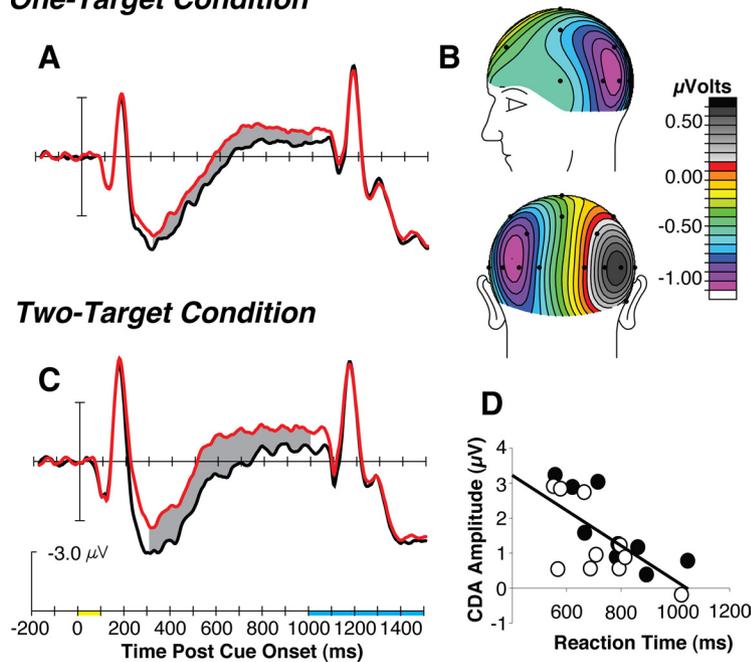


Figure 2. The ERP results of experiment 1. *A*, Grand average waveforms from lateral occipital-temporal electrodes contralateral (red) and ipsilateral (black) to the location of the cue on each trial of the one-target condition in experiment 1. The gray region shows the epoch (300–1000 ms postcue onset) in which a significant CDA component was measured. The shading on the timeline (in *C*, applies to *A* and *C*) indicates the duration of the cue (yellow) and the search array (blue). *B*, The voltage distribution (spherical spline interpolation) of the CDA during the measurement window for right visual field cues minus left visual field cues. *C*, The grand average waveforms from the two-target condition. *D*, The relationship between individual subject's CDA amplitude and behavioral RT on target-present (empty circles) and target-absent trials (filled circles).

waveforms were binned according to number of trials that occurred since a change in target identity (i.e., 1–2, 3–4, or 5–7 consecutive trials searching for that particular target shape). The last bin consisted of the last three serial positions in the repetition sequence because the longer same-target runs occurred less frequently because each run length was equally likely. This three-level factor of repetition bin was added to the ANOVA for experiment 3.

Results

Experiment 1

The goal of experiment 1 was to determine whether humans store attentional templates in visual working memory during a visual search task similar to that of Chelazzi and colleagues (1993), in which the search target is cued at the beginning of each trial (Fig. 1*A*). We also included a parametric manipulation of the number of possible targets to further test the hypothesis that the CDA could be used to directly measure the attentional templates in visual working memory.

One of the defining characteristics of the CDA component is that it increases in amplitude with each additional item that is stored in visual working memory, reaching an asymptotic amplitude at each individual's visual working memory capacity (Vogel and Machizawa, 2004). In experiment 1, we exploited this characteristic to further test the attentional template hypothesis. Specifically, if the CDA is a direct index of the visual working memory representations of the targets that drive attentional selection, then if we cue observers to search for two possible targets, the CDA amplitude should be twice what it is when people search for a single target. Thus, in experiment 1, we had half of our observers search for two possible targets in the two-target condition (Fig. 1*B*) and compared the amplitude of the CDA in this condition to that of the one-target condition (Fig. 1*A*).

Behavior

Subjects' response accuracy was not significantly different between target-present and target-absent trials in the one-target condition (97.7% vs 97.3% correct, respectively; $p > 0.30$) or two-target condition (92.4% vs 93.7% correct, respectively; $p = 0.30$). However, reaction times (RTs) were significantly faster when the target was present in the one-target condition (576.9 ms) compared with when it was absent (652.5 ms, $F_{(1,9)} = 81.56$, $p < 0.001$). Search RTs were also significantly faster on target-present trials in the two-target condition (718.8 ms) than target-absent trials (772.7 ms, $F_{(1,9)} = 19.09$, $p < 0.01$). This replicates a pattern observed across decades of behavioral studies of visual search performance in which target-present responses are consistently faster than target-absent responses (Neisser, 1964; Treisman and Gelade, 1980; Wolfe, 1998). As expected based on previous research (Sperling and Melchner, 1978; Menneer et al., 2009), RTs in the two-target condition were significantly slower than those in the one-target condition (614.7 vs 745.8 ms, for the one and two-target conditions, respectively; $F_{(1,18)} = 7.48$, $p < 0.05$) despite the search arrays being identical in both conditions. The effect of target presence was significant in

this mix-model ANOVA ($F_{(1,18)} = 75.42$, $p < 0.001$), but this factor did not interact with number of targets (i.e., one vs two, $p > 0.15$).

ERPs

Figure 2 shows the waveforms recorded from a pair of lateral electrodes (T5/6) relative to the location of the task-relevant cue or cues with the measurement window for the CDA shaded in gray. In the one-target condition, cues elicited a clear CDA during the measurement window (mean contralateral-minus-ipsilateral amplitude = -0.8μ V) (Fig. 2*A*) and the scalp distribution of the CDA was similar to previous studies of this component (Fig. 2*B*) (McCollough et al., 2007). Of primary importance for the prediction that a CDA should be elicited by the cue indicating the shape of the visual search target, within the one-target condition we found a significant main effect of contralaterality of the electrode sites relative to the location of the cue during the 300–1000 ms postcue epoch ($F_{(1,9)} = 28.94$, $p < 0.001$). We also found main effects of electrode site ($F_{(3,27)} = 28.94$, $p < 0.001$), due to the parietal sites being generally more positive during the measurement window (i.e., the CDA temporally overlaps with the P3 component elicited by the cue stimulus) and of hemisphere ($F_{(1,9)} = 6.54$, $p < 0.05$), because the P3 that is also elicited during the CDA measurement window was largest over the right hemisphere. The only interaction we found was contralaterality \times electrode site due to the CDA being largest at occipital-parietal and temporal electrode sites ($F_{(3,27)} = 10.06$, $p < 0.0001$) due to the anticipated distribution of the CDA (Fig. 2*B*).

In the two-target condition, we measured a CDA that was essentially twice the amplitude of that found in the one-target condition (i.e., -1.5μ V in the two-target condition compared

with $-0.8 \mu\text{V}$ in the one-target condition). These two-target cues elicited a CDA that resulted in significant main effects of contralaterality ($F_{(1,9)} = 24.90, p < 0.001$), electrode site ($F_{(3,27)} = 4.17, p < 0.05$), and a significant interaction of contralaterality \times electrode site ($F_{(3,27)} = 7.65, p < 0.001$) due to the expected scalp distribution of the CDA across electrode sites.

In the next analysis, we entered the voltage data from the one- and two-target conditions into a mixed-model ANOVA. As expected based on the amplitude effects we observed (compare Fig. 2, A and C), this yielded a significant main effect of condition ($F_{(1,18)} = 4.58, p < 0.05$) and an interaction of condition \times contralaterality ($F_{(1,18)} = 5.08, p < 0.05$). The ANOVA also yielded main effects of hemisphere ($F_{(1,18)} = 6.72, p < 0.05$), contralaterality ($F_{(1,18)} = 46.15, p < 0.001$), electrode site ($F_{(3,54)} = 8.15, p < 0.001$), and interactions of hemisphere \times contralaterality ($F_{(1,18)} = 4.96, p < 0.05$) and contralaterality \times electrode site ($F_{(3,54)} = 14.36, p < 0.001$), largely due to the CDA riding on top of P3 components that differed between conditions. These results provide statistical support for the observation that the CDA measured in the two-target condition was essentially twice as large as that found in the one-target condition. These findings support the hypothesis that representations of each searched-for target object are maintained in working memory when the targets changed on each trial.

Finally, we thought that if the CDA was measuring the representations that were being used to guide attention, then we should be able to find a relationship between an individual subject's CDA amplitude and the efficiency of their subsequent visual search. We focused on the CDA and the RT effects in the two-target condition in experiment 1 because this was where the variability of both the electrophysiological and behavioral measures was greatest. As Figure 2D shows, we found that the amplitude of an individual subject's CDA across both target-present and target-absent trials was significantly correlated with the speed of their correct search responses ($r = -0.68$, or $r^2 = 0.46, p < 0.05$). That is, a large amplitude CDA following the cue could be used to predict that the subject's subsequent search response was likely to be fast. This provides another piece of evidence that the CDA is measuring the representations that are being used to control attention in the search tasks of experiment 1.

In summary, the findings from experiment 1 conformed to what we predicted if we maintain attentional templates in visual working memory while operating in this dynamic task environment in which the targets of our visual search change from trial-to-trial. Our parametric manipulation of the number of targets on each trial resulted in a systematic increase in the amplitude of the CDA. In the two-target condition, there was sufficient variability in the search RTs and CDA amplitude to observe that the amplitude of the CDA following the cue predicted the efficiency of the subject's search. Thus, our findings supported the hypothesis that representations of the searched-for target objects are maintained in visual working memory on each trial.

Experiment 2

Our findings from experiment 1 show that attentional templates are maintained in visual working memory when participants perform visual search in a task environment in which the searched-for targets change from moment-to-moment. However, in most studies of visual search with human subjects, the searched-for target remains the same across all of the trials in the experiment with few studies using a trial-by-trial cuing method similar to single-unit studies with monkeys (Chelazzi et al., 1993) and the present work (Wolfe et al., 2004; Vickery et al., 2005; Woodman

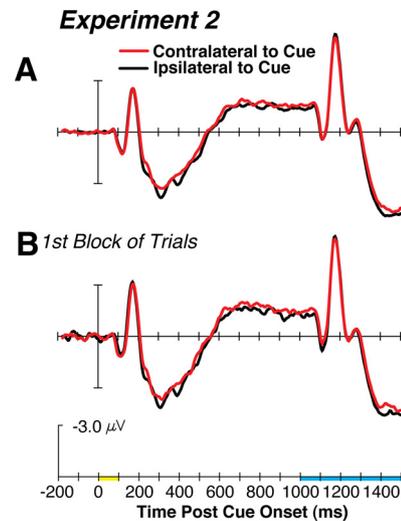


Figure 3. A, The grand average ERP results of experiment 2 using the same conventions defined in Figure 2. B, Waveforms averaged across subjects from just the first block of trials.

et al., 2007). Theories of learning and automaticity suggest that under such stable task demands, long-term memory representations and not working memory representations should come to drive the attentional selection of task-relevant items in the visual field (Logan, 1988, 2002; Anderson, 2000, 2007; Poldrack and Gabrieli, 2001). We sought to determine the generality of our findings from experiment 1 by using a more typical laboratory version of visual search in which human subjects searched for the same target object across trials. We predicted that under these conditions, the CDA elicited by the target cue should disappear, as long-term memory representations are used to control the deployment of attention to task-relevant objects.

Behavior

Search behavior was significantly more accurate on target-absent than target-present trials (98.0% vs 97.1% correct, respectively; $F_{(1,9)} = 11.42, p < 0.01$), although performance was near ceiling levels for both trial types. Search RTs were faster when the search target was present (572.0 ms) than when it was absent (610.3 ms, $F_{(1,9)} = 24.87, p < 0.001$). The stimuli in experiment 2 were identical to those used in the one-target condition of experiment 1, with the exception that the same target shape was cued on each trial of experiment 2 for a given subject. Although RTs were slower in the one-target condition of experiment 1 (mean = 614.7 ms) than in experiment 2 (mean = 591.2 ms), this difference was not significant ($F < 1.0$). In addition, accuracy in experiment 2 and the one-target condition of experiment 1 did not significantly differ (97.5% and 97.5% correct, respectively; $F < 1.0$). These findings support the conclusion that the search tasks in experiment 2 and the one-target condition of experiment 1 were both performed quickly and accurately. However, the ERP findings showed that the targets were represented in qualitatively different ways in experiments 1 and 2.

ERPs

The waveforms time locked to the target cues are shown in Figure 3A. The statistical analyses of the ERP data supported our observation that no significant CDA was elicited by the constant-target cues in experiment 2. Specifically, the contralaterality factor did not yield any significant effects or interactions ($F_s < 1.0$). Consistent with our observation that qualitatively different patterns of contralateral activity occurred in experiment 2 compared with

the one-target condition of experiment 1, the ANOVA of the ERP data yielded significant effects of experiment ($F_{(1,18)} = 4.25, p = 0.05$) and contralaterality ($F_{(1,18)} = 20.27, p < 0.001$), a significant interaction of these factors ($F_{(1,18)} = 12.02, p < 0.01$), as well as a significant three-way interaction of experiment \times contralaterality \times electrode site ($F_{(3,54)} = 2.91, p < 0.05$). This analysis also yielded a significant effect of electrode site ($F_{(3,54)} = 5.27, p < 0.01$) and significant interactions of contralaterality \times electrode site ($F_{(3,54)} = 14.20, p < 0.001$) and experiment \times hemisphere \times contralaterality \times electrode site ($F_{(3,54)} = 4.69, p < 0.01$). These effects were driven by the absence of the CDA following the cue in experiment 2.

To examine whether we could initially observe a CDA that then disappeared as subjects were given further practice searching for the same target, we also confined an analysis to the measurements of the CDA on the first block of trials relative to the remaining three blocks. Figure 3B shows that even during the first block of trials the CDA was absent. This observation is consistent with the power law of learning (Newell and Rosenbloom, 1981), which states that most of the learning occurs within the first 10–20 trials after a task begins, and this results in the reliance on long-term memory representations over those in working memory (Logan, 1988, 2002). In experiment 2, our subjects were given 24 practice trials to familiarize themselves with the task before ERP data collection began and it seems possible that the reliance on working memory templates was past even after this moderate amount of practice. In experiment 3, we tested the hypothesis that the handoff of the attentional template from working memory to long-term memory occurs fairly quickly during visual search (i.e., within the first 5–10 trials), as one might predict based on the RT functions from studies of the effects of task practice (Kolers, 1976; Logan, 1988).

Experiment 3

In experiment 3, we tested the hypothesis that attentional templates are initially represented in working memory, but then handed off to long-term memory with continued search for the same target, by having subjects search for the same target across runs of trials. That is, subjects searched for a given Landolt C stimulus (Fig. 1C) across consecutive trials with the randomly selected length of the run being 3, 5, or 7 trials long. We predicted that if the attention mechanisms shift between being driven by working memory representations of the target to long-term memory attentional templates, then the CDA should rapidly decrease across consecutive trials searching for the same object, perhaps mirroring the rapid decrease in the search RT functions.

Behavior

Search accuracy did not significantly differ between trials in which the target shape was present (96.9% correct) versus absent (97.3% correct, $p = 0.36$), and was similar for each repetition value ($p = 0.17$). Search RTs were faster when the search target was present (668.6 ms) than when it was absent (707.4 ms; $F_{(1,17)} = 10.83, p < 0.01$). As can be seen in Table 1, search RTs decreased as the number of trials searching for the same target increased, with the exception of the final target repetition. This led to a significant main effect of repetition on search RT ($F_{(6,102)} = 3.76, p < 0.01$).

ERPs

Figure 4, A–C, shows the waveforms recorded relative to the location of the task-relevant cue with the measurement window for the CDA shaded in gray. The CDA for the first trial with a new

Table 1. CDA amplitude and search RT for each repetition value

Target repetition	CDA amplitude (μV)	RT (ms)
1	−0.80	701
2	−0.55	696
3	−0.53	683
4	−0.46	682
5	−0.21	678
6	−0.24	676
7	−0.44	685

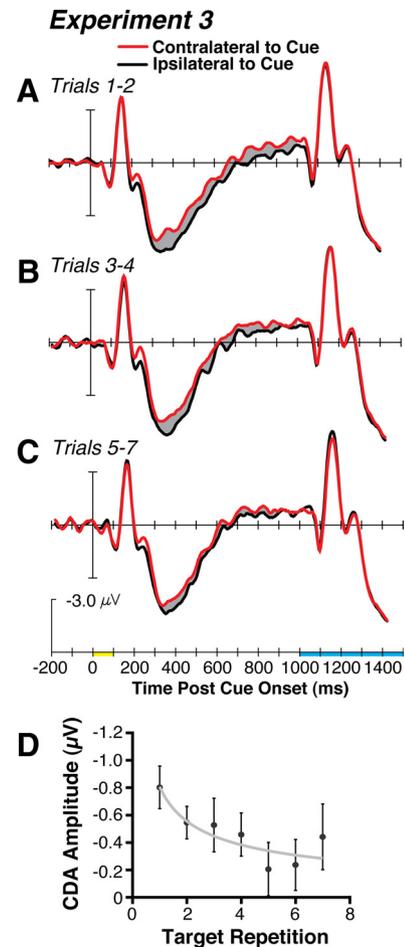


Figure 4. A–C, The grand average ERP results from the three trial-wise bins of experiment 3 using the same conventions defined in Figure 2. D, Plot of the CDA amplitude across consecutive trials with the same search target. The gray line shows the power-function fit and the error bars represent ± 1 SEM.

target was similar for experiment 3 (mean = $-0.80 \mu\text{V}$) and for the one-target condition of experiment 1 (mean = $-0.80 \mu\text{V}$), replicating the size of the effect shown immediately after a target-identity switch. By combining the ERP data from two or three consecutive trials, we created repetition bins to increase statistical power across these same-target runs, allowing us to measure the CDA as experience accrued searching for a given item. The CDA decreased across the repetition bins of one to two, three to four, and five to seven after a target identity switch (mean = $-0.66 \mu\text{V}$, mean = $-0.50 \mu\text{V}$, and mean = $-0.30 \mu\text{V}$, respectively). This decrease across repetition led to a significant interaction of contralaterality and repetition bin ($F_{(2,34)} = 3.83, p < 0.05$). In addition to this critical interaction, we also found a significant main effects of repetition bin ($F_{(2,34)} = 4.04, p < 0.05$), contralaterality ($F_{(1,17)} = 14.79, p < 0.01$), and a significant interaction of con-

trality \times electrode ($F_{(3,51)} = 16.25, p < 0.0001$) due to the same distribution of the CDA observed in the previous experiments.

One potential alternative explanation for these results is that the participants stopped attending to the cues, and would therefore not show a CDA to the location of the cue item. To rule out this possibility, we measured the N2pc, a negative-going potential contralateral to the attended hemifield. If participants stopped attending to the cue, we would expect the N2pc to decrease across repetition bins. To assess this alternative explanation, we conducted an ANOVA with the within-subjects factors of repetition bin, hemisphere (left vs right), contralaterality with respect to the cue (ipsilateral vs contralateral), and electrode site (O1/2 and OL/R, vs T5/6) from 200 to 275 ms after cue onset. We found a significant N2pc to the relevant cue, as indicated by a significant main effect of contralaterality ($F_{(1,17)} = 8.32, p < 0.05$). However, there was no interaction of repetition bin \times contralaterality ($p = 0.21$), indicating that participants were consistently attending to the cue across repetitions. The analysis also yielded an interaction of hemifield \times electrode ($F_{(2,34)} = 4.26, p < 0.05$) and a higher-order interaction of repetition bin \times hemifield \times contralaterality \times electrode ($F_{(4,68)} = 2.59; p < 0.05$).

As another way of quantifying the relationship between the number of trials searching for the same target and the amplitude of the CDA, we plotted the trial-by-trial CDA amplitude data in Figure 4D and fit it with a power function. This shows that the CDA rapidly decreased in amplitude with experience searching for the same target, as predicted by theories of learning and automaticity. We performed planned comparisons to determine which repetition values led to a CDA amplitude that was significantly different from zero. The CDA was significantly different from zero for repetitions 1, 2, 3, and 4 ($p < 0.0001, p < 0.001, p < 0.05$, and $p < 0.01$, respectively) but not for repetitions 5, 6, and 7 ($p = 0.31, p = 0.22$, and $p = 0.08$, respectively). The rapid decrease in the amplitude of the working-memory-related CDA was particularly surprising given the previous single-unit work on this topic. The recordings from monkey IT cortex during a similar cued-search task in which the target changed from trial-to-trial or was the same across 10–30 trial blocks showed little or no difference across these cuing conditions in the strength of the firing-rate modulation during the interval between the target cues and the search arrays (Chelazzi et al., 1998). We speculate that humans may have a more rapid transition between using working memory and long-term memory representations to guide attention. This might underlie the differences between our findings and the previous single-unit work, as well as learning rates in general between humans and nonhuman primates.

Next, it is noteworthy that the specific target shapes were repeatedly presented during the course of the experiment. We wondered if the averaged CDA amplitudes measured above might be due to large amplitude activity on the first encounters with a specific target, which became progressively smaller on the subsequent runs of trials of searching for that item. To test for this, we performed an analysis of the CDA for the first run of trials searching for each target (collapsing across all trial repetitions during the run) and compared the CDA amplitude measured to the second, third, fourth, and fifth or higher runs. This analysis showed overall CDA amplitudes across the target object repetitions that were neither significantly different nor showed a clear linear trend (CDA amplitude = 0.26, 0.40, 0.61, 0.31, 0.39, for first, second, third, fourth, and fifth or higher runs, respectively, with the critical interaction of target object repetition \times contralaterality resulting in $p = 0.12$). These findings do not show the

pattern that we would expect if the CDA was getting smaller and smaller across runs of trials as subjects searched for the same target objects. Subsequent work will be needed to determine whether these findings are due to the varied mapping between the target and distractor objects across runs of trials (Schneider and Shiffrin, 1977) or because the transition from a working memory template to a long-term memory template occurs anew each time a new search target is designated.

Finally, we assessed whether the reduction in CDA amplitude mirrored the reduction in search RT. Table 1 shows that we found a strong correlation in the mean CDA amplitude and the mean search RT ($r = 0.90$ or $r^2 = 0.81, p < 0.0001$). This finding conformed to our predictions based on the idea that both measures are influenced by the increased processing efficiency as the attentional template is handed off from the more cognitively demanding working memory representation to long-term memory (Logan, 1988; Anderson, 2000). Thus, by measuring the CDA index of the attentional template in visual working memory across trials, we were able to test a fundamental prediction of theories of attention and memory regarding the shifting locus of the memory representations that control processing during a task like visual search.

Discussion

Our study demonstrates how the CDA component of subjects' ERP waveforms can be used to directly measure the target representations in visual working memory that guide attention to targets in cluttered visual scenes. When subjects searched for two possible target objects, the amplitude of the CDA was approximately twice as large as when they searched for a single target object. The CDA disappeared when subjects searched for the same possible target during an entire experimental session, consistent with theories of learning and automaticity that propose that search performance will only depend upon visual working memory representations for a short period (i.e., <30 trials) after the identity of a new search target is introduced (Schneider and Shiffrin, 1977; Logan, 1988, 2002; Anderson, 2000). Finally, we showed how the amplitude of the CDA could be used to directly track the handoff of the attentional template from visual working memory to long-term memory as subjects searched for the same target object across runs of trials. Our findings not only establish a general method for accessing how working memory representations are used to control an individual's attentional set for the features of task-relevant objects, but provide new and direct tests of several key aspects of theories of attention and memory that have only been inferred indirectly until now.

The idea that working memory representations drive attentional selection is a very old one. For example, James (1890) noted that the connection between attention and memory was a common topic of previous authors (p. 427). He discussed the idea that the neurons that implement attentional selection are sensitive to bottom-up influences as well as top-down inputs from within the brains of the observers, making voluntary attentional selection possible. "While the object excites it [attention] from without, other brain-cells . . . arouse it from within" (p. 441). Current theories of attention now explicitly propose that working memory representations are the source of this top-down attentional control, as foreshadowed more than a century ago. For example, previous theories (Duncan and Humphreys, 1989; Bundesen, 1990; Desimone and Duncan, 1995; Desimone, 1996; Duncan, 1996; Bundesen et al., 2005) proposed that attentional template representations stored in working memory provide the top-down inputs to visual areas of cortex that implement atten-

tional selection of task-relevant stimuli (e.g., V4 and IT). These top-down biases allow us to focus attention on the task-relevant items in our visual field and not just any input with strong bottom-up input (i.e., a particularly bright or fast moving item). Thus, the idea that the currently active working memory representations determine what is processed by perceptual mechanisms of attentional selection is as old as cognitive science and no less important for neuroscientific theories of attention today.

Findings from neurophysiological studies have been the primary evidence in favor of the proposal that working memory representations play a critical role in the control of attention by establishing a specific attentional set. However, before the present study, only single-unit recordings were able to infer what was maintained in monkeys' working memory based on the delay period activity of neurons (Chelazzi et al., 1993, 1998; Desimone, 1996). Moreover, evidence for template-like neural activity has not been consistently observed (Chelazzi et al., 2001; Kusunoki et al., 2009). For example, no such template-like activity was found following target cues in the activity of V4 cells during the same memory-guided visual search task used during the recordings from IT that launched the attentional template idea (Chelazzi et al., 2001). The electrophysiological methods we used here allow us to generalize to the human brain the type of delay-period activity that Chelazzi and colleagues (1998) observed (see also Woodman and Arita, 2011). In addition, the electrophysiological index of the CDA allowed us to measure when attention was controlled by working memory representations and when that control was turned over to long-term memory representations.

Studies of learning and automaticity have long shown that behavioral performance of a task across practice is fairly well described by a power function (Newell and Rosenbloom, 1981). That is, as subjects perform visual search for the same target trial after trial, their RT speeds up considerably across the first few trials, but then the benefit of further practice rapidly diminishes (Schneider and Shiffrin, 1977; Shiffrin and Schneider, 1977; Logan, 1979). This rapid improvement followed by a slow march toward asymptotic behavioral performance led theorists to propose that task performance is initially supported by working memory representations, with later performance being a result of the necessary representations that control processing being stored in long-term memory (Logan, 1988, 2002; Anderson, 2000, 2007; Bundesen and Habekost, 2008). Evidence from functional magnetic resonance imaging (Poldrack, 2000; Poldrack and Gabrieli, 2001; Cole and Schneider, 2007), lesion studies with monkeys (Rossi et al., 2007), and dual-task experiments (Woodman et al., 2007) have been generally consistent with this proposal, but lacked the ability to isolate and track the specific contributions of visual working memory to the trial-by-trial performance during a task like visual search. For example, previous research using dual-task experiments and measuring the slope of the functions relating search RT to set size with and without a concurrent visual working memory task suggested that visual working memory representations of objects played a minimal role in visual search tasks (Logan, 1979; Woodman et al., 2001). However, these previous experiments had participants search for the same target trial after trial and the present findings show that a lack of interference would not be surprising given that long-term memory was performing the job of guiding attention to the target objects, thus minimally interacting with the concurrent visual working memory load (Woodman et al., 2007; Woodman and Luck, 2010). The methods we developed here allowed us to measure, with unprecedented precision, when long-term memory representations took over the guidance of attention from the

visual working memory representations as experience accrued searching for the same target object. Our estimate that visual working memory representations are only exclusively used to guide attention during the first one or two trials of a visual search task provides a level of detail that will greatly constrain biologically plausible models of cognitive control, attention, and memory.

Finally, the methods we have developed here for measuring the use of visual working memory representations in the performance of a task could be applied widely to understand the nature of cognitive control during a variety of other tasks and paradigms. We believe that visual search tasks were a natural place to begin because the findings allow us to test specific theories of attention; however, it is fairly easy to adapt our methods to other laboratory paradigms to understand the nature of memory representations underlying task performance.

References

- Anderson JR (2000) Learning and memory. New York: Wiley.
- Anderson JR (2007) How can the human mind occur in the physical universe? New York: Oxford UP.
- Bundesden C (1990) A theory of visual attention. *Psychol Rev* 97:523–547.
- Bundesden C, Habekost T (2008) Principles of visual attention: linking mind and brain. New York: Oxford UP.
- Bundesden C, Habekost T, Kyllingsbaek S (2005) A neural theory of visual attention: bridging cognition and neurophysiology. *Psychol Rev* 112:291–328.
- Chelazzi L, Miller EK, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex. *Nature* 363:345–347.
- Chelazzi L, Duncan J, Miller EK, Desimone R (1998) Responses of neurons in inferior temporal cortex during memory-guided visual search. *J Neurophysiol* 80:2918–2940.
- Chelazzi L, Miller EK, Duncan J, Desimone R (2001) Responses of neurons in macaque area V4 during memory-guided visual search. *Cereb Cortex* 11:761–772.
- Cole MW, Schneider W (2007) The cognitive control network: integrated cortical regions with dissociable functions. *Neuroimage* 37:343–360.
- Desimone R (1996) Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci U S A* 93:13494–13499.
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Ann Rev Neurosci* 18:193–222.
- Duncan J (1996) Cooperating brain systems in selective perception and action. In: Attention and performance XVI: information integration in perception and communication. (Inui T, McClelland JL, eds), pp 549–578. Cambridge, MA: MIT.
- Duncan J, Humphreys GW (1989) Visual search and stimulus similarity. *Psychol Rev* 96:433–458.
- Hillyard SA, Picton TW (1987) Electrophysiology of cognition. In: Handbook of physiology, section 1: the nervous system, volume 5: higher functions of the brain, part 2 (Plum F, ed), pp 519–584. Bethesda, MD: Waverly.
- James W (1890) The principles of psychology. New York: Holt.
- Jennings JR, Wood CC (1976) The e-adjustment procedure for repeated-measures analyses of variance. *Psychophysiology* 13:277–278.
- Kolers PA (1976) Reading a year later. *J Exp Psychol Hum Learn Mem* 2:554–565.
- Kusunoki M, Sigala N, Gaffan D, Duncan J (2009) Detection of fixed and variable targets in the monkey prefrontal cortex. *Cereb Cortex* 19:2522–2534.
- Logan GD (1979) On the use of a concurrent memory load to measure attention and automaticity. *J Exp Psychol Hum Percept Perform* 5:189–207.
- Logan GD (1988) Toward an instance theory of automatization. *Psychol Rev* 95:492–527.
- Logan GD (2002) An instance theory of attention and memory. *Psychol Rev* 109:376–400.
- McCullough AW, Machizawa MG, Vogel EK (2007) Electrophysiological measures of maintaining representations in visual working memory. *Cortex* 43:77–94.
- Menneer T, Cave KR, Donnelly N (2009) The cost of searching for multiple

- targets: effects of practice and target similarity. *J Exp Psychol Appl* 15:125–139.
- Neisser U (1964) Visual search. *Sci Am* 210:94–102.
- Newell A, Rosenbloom PS (1981) Mechanisms of skill acquisition and the law of practice. In: *Cognitive skills and their acquisition* (Anderson JR, ed), pp 1–55. Hillsdale, NJ: Erlbaum.
- Nunez PL (1981) *Electric fields of the brain*. New York: Oxford UP.
- Poldrack RA (2000) Imaging brain plasticity: conceptual and methodological issues—a theoretical review. *Neuroimage* 12:1–13.
- Poldrack RA, Gabrieli JD (2001) Characterizing the neural mechanisms of skill learning and repetition priming: evidence from mirror reading. *Brain* 124:67–82.
- Rossi AF, Bichot NP, Desimone R, Ungerleider LG (2007) Top down attentional deficits in macaques with lesions of lateral prefrontal cortex. *J Neurosci* 27:11306–11314.
- Schneider W, Shiffrin RM (1977) Controlled and automatic human information processing. I. Detection, search and attention. *Psychol Rev* 84:1–66.
- Shiffrin RM, Schneider W (1977) Controlled and automatic human information processing. II. Perceptual learning, automatic attending, and a general theory. *Psychol Rev* 84:127–190.
- Sperling G, Melchner MJ (1978) The attention operating characteristic: examples from visual search. *Science* 202:315–318.
- Treisman AM, Gelade G (1980) A feature-integration theory of attention. *Cogn Psychol* 12:97–136.
- Vickery TJ, King LW, Jiang Y (2005) Setting up the target template in visual search. *J Vis* 5:81–92.
- Vogel EK, Machizawa MG (2004) Neural activity predicts individual differences in visual working memory capacity. *Nature* 428:748–751.
- Vogel EK, McCollough AW, Machizawa MG (2005) Neural measures reveal individual differences in controlling access to working memory. *Nature* 438:500–503.
- Wolfe JM (1998) Visual search. In: *Attention* (Pashler H, ed), pp 13–73. London, UK: University College London.
- Wolfe JM, Horowitz TS, Kenner N, Hyle M, Vasan N (2004) How fast can you change your mind? The speed of top-down guidance in visual search. *Vision Res* 44:1411–1426.
- Woodman GF (2010) A brief introduction to the use of event-related potentials (ERPs) in studies of perception and attention. *Atten Percept Psychophys* 72:2031–2046.
- Woodman GF, Arita JT (2011) Direct electrophysiological measurement of attentional templates in visual working memory. *Psychol Sci* 22:212–215.
- Woodman GF, Luck SJ (2003) Serial deployment of attention during visual search. *J Exp Psychol Hum Percept Perform* 29:121–138.
- Woodman GF, Luck SJ (2010) Why is information displaced from visual working memory during visual search? *Vis Cogn* 18:275–295.
- Woodman GF, Vogel EK, Luck SJ (2001) Visual search remains efficient when visual working memory is full. *Psychol Sci* 12:219–224.
- Woodman GF, Luck SJ, Schall JD (2007) The role of working memory representations in the control of attention. *Cereb Cortex* 17:i118–i124.