

Searching while loaded: Visual working memory does not interfere with hybrid search efficiency but hybrid search uses working memory capacity

Trafton Drew¹ · Sage E. P. Boettcher² · Jeremy M. Wolfe^{3,4}

© Psychonomic Society, Inc. 2015

Abstract In “hybrid search” tasks, such as finding items on a grocery list, one must search the scene for targets while also searching the list in memory. How is the representation of a visual item compared with the representations of items in the memory set? Predominant theories would propose a role for visual working memory (VWM) either as the site of the comparison or as a conduit between visual and memory systems. In seven experiments, we loaded VWM in different ways and found little or no effect on hybrid search performance. However, the presence of a hybrid search task did reduce the measured capacity of VWM by a constant amount regardless of the size of the memory or visual sets. These data are broadly consistent with an account in which VWM must dedicate a fixed amount of its capacity to passing visual representations to long-term memory for comparison to the items in the memory set. The data cast doubt on models in which the search template resides in VWM or where memory set item representations are moved from LTM through VWM to earlier areas for comparison to visual items.

Keywords Visual search · Working memory · Dual-task performance

Imagine that you are in a grocery store, searching for the items on your memorized shopping list. With luck, the list resides

fairly stable in your long-term memory. Your shopping task is a “hybrid search”, combining visual and memory search (Schneider & Shiffrin, 1977; Wolfe, 2012). In the midst of your search, you meet someone and exchange phone numbers but you don’t have a pen, so you must hold the number in Working Memory while continuing your search. Does that Working Memory load interfere with your ongoing hybrid search and if so, what aspect of the search is perturbed? We know that working memory does interact with visual search. Two influential papers demonstrated strong evidence that under some circumstances a VWM load increases the slope of the reaction time (RT) x set size function, indicating that a VWM load reduces search efficiency (Oh & Kim, 2004; Woodman & Luck, 2004). The implication is that the act of holding information in VWM slows the rate with which we evaluate potential targets. This suggests that VWM plays a vital role in our ability to determine whether an object is a target or a distractor. More generally, it supports the idea, outlined in a number of important models of visual attention (Bundesen, 1990; Desimone & Duncan, 1995; Logan & Gordon, 2001; Miller & Cohen, 2001; Wolfe, Cave, & Franzel, 1989), that working memory plays an important part in the ability to effectively deploy visual attention. The most common notion is that the ‘template’ for visual search, the representation of the target, resides in VWM. This idea is supported by a growing line of research that demonstrates that attention tends to be automatically drawn to information being actively held in VWM (Downing & Dodds, 2004; Soto, Hodsoll, Rotshtein, & Humphreys, 2008). These results seem to predict that loading VWM would disrupt hybrid visual and memory search.

Even basic visual search requires memory to specify the current target. The target can be thought of as an attentional “template” (Olivers, Peters, Houtkamp, & Roelfsema, 2011) or “set” (Wolfe, 1994) that must be maintained in some

✉ Trafton Drew
trafton.drew@psych.utah.edu

¹ University of Utah, Salt Lake City, UT, USA

² Goethe Universitat Frankfurt, Frankfurt, Germany

³ Brigham & Women’s Hospital, Boston, MA, USA

⁴ Harvard Medical School, Boston, MA, USA

durable memory format while the search is being performed. A deluge of recent research has implicated WM as the mechanism that holds the attentional template in an activated state so that we recognize when an item matches the template. For instance, Lavie and colleagues have shown that observers are worse at ignoring distractors under WM load (Lavie, 2005). Olivers et al. demonstrated that visual search is reliably slowed if an item that is being held in working memory also appears as an irrelevant distractor in the search array (Olivers, Meijer, & Theeuwes, 2006). Monkey neurophysiology appears to further strengthen this view. The firing rates of neurons that are sensitive to a target's features increase while that specific object is the target of a search task (Chelazzi, Duncan, Miller, & Desimone, 1998). Finally, search efficiency decreases in the face of certain types of WM loads (Oh & Kim, 2004; Woodman & Luck, 2004).

Interestingly, Woodman and Luck found that WM for non-spatial information interfered with Visual Search but only when the target of search varied from trial to trial (2007). When the target template did not change, there was no effect of holding colored squares in WM on search efficiency (Woodman & Luck, 2007). Apparently, changing the search template on each trial tapped into WM resources in a way that did not occur when the target was consistent. What would happen if, rather than changing the identity of the target template from one trial to the next, it was necessary for the observer to change the target template multiple times in a single trial? In the current set of hybrid search experiments, observers had to search for one of up to 64 possible targets. The 64-item set was held constant for a block of trials. Given the limited capacity of VWM (Luck & Vogel, 1997; Luck & Vogel, 2013), the full set of target templates must be held in activated long-term memory (ALTM: Cowan, 1995; Cowan, 2001), rather than working memory. Thus, if VWM is a mandatory component in visual search, then one might predict that items from the 64-member set in ALTM would be shuttled in and out of VWM during the course of a hybrid search. Following this logic, the larger the target list in ALTM, the greater would be the ALTM – VWM interaction and the greater would be the adverse effects of WM load. Forrin and Morin (1969) offer a hint to the contrary, reporting that LTM and WM do not interact in memory search. However, this work used a relatively small LTM manipulation of between 1 and 3 items and no visual search component.

In the present hybrid search experiments, involving visual and memory search, we used photorealistic objects, allowing us to employ much larger target sets, known to produce large costs in search efficiency (Wolfe, 2012). Given previous findings that VWM interferes with visual search and given the theory that the current search template resides in VWM, it comes as a surprise that, in seven experiments, we found little or no effect of loading VWM with task irrelevant material on hybrid search. However, we also found that performing a

hybrid search exacted a cost on VWM performance. This cost was independent of the size of the memory set in the hybrid task.

We will use the results of these experiments to propose quite a different account of the role of VWM in visual search. We suggest that the representation of the target (or targets) of a search resides in LTM (more precisely, in ALTM). In visual search, those target representations must be compared with items in the visual stimuli. We argue that the role of WM is not to hold the templates or transfer them from LTM to earlier visual processes. Instead, we hypothesize that WM acts to pass a representation of the current object of attention from earlier visual processes to LTM where it can be compared to the target template(s).

Materials and methods

In each block of each experiment, observers were asked to memorize 2, 8, 16, or 64 photos of real-world objects (Brady, Konkle, Alvarez, & Oliva, 2008). They then searched for the presence of a target item in visual search arrays that contained either 8 or 16 objects. On half of the trials, prior to the visual search array, observers were given a visual working memory load of three items. After responding to the search array, observers were again shown the visual working memory objects. On half of those trials, one of these items changed and the observers were asked to identify these change trials. A schematic of trials with and without a VWM load can be seen in Fig. 1. We were surprised to find no influence of a VWM load on search efficiency in Experiment 1. In the subsequent 6 experiments, we varied the nature of the VWM load in the hope and expectation of finding a situation where VWM load did interact with hybrid search. The possibility that the observers did not have enough time to encode information into the VWM in Experiment 1 led to Experiment 2, where the encoding display time was increased from 0.5 to 3 s. Experiment 3 employed simple color square stimuli that were distinct from the visual search stimuli. Experiments 4 and 5 asked observers to perform a spatial WM task that was modeled after well-known experiments that have demonstrated a strong interaction between WM load and search efficiency (Oh & Kim, 2004; Woodman & Luck, 2004). Experiment 6 examined the role of the phonological loop in maintaining VWM by asking the observers to perform the task while engaged in an articulatory suppression task. Finally, Experiment 7 explored the effect of hybrid search on VWM performance, rather than the other way around, by including trials that did not involve the hybrid search task.

Observers gave informed consent and were compensated \$10 per hour or through course credit. All observers had, at least, 20/25 acuity with correction, passed the Ishihara Color Blindness Test and were fluent speakers of English. Observers

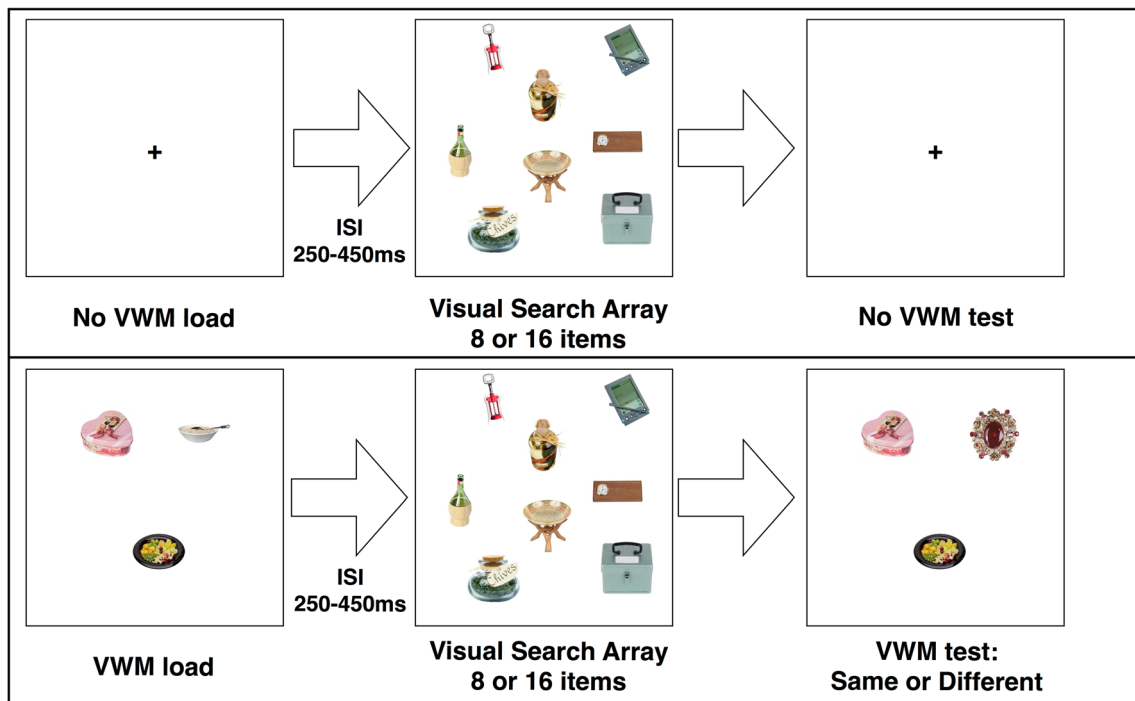


Fig. 1 Experimental paradigm schematic. Observers searched the visual search array for the presence of any of the previously memorized target set items. In Experiments 1-4 and 6, trials with and without a VWM load

were interleaved. One item changed in the VWM test on 50 % of the trials. The correct response in the lower example above would be “different,” because one of the items changed

in Experiments 1-6 were recruited from the Cambridge, MA, community. Observers in Experiment 7 were recruited from the University of Utah Human Subject Pool and participated for course credit.

Experiments 1-7 all had a similar design. This consisted of a memorization portion at the beginning of each block followed by the experimental trials. Each trial consisted of three parts: the working memory initial display, the search array, and the memory test array. For Experiments 1-6, observers were seated so that their eyes were 57 cm from a 20” CRT monitor with an 85-Hz refresh rate. For Experiment 7, sat 66 cm from a 19” LCD monitor with a 60-Hz refresh rate. Stimuli were either real-world objects that subtended 3.25° (2.81° in Experiment 7), colored squares (1.9° ; 1.65°), or smaller black squares (1.6° ; 1.39°). All experiments were written in MATLAB using Psychophysics Toolbox.

Based on our previous work using the hybrid search paradigm and the existing Working Memory – Visual Search, dual task literature (Oh & Kim, 2004; Woodman & Luck, 2004, 2007), sample size for each of the experiments was held to between 12-15 observers. Observers were run in groups of up to 10 at a time, with an unpredictable no-show rate. Data collection was stopped after >11 observers completed the experiment. Data were collected from a total of 101 observers. Of these, four individuals were eliminated from data analysis due to poor performance. Two of these observers performed below the chance level on the WM task for at least one block of the experiment. Error rate on the visual search task

exceeded 30 % for two other observers. One more observer was unable to finish the experiment in the allotted time.

All of the experiments began each block of trials with an identical memorization procedure similar to those described elsewhere (Drew & Wolfe, 2014; Wolfe, 2012). Observers were asked to memorize 2, 8, 16, or 64 real-world objects, which were presented individually for 3 seconds at a time. Each observer experienced all experimental conditions. Order of memory set size was randomized. All other conditions were randomized within block unless otherwise noted (as in Experiment 5). All objects were taken from a heterogeneous set of 3,000 unique photorealistic objects provided by Brady et al. (2008). During the recognition test that followed, a single object was displayed in the center of the screen and observers made “old” or “new” responses to either targets or distractor objects. Targets appeared 50 % of the time. The memory test contained twice the number of trials as targets in that block. Observers were required to perform this task with at least 90 % accuracy on two consecutive tests before being allowed to proceed. If performance fell below threshold, observers were retrained and retested. Block order was randomized between subjects.

After successfully learning the target set, observers completed 12 practice trials followed by 208 experimental trials. Each trial consisted of three parts. Half of the trials began with the display of the VWM load. For those trials that did not contain a VWM load, observers were shown a blank screen for the same duration as the VWM load. This was followed by

a random blank interval between 250–450 ms, followed by onset of the visual search array of either 8 or 16 real-world objects. Half of these arrays contained one of the target items. Nontarget “distractor” items and target items were drawn from separate sets so that an item that served as a target could never appear as a distractor. Observers indicated by key press whether a target was “present” or “absent” and were encouraged to answer as quickly and accurately as possible. After the response to the hybrid search, on trials with no VWM load, an on-screen message asked the observer to move on to the next trial. For VWM trials, the VWM test screen appeared immediately after the response to the hybrid search. The test consisted of three items, presented in the same locations as the original VWM load. For half of the trials, the identity of one of the objects changed. Throughout all of the experiments, object changes were categorical (either changing from one of seven distinct colors to another, or one unique item to another) and location changes necessitated a change from one of eight locations equally spaced around a centered circular area with a 6° diameter. Observers indicated using a key press whether the test array was the “same” or “different” from the initial display. The “present” and “absent” keys were located next to each other on either the left or right side of the keyboard and the “same” and “different” keys were located together on the opposite side of the keyboard. Key assignment was counterbalanced across observers.

Experiments 1 and 2: Object memory

Thirteen observers participated in Experiment 1 and 11 participated in Experiment 2 (average age 30.7 years, 15 females). The WM displays consisted of photos of real-world objects drawn from the same superset as the targets and distractors in the hybrid search. Targets, distractors, and VWM objects were all drawn from distinct subsets of that superset. In Experiment 1, objects in the WM initial display were displayed for 0.5 seconds, and in Experiment 2 objects were displayed for 3 seconds.

Experiment 3: Colored square memory

Twelve observers (average age 29.3 years, 6 females) participated in Experiment 3. Working memory was loaded with three colored squares that were presented for 500 ms. The colors of the squares were randomly selected from a set of seven colors (red, green, blue, yellow, black, gray, and light blue). As in Experiments 1 and 2, the working memory trials were randomly placed within a particular hybrid search block.

Experiments 4 and 5: Location memory

Fourteen different observers participated in Experiments 4 and 5 (average age 29.2 years, 14 females). Following the methods of Oh and Kim (2004), in these experiments, the initial working memory displays consisted of three black squares (500 ms), and observers were told to remember the locations. In the test array, only a single square was presented. Observers were asked if this square was in a location that matched any of the three locations in the initial display. In Experiment 4, working memory trials were randomized among nonworking memory trials, whereas in Experiment 5, the working memory trials were blocked. The order of the WM load and no load blocks was counterbalanced between observers in Experiment 5.

Experiment 6: Articulatory suppression

Fourteen observers participated in Experiment 6 (average age 32.6, 6 females). Experiment 6 replicated Experiment 3 (colored square memory) with the following exceptions. Observers in this experiment were asked to memorize 2 or 16 objects in different blocks of the experiment. As in Experiment 5, WM load and no-load trials were blocked. The order of WM load and no load blocks was counterbalanced across observers. At the beginning of each block of visual search trials, observers were instructed to recite “abcd,” “wxyz,” “1234,” or “6789” throughout that block of trials. The experimenter sat in the room with the observers to ensure that this phrase was audibly repeated throughout the experiment. The phrase for a given block was randomly permuted for each observer.

Experiment 7: What is the cost of visual search on the VWM performance?

The focus of Experiments 1–6 is the effect of WM load on hybrid search. Experiment 7 allowed us to look at the influence of hybrid search on WM capacity. Twenty-one observers participated in Experiment 7. One observer did not finish in the allotted time and was removed from further analyses (average age 26.6, 10 females). Experiment 7 replicated Experiment 3 (colored square memory) with the following exceptions. Observers in this experiment were asked to memorize 8 or 16 objects in different blocks of the experiment. There were three trial types in this experiment. Dual-task and No-Memory conditions were identical to trials in the previous experiments. Observers were also shown no-search trials. For these trials, the working memory load screen (three colored squares) preceded a screen that informed the observer that there was no search task on this trial and to press a button when ready to continue. All trial types were randomly interleaved with the block.

Results

The results for Experiments 1-5 can be seen in Fig. 2 and Tables 1 and 2. During the course of seven experiments, we found very little evidence of an interaction between VWM load and target set size. For trials with a WM load, we focused our analyses on those trials where the WM task was answered correctly. We assessed three primary dependent measures of hybrid search performance: mean reaction time (RT), search efficiency (slope), and d' (sensitivity). Search efficiency was the measure of primary interest; RT and d' were essentially used as controls to ensure that if we did not find an effect of VWM load on slope of the RT x set size functions, this was not due to a speed accuracy trade-off, or an overall accuracy decrement in the face of the VWM load. Previous evidence has suggested that RT increases as a linear function of the log of the memory set size (Wolfe, 2012; Cunningham & Wolfe, 2014). Accordingly, we use a \log_2 scale for memory set size on the X-axis on graphs throughout the paper.

















To summarize, in each experiment we computed a 2×4 repeated measures ANOVA with WM load and target set size as factors. Not surprisingly, there was a large, reliable main effect of the number of targets held in ALTM (target set size) in each experiment. The VWM load had a less consistent effect. In Experiment 1 (the first object experiment) and Experiment 3 (the color experiment), there was a significant effect of WM load on reaction time, but only on absent trials (Exp1: 78-ms difference, $F(1,12) = 7.43$, $p = 0.018$; Exp3: 275-ms difference, $F(1,11) = 26.151$, $p < 0.001$). There was no main effect of WM load on RT in any of the other experiments. When present, the observed main effects on RT are broadly consistent with previous results (Oh & Kim, 2004; Woodman & Luck, 2004, 2007; Woodman, Vogel, & Luck, 2001), although they are smaller and less reliable. With the exception of present trials in Experiment 3 (colored squares), search efficiency was unaffected by the presence or absence of a VWM load. This is in contrast to previous work that suggested that when VWM loads require observers to encode location, there is a decrease in search efficiency (Oh & Kim, 2004; Woodman & Luck, 2004). If the effect in Experiment 3 is reliable, perhaps it is the use of colored squares that is critical and was missing in the Experiments that failed to detect an interaction between search efficiency and VWM load. To test this hypothesis, observers in Experiments 6 and 7 were asked to memorize color squares. The efficiency X VWM load interaction did not replicate in either experiment. Thus, we found that there was a significant main effect of VWM load on slope for just 1 of the 6 (present or absent) possibilities using colored square memory and only 1 of the 14 total possibilities for all experiments. On balance, VWM load does not appear to reliably alter search efficiency in these experiments.

Finally, we found very little evidence for an interaction between memory set size in the hybrid search task and WM load. Of the five experiments for both present and absent trials, there was just one instance where the interaction was modestly statistically significant: present trials in Experiment 5: $F(3,33) = 4.28$, $p = 0.043$. We address this finding in more detail in the discussion of Experiments 6 and 7.





In each experiment, we computed both d' and c (criterion) to examine whether overall task performance or response bias varied systematically as a function of VWM load or target set size. There was a reliable effect of target set size (all $F_s > 19$, all $p_s < 0.001$) but not VWM load (all $F_s < 1.8$, all $p_s > 0.2$) on d' in Experiments 1-5. The c results are similar: there was a reliable effect of target set size in every experiment (all $F_s > 2.9$, all $p_s < 0.05$), and no effect VWM load (all $F_s < 2.3$, all $p_s > 0.16$) except in Experiment 1 ($F(1,12) = 9.6$, $p < 0.01$). On the whole, these results suggest that while target set size reliably influenced both task difficulty and response bias (larger target set size led to a greater likelihood of “present” responses). VWM did not consistently influence either of these factors. This can be seen clearly in Fig. 3, which plots each observer's d' and criterion scores with and without a VWM load for Experiments 1-5.

Bayes Factor

Our primary interest was to determine whether holding information in working memory led to greater search difficulty when memory set size was increased. Whereas none of the interactions for this effect were significant for any of the three dependent variables (RT, slope, and d') in Experiments 1-5, it can be difficult to interpret null effects using traditional approaches (Wagenmakers, 2007). Accordingly, we used Bayes Factor (BF) calculations to evaluate whether the interaction between WM and memory set size was or was not contributing meaningfully to our results relative to a model that only included the main effects. We used the BayesFactor 0.9.6 package (Rouder & Morey, 2012) in R, which implements the Jeffreys–Zellner–Siow (JZS) default on effect sizes (Rouder, Morey, Speckman, & Province, 2012). This analysis yields a likelihood ratio that reflects the relative probability of the data arising from a pure *main effect model* compared to the *interactive model*. Thus, a BF of 4 can be interpreted as meaning that the model A is 4 times more likely to be supported than model B. Results are shown in Tables 1 and 2; large numbers indicate support for the *main effect model*. The preponderance of evidence favors the *main effect model*. In every case, the *main effect model* does a better job of explaining the data than the *interactive model*. This includes the present trials in Experiment 5, where the BF was 4.33 even though the interaction was modestly “significant” (0.043) in the ANOVA.



Interleaved WM Task



Blocked WM Task

◀ **Fig. 2** Results from Experiments 1-5 denoting (from left to right) the materials used in the WM test, search efficiency (as measured by the slopes of the RT x visual set size for each memory set size), RT (averaged over visual set size), and d' as a function of memory set size. Note that the X-axis scale is logarithmic, rather than linear. See text for additional details. VWM performance (black line) is overlaid on the RT graph with proportion incorrect on the right axis. Critically, performance is very similar with (solid) and without (dashed) a working memory load. Error bars here and throughout the paper represent standard error of the mean

Visual working memory performance

The difficulty of the VWM task varied significantly across experiments (Fig. 2, second column). As expected, increasing the duration of the memory array from 0.5 to 3 s resulted in improved performance in Experiment 2 ($t(24) = 8.23, p < 0.001$). However, this difference in the amount of information that was successfully encoded into VWM had no influence on the interaction between the VWM load and the target set size. Similarly, although the color squares in Experiment 3 were easier to memorize than the locations in Experiment 4 ($t(24) = 5.32, p < 0.001$), the additive model was preferred to the model with the interaction effect for all three of our dependent measures in both cases (Tables 1 and 2). Importantly, for our later discussion, there were no significant effects of memory set size on WM capacity (all $p > 0.05$).

Control Experiment 6: The role of verbal encoding

Based on previous results, we were surprised to find that VWM load did not have a reliable effect on search slope in Experiments 1-5. Perhaps the absence of an effect was due to observers verbally encoding VWM information, thereby leaving the hybrid search task unaffected by this manipulation. If this were the case, we would expect that the addition of an articulatory suppression task (repeating a simple set of letters or numbers throughout the trials) would result in a large effect of VWM load on the search efficiency. One also might expect that this effect would interact with memory set size such that there would be a larger effect of VWM load when searching for a larger number of possible items. Experiment 6 examined this question using the colored square VWM stimuli. Results are shown, compared with the equivalent Experiment 3, in Fig. 4. Consistent with previous results, we found that although there was a large effect of memory set size on search efficiency (absent: $F(1,13) = 9.68, p < 0.001$; present: $F(1,13) = 22.45, p < 0.001$), there was no effect of the presence of a VWM load (absent: $F(1,13) = 0.03, p = 0.85$; present: $F(1,13) = 0.18, p = 0.68$) and the two factors did not interact reliably (absent: $F(1,13) = 0.6, p = 0.44$; present: $F(1,13) = 0.78, p = 0.39$). These results suggest that the lack of influence of the VWM load on search efficiency in Experiments 1-5 was not driven by a verbal encoding strategy.

Table 1 Absent trials and d' results. Results of repeated measures ANOVAs and Bayes Factor estimation for Experiments 1-5

	Slope	RT	d'
Object (0.5 s)	WM: $F(1,12) = 3.94, p = 0.071$ Mset: $F(3,36) = 6.64, p = 0.001$ Interaction: $F(3,36) = 0.8411, p = 0.480, b.f. = 7.64$	WM: $F(1,12) = 7.43, p = 0.018$ Mset: $F(3,36) = 57.94, p < 0.001$ Interaction: $F(3,36) = 0.05, p = 0.983, b.f. = 9.48$	WM: $F(1,12) = 1.77, p = 0.21$ Mset: $F(3,36) = 19.33, p < 0.001$ Interaction: $F(3,36) = 0.69, p = 0.56, b.f. = 6.71$
Object (3 s)	WM: $F(1,10) = 0.114, p = 0.742$ Mset: $F(3,30) = 4.92, p = 0.007$ Interaction: $F(3,30) = 0.61, p = 0.613, b.f. = 4.88$	WM: $F(1,10) = 1.22, p = 0.294$ Mset: $F(3,30) = 53.19, p < 0.001$ Interaction: $F(3,30) = 1.12, p = 0.354, b.f. = 7.77$	WM: $F(1,10) = 0.008, p = 0.92$ Mset: $F(3,30) = 26.46, p < 0.001$ Interaction: $F(3,30) = 0.419, p = 0.741, b.f. = 7.10$
Color (0.5 s)	WM: $F(1,11) = 0.227, p = 0.642$ Mset: $F(3,33) = 22.878, p < 0.001$ Interaction: $F(3,33) = 1.404, p = 0.258, b.f. = 5.708$	WM: $F(1,11) = 26.151, p < 0.001$ Mset: $F(3,33) = 53.25, p < 0.001$ Interaction: $F(3,33) = 1.41, p = 0.258, b.f. = 7.87$	WM: $F(1,11) = 0.273, p = 0.611$ Mset: $F(3,33) = 43.86, p < 0.001$ Interaction: $F(3,33) = 0.156, p = 0.924, b.f. = 7.891$
Location (0.5 s) interleaved	WM: $F(1,13) = 0.003, p = 0.955$ Mset: $F(3,39) = 16.76, p < 0.001$ Interaction: $F(3,39) = 2.54, p = 0.0706, b.f. = 2.50$	WM: $F(1,13) = 1.628, p < 0.224$ Mset: $F(3,39) = 51.50, p < 0.001$ Interaction: $F(3,39) = 1.23, p = 0.311, b.f. = 8.49$	WM: $F(1,13) = 0.004, p = 0.951$ Mset: $F(3,39) = 46.52, p < 0.001$ Interaction: $F(3,39) = 1.189, p = .326, b.f. = 5.51$
Location (0.5 s) blocked	WM: $F(1,11) = 1.99, p = 0.18$ Mset: $F(3,33) = 8.05, p < 0.001$ Interaction: $F(3,33) = 2.01, p = 0.13, b.f. = 1.05$	WM: $F(1,11) = 0.0, p = 0.998$ Mset: $F(3,33) = 5.88, p < 0.001$ Interaction: $F(3,33) = 3.08, p = 0.082, b.f. = 8.12$	WM: $F(1,11) = 0.977, p = 0.344$ Mset: $F(3,33) = 25.08, p < 0.001$ Interaction: $F(3,33) = 0.85, p = 0.476, b.f. = 6.77$

Table 2 Present trial results. Results of repeated measures ANOVAs and Bayes Factor estimation for Experiments 1-5

	Slope	RT
Object (0.5 s)	WM: $F(1,12) = 0.410, p = 0.534$ Mset: $F(3,36) = 44.7, p < 0.001$ Interaction: $F(3,36) = 0.402, p = 0.752, b.f. = 5.24$	WM: $F(1,12) = 1.43, p = 0.254$ Mset: $F(3,36) = 70.6, p < 0.001$ Interaction: $F(3,36) = 0.104, p = 0.957, b.f. = 8.85$
Object (3 s)	WM: $F(1,10) = 0.534, p = 0.481$ Mset: $F(3,30) = 16.91, p < 0.001$ Interaction: $F(3,30) = 1.55, p = 0.220, b.f. = 5.41$	WM: $F(1,10) = 0.03, p = 0.865$ Mset: $F(3,30) = 56.72, p < 0.001$ Interaction: $F(3,30) = 1.015, p = 0.399, b.f. = 5.79$
Color (0.5 s)	WM: $F(1,11) = 9.766, p = 0.0096$ Mset: $F(3,33) = 27.23, p < 0.001$ Interaction: $F(3,33) = 0.774, p = 0.516, b.f. = 3.488$	WM: $F(1,11) = 19.6, p = 0.001$ Mset: $F(3,33) = 55.88, p < 0.001$ Interaction: $F(3,33) = 1.70, p = 0.184, b.f. = 6.23$
Location (0.5 s) Interleaved	WM: $F(1,13) = 0.472, p = 0.504$ Mset: $F(3,39) = 28.689, p < 0.001$ Interaction: $F(3,39) = 2.104, p = 0.115, b.f. = 0.811$	WM: $F(1,13) = 1.23, p = 0.287$ Mset: $F(3,39) = 60.85, p < 0.001$ Interaction: $F(3,39) = 0.843, p = 0.478, b.f. = 7.49$
Location (0.5 s) Blocked	WM: $F(1,11) = 0.07, p = 0.796$ Mset: $F(3,33) = 36.54, p < 0.001$ Interaction: $F(3,33) = 4.28, p = \mathbf{0.043}, b.f. = 4.33$	WM: $F(1,11) = 0.18, p = 0.681$ Mset: $F(3,33) = 49.4, p < 0.001$ Interaction: $F(3,33) = 3.14, p = 0.314, b.f. = 2.29$

Note that the bolded p-value represents the only instance where slope significantly interacted with WM and Mset

Results from the RT and d' data from this Experiment 6 also were consistent with the general pattern in Experiments 1-5. There was a strong effect of memory set size for both measures (all $p < 0.001$) but no main effect of VWM load (all $p > 0.2$) or interaction between the two factors (all $p > 0.1$).

Experiment 7: Effects of hybrid search on working memory

Experiments 1-6 make it quite clear that WM load does not have any consistent effect on hybrid search performance, but what about the other way around? Does performing a hybrid search during the retention period influence WM? It is notable

that VWM in Experiment 1-5 is consistently lower than in previous published data involving three simple items (Luck & Vogel, 1997). Moreover, it is interesting that WM performance is basically constant across memory set sizes in the hybrid search task. To determine whether this apparent decrement in VWM performance was caused by the presence of a hybrid search task during the retention interval, we performed an additional experiment that contained some trials that did not contain the hybrid search task.

In Experiment 7, one-third of trials did not contain a hybrid search Task, one-third did not have a VWM task, and one-third had both. When there was no hybrid search task, after encoding the VWM material, observers were instructed to

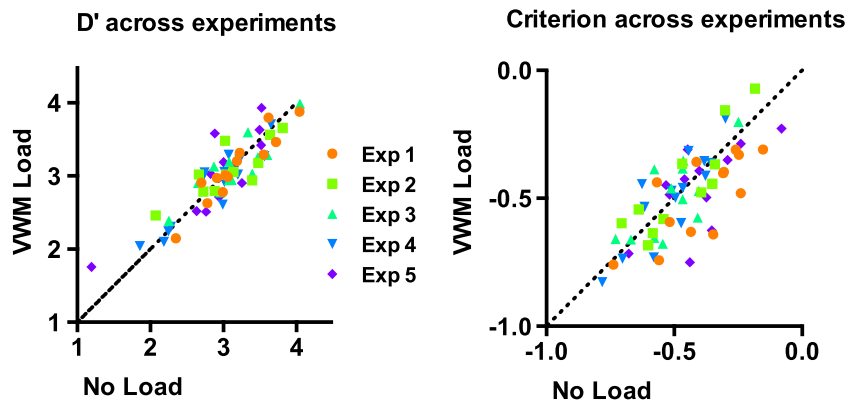


Fig. 3 D-prime and criterion data for Experiments 1-5. Across experiments, VWM load did not have a consistent effect on task difficulty or response bias

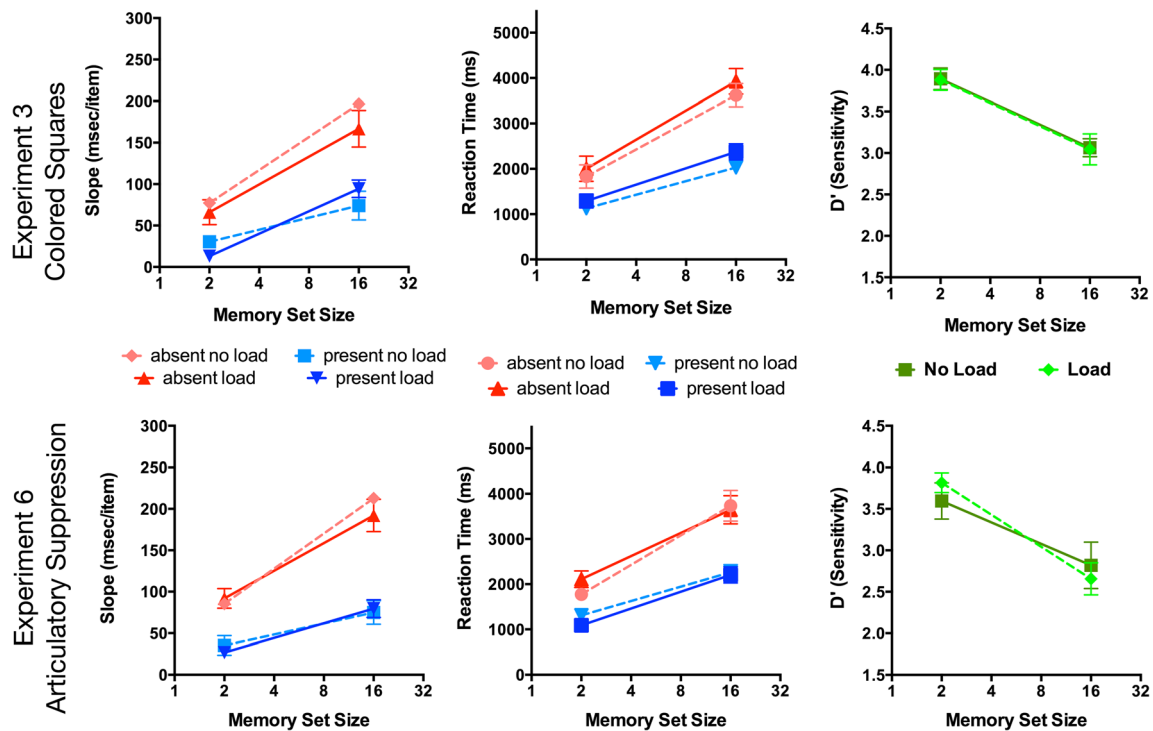


Fig. 4 Comparison between results for Experiments 3 and 6. The experiments were identical (including a VWM load of colored squares in both cases) aside from the articulatory suppression task observers

performed in Experiment 6. In addition, VWM trials were blocked, rather than interleaved, in Experiment 6

press a button to move on to the VWM test screen. All three trial types (Search Only, Memory Only, and Dual Task) were randomly interleaved. As in Experiments 3 and 6, the VWM load in this experiment was a set of colored squares.

The result from the Search Only and Dual Task trials are consistent with previous findings. There was a large effect of memory set size on search efficiency (absent: $F(1,19) = 4.61, p = 0.04$; present: $F(1,19) = 9.52, p = 0.006$), no effect of the presence of a VWM load (absent: $F(1,19) = 0.77, p = 0.39$; present: $F(1,19) = 1.95, p = 0.18$), and the two factors did not interact reliably (absent: $F(1,19) = 0.07, p = 0.79$; present: $F(1,19) = 2.17, p = 0.16$). However, VWM performance was strongly influenced by the presence of hybrid search task ($F(1,19) = 92.57, p < 0.001$). Performance in the absence of the search task was markedly higher (91 % correct, $k = 2.46$) than when the task was present (78 % correct, $k = 1.68$). Thus, while holding information in working memory does not appear to influence the efficiency with which we search through space during a hybrid search task, the act of performing a hybrid search task results in a the loss of about one “slot” worth of WM capacity (although this need not be a literal “slot.” We make no commitments in that debate: Suchow, Fougne, Brady, & Alvarez, 2014). While it is clear that, consistent with previous results (Woodman & Luck, 2010), there was dual-task interference, the interference appears to be a main effect on WM that does influence the efficiency with which the hybrid search task is performed.

General discussion

What do these results tell us about the role of WM in hybrid search and, perhaps, visual search more generally. Consider the outlines of a model, shown in Fig. 5.

In hybrid search, observers seek to determine if any of the items in the world are in the memory set of items, held in ALTM. Of course, the interaction of a visible item and its ALTM representation cannot be occurring in the world. The visible item must be selected by spatial attention. How and where is that representation compared to the target “templates” in ALTM? One possibility is that the current template is held in VWM and that the comparison is done in WM (Olivers et al., 2011). However, even if VWM can hold more than one template at a time (Gilchrist & Cowan, 2011), one of the defining characteristics of WM is a severely limited capacity. It is therefore implausible that a memory set of 8, 16, or 64 items could be loaded into VWM in one step. Furthermore, if the comparison between the visual world and the memory set is occurring in VWM, one would expect that loading VWM with additional information should strongly influence the efficiency with which this process is performed. In the seven experiments outlined, we found no evidence of this. Thus, this most basic model can be ruled out.

Perhaps, VWM is the arena for the comparison of visual item and search template, but in the case of a large memory set, templates are swapped in and out of VWM from ALTM.




Fig. 5 In a hybrid search, observers determine if an item in the world matches any item in activate long-term memory (ALTM). We hypothesize that items are selected by visual attentional mechanisms

and passed, one at a time, through visual working memory (VWM) to ALTM. In ALTM, the selected item is compared to the memory set in a process whose duration is a logarithmic function of the memory set size

The problem with this account is that the number of required templates swaps would seem to be linear with the size of the memory set. Thus, it should take twice as many swaps to decide if the currently attended visual item is one of 16 items in memory than if it is one of 8. It is hard to see why such a linear process would produce RTs that were a logarithmic function of memory set size (Wolfe, 2012). Furthermore, the number of swaps necessary should be influenced by the presence or absence of VWM load, which would predict a strong effect of VWM load on search efficiency; an effect we do not see in any of the experiments reported here.

One could propose that the comparison is done in LTM. A visual representation is passed up the visual hierarchy to LTM, where it is matched against the ALTM representations of the memory set, bypassing VWM altogether. This, however, ignores the evidence that WM contents do influence visual search and, in particular for present purposes, the evidence of Experiment 7 that hybrid search produces a clear reduction in the WM capacity available to the change detection WM task.

An account that seems consistent with the current data proposes that WM is the narrow bottleneck that must be traversed by the visual representation of an item as it is moved from the visual system to LTM, where it can be compared to the ALTM representations of the memory set. VWM is, of course, very limited in capacity, so when a visual item must be moved through an already loaded VWM, it must displace the representation of some VWM information, effectively reducing the capacity of VWM. Thus, performance is reduced on the subsequent test of VWM. If we assume that a one-item channel must be maintained to continue moving visual item after visual item to LTM, we see that the cost of the hybrid search – VWM interaction will not be dependent on either the visual set size or the memory set size. It is simply the fixed cost of a one-item path through VWM. Thus, this account is consistent with the fixed, reduced VWM capacity seen in Experiments 1-5 (WM performance data in Fig. 2). It is consistent with the lack of an effect of VWM load on hybrid search, because in this situation hybrid search simply commandeers the same one-item path in all conditions. This account predicts that hybrid search (or any visual search) would be impossible, if one could force the observer not to

relinquish any VWM capacity to search, although it is not currently clear how one would accomplish this goal.

In the present work, there is only exception to the general pattern that VWM load does not interfere with hybrid search efficiency. There was a modest effect on present (but not absent) trials in Experiment 5 ($p = 0.043$). The distinguishing manipulation in that experiment was that the VWM manipulation was blocked rather than interleaved. VWM trials were blocked in Experiment 7 as well, but there was no evidence of a statistically reliable interaction between VWM load and search efficiency in that experiment (p values for both present and absent trials >0.3). This suggests that the Experiment 5 finding may not be robust. However, there are prior studies where spatial working memory has been shown to interact reliably with visual search (Oh & Kim, 2004; Woodman & Luck, 2004). In the account offered, we propose that VWM was more resistant to surrendering the one-item pathway in those studies due to the inherent difference between hybrid search, where observers search for one potentially many possible targets amongst heterogeneous stimuli, and more traditional visual search, where observers learn to search for single target amongst relatively homogenous stimuli. This possibility needs further test.

Almost every model of visual search includes a role for WM (Bundesen, 1990; Desimone & Duncan, 1995; Logan & Gordon, 2001; Olivers et al., 2011; Wolfe, 1994). Typically, WM is thought to allow the observer to transfer small groups of items into WM for comparison with the target template. Contrary to this popular current view, the present results suggest that the target templates in hybrid search are not taking up space in VWM. As noted, this interpretation is implausible when there are large numbers of target templates as there must be in hybrid search with large memory set sizes. We have suggested that these templates are held in activated long-term memory (Cunningham & Wolfe, 2014; Drew & Wolfe, 2014; Wolfe, 2012). With the present results, we have clearly shown that search efficiency is not influenced by the presence or absence of a VWM load. We hypothesize that the role of WM in search is to serve as a conduit that passes visual representations to LTM where they can be compared to search template(s). In our experiments, that conduit remains open even when a VWM load is added. Apparently, the search task

took priority such that the cost was paid by the WM task. We would predict that, if WM were disabled, hybrid search would fail.

While little is known about the role of WM in searching through ALTM, the link between LTM search and WM is relatively well understood. Retrieval from LTM is thought to involve both random and directed components (Shiffrin, 1970; Unsworth, Brewer, & Spillers, 2013). Working memory seems to be important in selecting appropriate directed LTM search strategies (Raaijmakers & Shiffrin, 1981; Unsworth et al., 2013). While high WM capacity observers are able to recall more items from an LTM category than low WM capacity observers during free recall, when search strategy was minimized via cued-recall, the differences disappeared (Unsworth et al., 2013).

It would be interesting to see if a difference between high and low WM capacity is seen in hybrid search. The current studies were not designed to directly address this question. It is possible that individuals with higher WM capacity (WMC) would be able to search more efficiently if some observers can pass items through VWM at a higher rate than others. Recent work by Anderson and colleagues supported this prediction (Anderson et al., 2013). The authors suggest that higher WMC observers can compare more items in parallel than lower WMC observers. In future work, it would be interesting to examine if WMC interacted with memory set size in these experiments.

In summary, in the experiments reported here, loading VWM had virtually no effect on hybrid search. In contrast, hybrid search had a clear, fixed effect on VWM, independent of memory set size in the hybrid task. We argue that VWM is, indeed, required for search tasks but that what is required is a path through VWM to pass a representation of the current object of attention to LTM, where that representation can be compared to templates stored in ALTM.

References

- Anderson, D. E., Vogel E. K., & Awh E. (2013). "A common discrete resource for visual working memory and visual search." *Psychological Science*, 24(6), 929–938.
- Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences of the United States of America*, 105(38), 14325–14329.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, 97, 523–547.
- 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.
- Cowan, N. (1995). *Attention and memory: An integrated framework*. Oxford University Press on Demand.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87–185.
- Cunningham, C. A., & Wolfe, J. M. (2014). The role of object categories in hybrid visual and memory search. *Journal of Experimental Psychology: General*, 143(4), 1585.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222.
- Downing, P., & Dodds, C. (2004). Competition in visual working memory for control of search. *Visual Cognition*, 11(6), 689–703.
- Drew, T., & Wolfe, J. M. (2014). Hybrid search in temporal domain: Evidence for rapid, serial logarithmic search through memory. *Attention, Perception, & Psychophysics*, 76, 296–303.
- Forrin, B., & Morin, R. E. (1969). Recognition times for items in short- and long-term memory. *Acta Psychologica*, 30, 126–141.
- Gilchrist, A. L., & Cowan, N. (2011). Can the focus of attention accommodate multiple, separate items? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(6), 1484–1502.
- Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences*, 9(2), 75–82.
- Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, 108, 393–434.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17(8), 391–400.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24, 167–202.
- Oh, S. H., & Kim, M.-S. (2004). The role of spatial working memory in visual search efficiency. *Psychonomic Bulletin & Review*, 11(2), 275–281.
- Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology-Human Perception and Performance*, 32(5), 1243–1265.
- Olivers, C. N., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, 15(7), 327–334.
- Raaijmakers, J. G., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88(2), 93.
- Rouder, J. N., & Morey, R. D. (2012). Default Bayes factors for model selection in regression. *Multivariate Behavioral Research*, 47(6), 877–903.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356–374.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84(1), 1–66.
- Shiffrin, R. M. (1970). Memory search. In D. A. Norman (Ed.), *Models of human memory* (pp. 375–447). New York: Academic Press.
- Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*, 12(9), 342–348.
- Suchow, J., Fournie, D., Brady, T., & Alvarez, G. (2014). Terms of the debate on the format and structure of visual memory. *Attention, Perception, & Psychophysics*, 76(7), 2071–2079.
- Unsworth, N., Brewer, G. A., & Spillers, G. J. (2013). Working memory capacity and retrieval from long-term memory: The role of controlled search. *Memory & Cognition*, 41(2), 242–254.
- Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review*, 14(5), 779–804.
- Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, 1, 202–238.

- Wolfe, J. M. (2012). Saved by a log: How do humans perform hybrid visual and memory search? *Psychological Science*, *23*(7), 698–703.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 419–433.
- Woodman, G. F., & Luck, S. J. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin & Review*, *11*(2), 269–274.
- Woodman, G. F., & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology-Human Perception and Performance*, *33*(2), 363–377.
- Woodman, G. F., & Luck, S. J. (2010). Why is information displaced from visual working memory during visual search? *Visual Cognition*, *18*(2), 275–295. doi:[10.1080/13506280902734326](https://doi.org/10.1080/13506280902734326)
- Woodman, G. F., Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science*, *12*(3), 219–224.

A Soft Handoff of Attention between Cerebral Hemispheres

Trafton Drew,^{1,2,4,*} Irida Mance,³ Todd S. Horowitz,^{1,2} Jeremy M. Wolfe,^{1,2} and Edward K. Vogel³

¹Department of Surgery, Brigham and Women's Hospital, 64 Sidney Street, Suite 170, Cambridge, MA 02139, USA

²Harvard Medical School, 64 Sidney Street, Suite 170, Cambridge, MA 02139, USA

³Department of Psychology, University of Oregon, Eugene, OR 97403, USA

Summary

Each cerebral hemisphere initially processes one half of the visual world. How are moving objects seamlessly tracked when they traverse visual hemifields? Covert tracking of lateralized objects evokes a difference between slow-wave electrophysiological activity observed from contralateral and ipsilateral electrodes in occipitoparietal regions. This event-related potentials (ERP) waveform, known as contralateral delay activity (CDA) [1, 2], is sensitive to the number of objects tracked [1, 2] and responds dynamically to changes in this quantity [3]. When a tracked object crosses the midline, an inversion in CDA polarity revealed the dropping of the object's representation by one hemisphere and its acquisition by the other. Importantly, our data suggest that the initially tracking hemisphere continues to represent the object for a period after that object crosses the midline. Meanwhile, the receiving hemisphere begins to represent the object before the object crosses the midline, leading to a period in which the object is represented by both hemispheres. Further, this overlap in representation is reduced if the midline crossing is unpredictable. Thus, this process is sensitive to observer expectations and does not simply reflect overlapping receptive fields near the midline.

Results and Discussion

We recorded event-related potentials (ERPs) from healthy young adults as they covertly tracked a vertically or horizontally moving object while holding central fixation (see the [Supplemental Results and Discussion](#) available online for additional information on eye movements). As shown in [Figure 1A](#), on each trial, a pair of objects was presented in each quadrant. A brief (500 ms) cue informed the observer which object to track. ERP waveforms were time-locked to the onset of this cue. When the cue disappeared, all objects began to move either clockwise or counterclockwise, taking each pair over either the horizontal or vertical midline. Movement toward midlines was held constant so that all objects crossed their respective midline at the same time. Movement in the orthogonal direction was less constrained. For example, as a pair of objects moved to the right, they would vacillate up and down, allowing their paths to cross and making tracking difficult (see the [Supplemental Results and Discussion](#) and [Figures](#)

[S1](#) and [S2](#)). As the objects were otherwise identical, close attention was required in order to track the target. Objects moved for 2.55 s, crossing the midline 1.70 s after cue onset (1.20 s after motion onset). This design ensured that the number of objects in each visual hemifield was always balanced and the distance traveled by objects on vertical and horizontal trials was identical. For more information on experimental procedures and behavioral data, see the [Supplemental Information](#).

Each trial was categorized in terms of whether the tracked object crossed the vertical or the horizontal midline. We averaged across five pairs of occipitoparietal electrodes (selected based on prior work [1, 4]) and categorized the two resultant waveforms as contra- or ipsilateral with respect to the initial position of the tracked object (see [Figure 1B](#)). To simplify analysis, we collapsed across direction of motion and initial position. Only correct trials with no eye movements or blinks artifacts were included. In our analyses, we refer to activity over the initially contralateral hemisphere as the *source hemisphere activity* and activity from the initially ipsilateral hemisphere as *target hemisphere activity*. (Note that this naming convention is specific to the horizontal condition: in the vertical condition, the target hemisphere never receives the object information since it was confined to a single visual hemifield.)

On vertical trials, when the tracked object was *not* switching hemispheres, we observed a large contralateral delay activity (CDA) in the time window before the attended object crossed the horizontal meridian (800–1,200 ms: $t(13) = 7.76$, $p < 0.001$) and a similar CDA after the crossing (2,000–2,400 ms: $t(13) = 7.35$, $p < 0.001$). There were no differences between those time periods ($t(13) = 1.64$, $p = 0.123$). In contrast, on horizontal trials when the attended objects crossed the vertical midline and moved from one visual field to the other, we observed a large CDA prior to the crossing (800–1,200 ms: $t(13) = 11.03$, $p < 0.001$). The waveform then inverted in polarity shortly after the tracked object crossed the vertical meridian, such that ipsilateral activity was more negative than contralateral activity (2,000–2,400 ms: $t(13) = -3.54$, $p = 0.004$). As predicted, activity from the hemisphere contralateral to the *current* location of the tracked object was more negative than ipsilateral activity regardless of whether the tracked object stayed within a hemifield or crossed between fields, revealing a dynamic remapping of attended object information between hemispheres.

Like the handoff between cellular phone towers transferring a live call on a moving mobile device, the handoff between hemispheres can be decomposed into two events that could occur at different times. There is a moment when the target hemisphere picks up the attended object information and another when the source hemisphere drops the information. This hemispheric handoff is analogous to presaccadic remapping [5, 6], where two findings are consistent across a range of methodologies [7–10]. First, target information is typically picked up *prospectively*, meaning that cells at the new, post-saccadic position represent the object before the completion of the saccade [7]. Second, the remapping closely approximates a “hard handoff” in which the cells that code the object in its original eye position quickly truncate their activity once

⁴Present address: Psychology Department, University of Utah, 380 S. 1530 E. BEHS 502, Salt Lake City, UT 84112, USA

*Correspondence: traftondrew@gmail.com




Figure 1. Experimental Paradigm and Contralateral Waveforms for Experiment 1

(A) Schematic illustration of experiment 1 paradigm. Dotted lines and light gray lines were not visible in the experiment.

(B) Vertical and horizontal ERP waveforms broken down by object movement type. Note that negative is plotted up here and throughout the paper.

See also [Figure S1](#).

To examine the relative timing differences between pick up and release of attended object information, we used a bootstrapping procedure to determine whether the increase in target hemisphere activity became significant reliably before the decrease in source hemisphere activity. Specifically, we randomized condition labels (target or source) and then sampled with replacement as many waveform pairs as we had observers. We computed the difference between pickup and drop-off points for each bootstrapped data set and then repeated this procedure 10,000 times in order to create a sample distribution [14]. Z scores of the observed data relative to the bootstrap data served as a measure of the likelihood that the

new location begins to be represented, resulting in little temporal overlap between the activity of the two sets of neurons [11, 12].

To measure the pickup and release during the handoff of an attended item moving across the midline, we compared activity from the source and target hemispheres as a function of movement type (horizontal versus vertical; [Figure 2A](#)). In this analysis, the vertical condition served as a baseline against which to compare the relative difference evoked as the attended object moved from source to target hemifields. As the two conditions were otherwise matched, we infer that any difference observed between the two is due to the handoff of object information between hemispheres. We determined when the vertical and horizontal trials began to differ for each hemisphere by sliding a 4 ms window backward in time from the end of the trial until at least ten successive time windows did not differ significantly at $p = 0.001$ ([Figure 2B](#); see the [Supplemental Results and Discussion](#) for more details on analysis). This conservative analysis does not measure the time at which the handoff is first initiated, but instead provides an upper-bound estimate of the time by which the handoff has reliably occurred [13]. By this measure, the source hemisphere released object information at 2,372 ms (SD: 89 ms), or 672 ms after the attended object crossed the midline, having traveled 2.0° past the midline by that time. At this point, the source hemisphere showed a reduced negativity for horizontal trials relative to vertical trials, indicating the release of attended object information. Turning to pickup by the target hemisphere, divergence occurred at 1,292 ms (SD: 355 ms), roughly 400 ms (1.2°) prior to the attended object crossing the midline. Similar to saccadic remapping, the target hemisphere appears to prospectively pick up the incoming attended object well before it crosses the midline.

observed values occurred by chance. We found that the difference in target activity became significant reliably earlier than source activity (Z score = 4.05, $p < 0.001$). This is strong evidence that the target hemisphere began tracking the object prior to the source hemisphere releasing the object, resulting in $\sim 1,000$ ms when both hemispheres actively tracked the object during the exchange. In contrast to the hard handoff observed for saccadic-based remapping, which is characterized by the abrupt truncation of activity in the source hemisphere once the target hemisphere picks up the attended object, these data support a soft-handoff model in which both hemispheres share object information for an extended period of time.

It is important to note that our statistical approach focuses on the relative differences between the two conditions, rather than absolute time points. As such, it is not surprising that the estimated time when the object was dropped falls several hundred milliseconds after crossing the midline, although the waveform appears to differ from zero prior to this time point. More observers would almost certainly result in earlier time estimates. However, there is no reason to assume that the overall pattern of pickup occurring before drop-off would change with more observers.

The pickup-before-drop-off results of experiment 1 might reflect an anatomical constraint based on the overlap of receptive fields from the two hemispheres at the vertical midline [15]. Once an attended object comes within a certain distance of the vertical meridian, receptive fields in both hemispheres might passively monitor the attended object without any need for active coordination. Experiment 2 tested this hypothesis by manipulating the predictability of the attended object crossing the vertical or horizontal meridian. If the handoff reflects receptive field overlap around the midline, only the physical position

A Soft Handoff of Attention

3




Figure 2. Source and Target Hemisphere Waveforms for Experiment 1

(A) Waveforms from the source and target hemispheres. These data are replotted from Figure 1B to emphasize hemisphere rather than object motion.

(B) Difference waveforms for the target and source hemispheres. Asterisks denote the point in time when the waveform reliably differed from zero for at least ten successive time windows for the first time. Error bars represent the SD. See also Figure S2.

of the object relative to the midline should affect handoff time. However, if the handoff reflects active coordination between the hemispheres, we might expect contextual factors such as predictability to come into play. In experiment 1, midline crossing was highly predictable, allowing the target hemisphere to prepare for the crossing in advance. If the crossing were less predictable, the handoff might be delayed. The *predictable* block of experiment 2 replicated experiment 1. Initial object motion was perfectly predictive of the final hemifield position. In the *unpredictable* block, on 50% of trials, all objects reversed trajectory back toward their starting position when they reached the vertical or horizontal midline (see Figure 3A). These trials were randomly interleaved with trials where the objects did not reverse directions thus making the movement in this block of the experiment more unpredictable than in the first block. For present purposes, we focused on trials in the unpredictable block that did not reverse trajectories. These trials were visually identical to the trials in the predictable block; they differed only in context. In order to assess the timing of the interhemispheric cooperation and focus on the unique impact that an attended object switching hemifields had on the waveforms, we again subtracted horizontal activity from matched vertical. For additional details and analysis on reversed trajectory trials, see the [Supplemental Results and Discussion](#).

The predictable block of experiment 2 replicated the basic effect from experiment 1: horizontal and vertical activity diverged *after* the attended object crossed the midline for the source hemisphere (2,381 ms; SD: 141 ms, 2.9° past midline) and *before* the attended object crossed the midline for the target hemisphere (1,222 ms; SD: 158 ms; 0.9° prior to midline; see Figure 3B). In the unpredictable block, the source hemisphere released object information at roughly the same relative time as in the predictable block (2,168 ms; SD: 293 ms; 2.2° past midline; see Figure 3C). However, the target hemisphere picked up the attended object information much later when the object's motion was unpredictable than when it was predictable. In the unpredictable block, by

our conservative timing measure, the tracked object was picked up after an average 2,068 ms (SD: 315 ms, 1.9° past midline). Our simple manipulation of motion predictability led to approximately an 800 ms delay in remapping the attended object to the target hemisphere (Z score = 2.58, $p = 0.010$). In contrast, the source hemisphere was unaffected: activity in this hemisphere continued to represent attended object information after it crossed the midline

for an equivalent duration in both cases (Z score = 0.56, $p = 0.576$; for converging evidence, see the [Supplemental Results and Discussion](#) and Figure S3). Further, while pickup time was significantly earlier than release time in the predictable condition (Z score = 3.5; $p < 0.001$), it was not significantly different in the unpredictable condition (Z score = 0.50; $p = 0.617$). If the shared object representation in experiment 1 reflected source and target hemispheres both representing a strip of the visual field straddling the midline, then motion predictability should not have influenced the time course of the handoff. Thus, while proximity to the midline certainly has an influence on object representation, we argue that the handoff is not an entirely hardwired consequence of receptive field overlap at the midline. For converging evidence that context influences the timing of the handoff using a fractional area latency measure, see the [Supplemental Results and Discussion](#).

These experiments reveal three core properties of the remapping process between hemispheres during covert attentional tracking of objects that travel between visual hemifields. First, similar to saccadic remapping, we find that remapping for predictable movement occurs *prospectively*, with the target hemisphere beginning to represent the attended object at least 300 ms prior to the object crossing the midline. These results echo those of Assad and Maunsell [16], who found that the spiking rate of neurons in the posterior parietal cortex increased 20–120 ms prior to the predictable occlusion of a tracked object. Second, our data suggest that attentional remapping approximates a soft handoff in which the source hemisphere holds on to object activity after the target hemisphere begins to represent the object, and well after the object has crossed the midline. This finding is broadly similar to the recent observation that unit activity for a newly attended object increases prior to the decrease in activity for the previously attended object [17]. Third, the nature of interhemispheric cooperation is context dependent. If the midline crossing is unpredictable, the target hemisphere does not reliably pick up object information until after the object crosses the midline. Thus, predictable object movement leads to a




Figure 3. Experiment 2 Trial Schematic and Difference Waveforms

(A) Schematic for horizontal trial movement in experiment 2. Vertical trials followed the same pattern.

(B) Difference waveforms for experiment 2. Unpredictable data are taken from the 50% of trials in which the objects crossed the midline. Asterisks denote the point in time when the waveform reliably differed from zero for at least ten successive time windows for the first time. Error bars represent the SD.

See also [Figure S3](#).

longer period of time when the two hemispheres appear to simultaneously represent object information.

We speculate that remapping is not specific to transsaccadic perception, but rather is a general property of the visual processing system. We often need to keep track of objects that we are not directly fixating. When a covertly attended object (such as a car in a rearview mirror) travels from one visual hemifield to another, it is necessary for object information to be sent from one cortical hemisphere to the other. The present results provide important initial steps toward characterizing how dynamic cooperation between the two hemispheres allows this computationally complex handoff to be completed. Perception of a stable visual world despite continuous changes to retinal inputs is a fundamental problem in cognitive neuroscience. The current work reveals that this stability involves complex cooperation between the two hemispheres to allow dynamic remapping of attended object representations. This interaction is sensitive to the predictability of object motion, with greater predictability producing more prospective tracking by the target cerebral hemisphere. The collaboration between the hemispheres may be responsible for stitching together the perceived visual world across two halves of the brain, and the paradigm outlined here is a promising method for studying this process.

Supplemental Information

Supplemental Information includes three figures, Supplemental Experimental Procedures, and Supplemental Results and Discussion and can be found with this article online at <http://dx.doi.org/10.1016/j.cub.2014.03.054>.

Acknowledgments

The authors would like to acknowledge funding support from NIH to T.D.: 1F32EB011959 and R01MH077105 to E.V. and R01MH65576 to T.S.H. and

J.M.W. E.V. would like to thank the Office of Naval Research (N00014-12-1-0972). This work was also partially funded by grants from ONR (N000141010278) and the National Geospatial Agency (HM0177-13-1-0001_P00001) to J.W.M. The authors would also like to thank Richard Matullo for assistance in running the experiments.

Received: September 25, 2013

Revised: February 7, 2014

Accepted: March 19, 2014

Published: April 24, 2014

References

- Drew, T., and Vogel, E.K. (2008). Neural measures of individual differences in selecting and tracking multiple moving objects. *J. Neurosci.* 28, 4183–4191.
- Vogel, E.K., and Machizawa, M.G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature* 428, 748–751.
- Drew, T., Horowitz, T.S., Wolfe, J.M., and Vogel, E.K. (2012). Neural measures of dynamic changes in attentive tracking load. *J. Cogn. Neurosci.* 24, 440–450.
- Drew, T., Horowitz, T.S., Wolfe, J.M., and Vogel, E.K. (2011). Delineating the neural signatures of tracking spatial position and working memory during attentive tracking. *J. Neurosci.* 31, 659–668.
- Cavanagh, P., Hunt, A.R., Afraz, A., and Rolfs, M. (2010). Visual stability based on remapping of attention pointers. *Trends Cogn. Sci.* 14, 147–153.
- Colby, C.L., and Goldberg, M.E. (1999). Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 22, 319–349.
- Duhamel, J.R., Colby, C.L., and Goldberg, M.E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255, 90–92.
- Golomb, J.D., Chun, M.M., and Mazer, J.A. (2008). The native coordinate system of spatial attention is retinotopic. *J. Neurosci.* 28, 10654–10662.
- Parks, N.A., and Corballis, P.M. (2008). Electrophysiological correlates of presaccadic remapping in humans. *Psychophysiology* 45, 776–783.
- Talsma, D., White, B.J., Mathôt, S., Munoz, D.P., and Theeuwes, J. (2013). A retinotopic attentional trace after saccadic eye movements: evidence from event-related potentials. *J. Cogn. Neurosci.* 25, 1563–1577.

A Soft Handoff of Attention

5

11. Berman, R., and Colby, C. (2009). Attention and active vision. *Vision Res.* 49, 1233–1248.
12. Ekiz, N., Salih, T., Kucukoner, S., and Fidanboyulu, K. (2005). An overview of handoff techniques in cellular networks. *Int. J. Inf. Technol.* 2, 132–136.
13. Luck, S. (2005). *An Introduction to the Event-Related Potential Technique*, First Edition (Cambridge: MIT press).
14. Efron, B., and Tibshirani, R. (1994). *An Introduction to the Bootstrap* (Boca Raton: CRC Press).
15. Hubel, D.H., and Wiesel, T.N. (1967). Cortical and callosal connections concerned with the vertical meridian of visual fields in the cat. *J. Neurophysiol.* 30, 1561–1573.
16. Assad, J.A., and Maunsell, J.H. (1995). Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature* 373, 518–521.
17. Khayat, P.S., Spekreijse, H., and Roelfsema, P.R. (2006). Attention lights up new object representations before the old ones fade away. *J. Neurosci.* 26, 138–142.

The Invisible Gorilla Strikes Again: Sustained Inattentional Blindness in Expert Observers

Psychological Science
XX(X) 1–6
© The Author(s) 2013
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0956797613479386
pss.sagepub.com


Trafton Drew, Melissa L.-H. Võ, and Jeremy M. Wolfe

Visual Attention Lab, Harvard Medical School, and Brigham and Women's Hospital, Boston, Massachusetts

Abstract

Researchers have shown that people often miss the occurrence of an unexpected yet salient event if they are engaged in a different task, a phenomenon known as *inattentional blindness*. However, demonstrations of inattentional blindness have typically involved naive observers engaged in an unfamiliar task. What about expert searchers who have spent years honing their ability to detect small abnormalities in specific types of images? We asked 24 radiologists to perform a familiar lung-nodule detection task. A gorilla, 48 times the size of the average nodule, was inserted in the last case that was presented. Eighty-three percent of the radiologists did not see the gorilla. Eye tracking revealed that the majority of those who missed the gorilla looked directly at its location. Thus, even expert searchers, operating in their domain of expertise, are vulnerable to inattentional blindness.

Keywords

visual attention, perception, selective attention

Received 8/1/12; Revision accepted 1/27/13

When one is engaged in a demanding task, attention can act like a set of blinders, making it possible for salient stimuli to pass unnoticed right in front of one's eyes (Neisser & Becklen, 1975). This phenomenon of sustained *inattentional blindness* (IB) is best known from Simons and Chabris's (1999) study in which observers attended to a ball-passing game while a human in a gorilla suit wandered through the field of play. Even though the gorilla walked through the center of the scene, a substantial portion of the observers did not report seeing it (the video can be viewed at <http://www.theinvisiblegorilla.com/videos.html>). Moving beyond such demonstrations, one might ask whether IB still occurs when the observers are experts who are highly trained on the primary task. There is some evidence that expertise mitigates the effect. For example, Memmert (2006) found a decreased rate of IB among basketball players who were asked to count the number of passes in an artificial basketball game. However, when Potchen (2006) asked radiologists to review cases as if for an annual exam and showed them chest x-rays with a clavicle (collarbone) removed, roughly 60% failed to notice the missing bone. Finally, a recent observational report

documented that a misplaced femoral line was not detected by a variety of health-care professionals who evaluated the case (Lum, Fairbanks, Pennington, & Zwemer, 2005).

Both of these instances of apparent IB in the medical setting occurred when single-slice medical images were viewed. Modern medical imaging technologies, such as MRI, computed tomography (CT), and positron-emission tomography (PET), are increasingly complex: The single image of a chest x-ray has been replaced with hundreds of slices in a chest CT scan. It is therefore important to study whether IB occurs in these modern imaging modalities. These situations are interesting because the observer actively interacts with the stimulus—for example, scrolling through a stack of images of the lung. This degree of control may ameliorate the effects of IB because the searcher is able to return to and further examine any images that appear unusual.

Corresponding Author:

Trafton Drew, Visual Attention Lab, Harvard Medical School, 64 Sidney St., Suite 170, Cambridge, MA 02139
E-mail: tdrew1@partners.org

Moreover, whereas Potchen (2006) showed that radiologists could miss the unexpected absence of a stimulus, we wanted to know if they would miss the presence of a readily detectable, highly anomalous item while performing a task within their realm of expertise. In an homage to Simons and Chabris's (1999) study, we made that item a gorilla. We compared the performance of radiologists with that of naive observers.

Design and Procedure

In CT lung-cancer screening, radiologists search a reconstructed "stack" of axial slices of the lung for nodules that appear as small light circles (Aberle et al., 2011). In Experiment 1, 24 radiologists (mean age = 48, range = 28–70) had up to 3 min to freely scroll through each of five chest CTs, searching for nodules as we tracked their eye position. The five trials contained an average of 10 nodules, and the observers were instructed to click on nodule locations with the computer mouse. In the final trial, we inserted a gorilla with a white outline into the lung (see Fig. 1). A typical stack of images from a chest CT contains 100 to 500 slices. In the current study, the stack that contained the gorilla had 239 slices.

Nine radiologists were tested at Brigham and Women's Hospital in Boston, Massachusetts, and 15 expert examiners from the American Board of Radiology were tested at a meeting of that organization in Louisville, Kentucky. The gorilla measured 29×50 mm. Because of equipment differences, the image size was slightly different at the two sites, and consequently the size of the gorilla differed slightly (Boston: 0.9×0.5 degrees of visual angle; Louisville: 1.3×0.65 degrees of visual angle). To avoid large onset transients, we had the gorilla fade into and out of visibility over five 2-mm-thick slices (Fig. 1). The total volume of the rectangular box that could hold the gorilla was more than $7,400 \text{ mm}^3$, roughly the size of a

box of matches. The gorilla was centered in depth near a lung nodule such that both were clearly visible when the gorilla was at maximum opacity. That is, if someone pointed at the correct location in the static image and asked you, "What is that?" you would have no trouble answering, "That is a gorilla." In the scans used in this study, which were taken from the Lung Image Database Consortium (Armato et al., 2011), the average volume of the lung nodules was 153 mm^3 . Thus, the gorilla was more than 48 times the size of the average nodule in the images (see Fig. 2a).

Experiment 2 replicated Experiment 1 with 25 naive observers (mean age = 33.7, range = 19–55), who had no medical training. Prior to the experiment, the experimenter spent roughly 10 min teaching these observers how to identify lung nodules. This experiment began with a practice trial, during which the experimenter took time to point out several nodules. The experimenter then encouraged the observer to try to find nodules on his or her own. Once the observer was able to detect at least one nodule, the practice trial was concluded, and the experimental trials began. As in Experiment 1, a subset (12) of observers completed the study on a slightly smaller screen. We observed no difference in gorilla or nodule detection as a result of equipment differences.

Experiment 3 was a control experiment intended to ensure that the gorilla was, in fact, visible. Twelve naive observers (mean age = 37.3, range = 21–54) were shown movies that progressed from the top to the bottom of the same chest CT case that was used as the final trial in Experiments 1 and 2. The gorilla was inserted into the movies in the same location on 50% of the 20 trials, and observers were asked to judge whether the gorilla was present or absent on each trial. A circular cue indicated the possible location of the gorilla on each trial. The movies were presented at a rate of 35 or 70 ms per frame (manipulated within subjects).

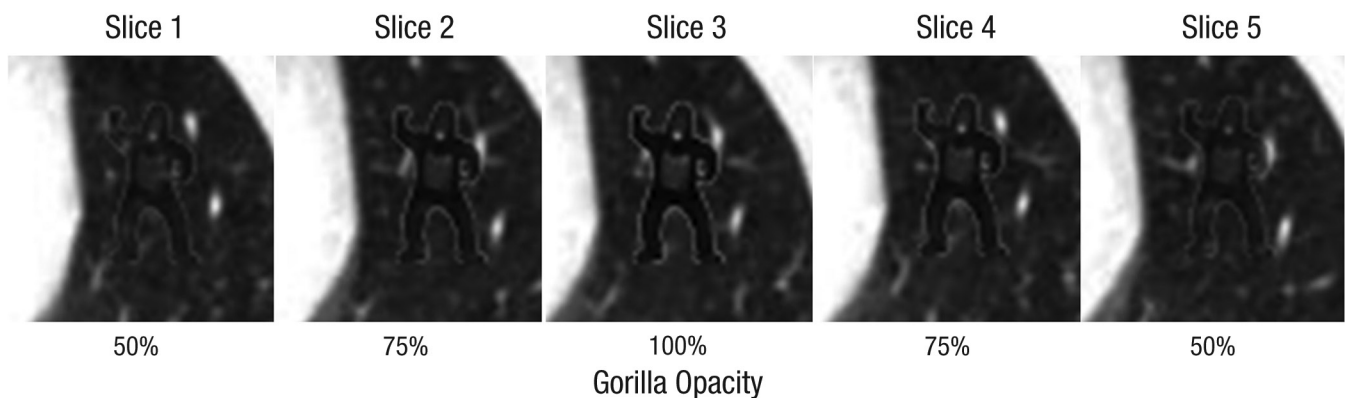


Fig. 1. Illustration of the slices showing the gorilla in the final trial of Experiments 1 and 2. The opacity of the gorilla increased from 50% to 100% and then decreased back down to 50% over the course of 5 slices within a stack of 239.




Fig. 2. Computed-tomography image containing the embedded gorilla (a) and eye-position plot of a radiologist who did not report seeing the gorilla (b). In (b), the circles represent eye positions recorded at 1-ms intervals.

Results

Experiment 1

The nodule detection task was challenging, even for expert radiologists. The overall nodule detection rate was 55%. While engaged in this task, the radiologists freely scrolled through the slices containing the gorilla an average of 4.3 times. At the end of the final case, we asked a series of questions to determine whether they had noticed the gorilla: “Did the final trial seem any different than any of the other trials?” “Did you notice anything unusual on the final trial?” and, finally, “Did you see a gorilla on the

final trial?” Twenty of the 24 radiologists failed to report seeing a gorilla. This was not due to the gorilla being difficult to perceive: All 24 radiologists reported seeing the gorilla when they were asked if they noticed anything unusual in Figure 1 after completing the experiment (see also the results for Experiment 3).

The radiologists had ample opportunity to find the gorilla. On average, those who missed the gorilla spent 5.8 s viewing the five slices containing it (range = 1.1–12 s) and spent an average of 329 ms looking at the gorilla’s location. Furthermore, eye tracking revealed that of the 20 radiologists who did not report the gorilla, 12 looked directly at the gorilla’s location when it was visible. The mean dwell time on the gorilla in this group was 547 ms. Figure 2b shows the eye positions of a radiologist who clearly fixated the gorilla but did not report it.

Experiment 2

None of the 25 naive observers reported noticing the gorilla. As was the case with the radiologists in Experiment 1, all of the naive observers reported seeing the gorilla when shown Figure 1. The results support the idea (Memmert, 2006) that experts are somewhat less prone to IB than novices are (Fisher’s exact test: $p = .0497$; see Fig. 3a). However, unlike in Memmert’s study, our two groups showed a sizable difference in performance on the primary task. As expected, radiologists were much better at detecting lung nodules (mean detection rate = 55%) than were naive observers (12%), $t(47) = 12.3$, $p < .001$ (see Fig. 3b).

Eye movement data followed the pattern seen with the radiologists. The naive observers spent an average of 4.9 s searching the frames in which the gorilla was visible and an average of 157 ms looking at the gorilla’s location. Although both measures showed that radiologists who missed the gorilla spent slightly more time searching in its vicinity than did the naive observers, neither difference was significant, $t(43) = 1.26$, $p = .22$, and $t(43) = 1.23$, $p = .22$, respectively. Of the 25 naive observers, 9 looked at the gorilla’s location. The mean dwell time on the gorilla in the latter group was 435 ms.

Experiment 3

Although all observers in Experiments 1 and 2 reported seeing the gorilla when shown Figure 1 at the end of the experiment, given the very high rate of IB in both studies, there was some concern that the gorilla was too difficult to detect when embedded within a stack of chest CT images. We tested this possibility in Experiment 3. The movies played at a fast or slower frame rate such that the gorilla was visible for 175 or 350 ms, respectively—substantially less time than the 4.9 s that the average

naive observer in Experiment 2 spent searching frames in which the gorilla was present. Despite this large difference, performance on the detection task was near ceiling (88% correct). Accuracy was not affected by the frame rate, $t(11) = 1.1$, $p = .18$ (see Fig. 3c).

Discussion

In Experiment 1, 20 of 24 expert radiologists failed to note a gorilla, the size of a matchbook, embedded in a stack of CT images of the lungs. This is a clear illustration that radiologists, though they are expert searchers, are not immune to the effects of IB even when searching medical images within their domain of expertise. Potchen (2006) showed that radiologists could miss the absence of an entire bone. Results from laboratory search tasks have shown that it is harder to detect the absence of something than to detect its presence (Treisman & Souther, 1985). Our data show that under certain circumstances, experts can also miss the presence of a large,

anomalous stimulus. In fact, there is some clinical evidence for errors of this sort in radiology. Lum et al. (2005) reported a case study in which multiple emergency radiologists failed to detect a misplaced femoral-line guide wire that was mistakenly left in a patient and was clearly visible on three different chest CT scans. Although these scans were viewed by radiologists, emergency physicians, internists, and intensivists, the guide wire was not detected for 5 days. Clearly, radiologists can miss an abnormality that is retrospectively visible when the abnormality is unexpected.

It is reassuring that our experts exhibited somewhat lower rates of IB than naive observers, as was reported by Memmert (2006). In that earlier study, expertise was defined as extensive basketball experience, and IB was measured during an artificial task in which two groups of individuals passed a ball back and forth while moving randomly about a small area. The observers were asked to count the number of passes completed by one group. In this rather abnormal basketball game, the rate of IB

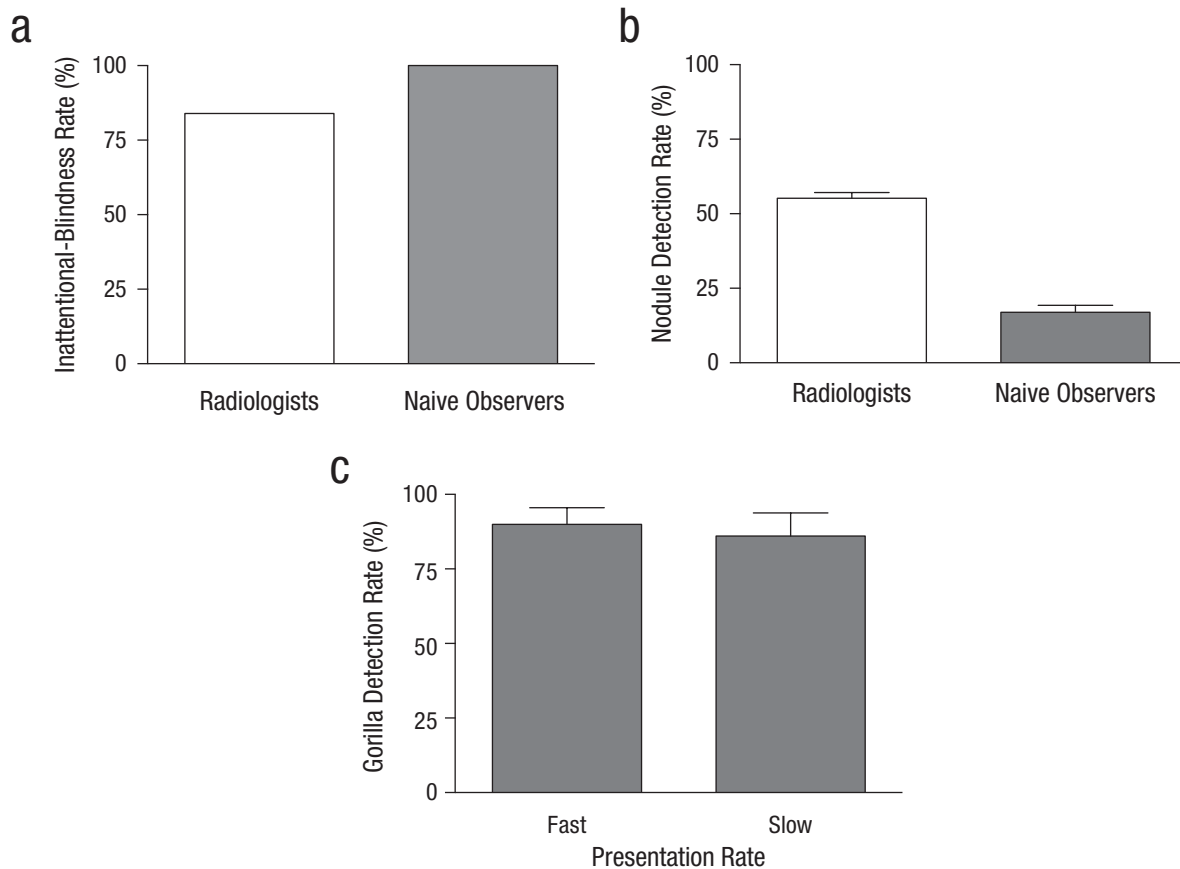


Fig. 3. Experimental results. The graph in (a) shows the rate of inattentional blindness (i.e., the percentage of observers who did not report seeing the gorilla) among the radiologists in Experiment 1 and the naive observers in Experiment 2. The graph in (b) shows the percentage of nodules that were correctly marked by these same observers. The graph in (c) shows the rate at which observers in Experiment 3 detected the gorilla as a function of presentation rate (fast: 35 ms/frame; slow: 70 ms/frame). Error bars represent standard errors of the mean.

was lower for the experts than for observers with less basketball experience. In the current study, high rates of IB were obtained with a task and stimulus materials that were very familiar to our expert observers: searching a chest CT scan for signs of lung cancer.

Experts may perform slightly better on this IB task than naive observers do because their attentional capacity is less completely occupied by the primary task. Simons and Jensen (2009) recently showed that the rate of IB decreased when the primary task (counting the number of times an object bounced) was made easier. Along similar lines, there is evidence that training on a specific task reduces the subsequent IB rate (Richards, Hannon, & Derakshan, 2010). In our study, the radiologists certainly had much more experience on the specific primary task, and were clearly better at it. Both factors are likely to have contributed to the reduced rate of IB observed in our experts. Nevertheless, even though the radiologists were slightly better than the naive observers, their miss rate of 83% indicates a striking level of IB.

Why do radiologists sometimes fail to detect such large anomalies? Of course, as is critical in all IB demonstrations, the radiologists were not looking for the unexpected stimulus. In most previous demonstrations of IB, observers engaged in a primary task that was unrelated to detection of the unexpected stimulus (e.g., counting the number of passes or bounces, as in Most et al., 2001; Richards et al., 2010; Simons & Chabris, 1999; Simons & Jensen, 2009). Here, too, though detection of aberrant structures in the lung would be a standard component of the radiologist's task, observers were not looking for gorillas. Presumably, they would have done much better at detecting the gorilla had they been told to be prepared for such a target. Moreover, the observers were searching for small, light nodules. Previous work with naive observers has shown that IB is modulated by the degree of match between the designated targets and the unexpected item (Most et al., 2001). This suggests that our observers might have fared better if we had used an albino gorilla that better matched the luminance polarity of the designated targets. Counterintuitively, perhaps a smaller gorilla would have been more frequently detected because it would have more closely matched the size of the lung nodules.

Our results could be seen as an example of a phenomenon known as *satisfaction of search*, in which detection of one stimulus interferes with detection of subsequent stimuli (e.g., Berbaum et al., 1998). We placed the gorilla on a slice that contained a nodule that was detected by 71% of the radiologists. Perhaps the observed rate of IB was inflated by the presence of this nodule. Without running an additional experiment examining the detection rate for the gorilla in the absence of the nodule, it

is difficult to be certain what role the presence of the nodule played. However, if satisfaction of search truly drove the IB effect, we would expect that radiologists who missed the nodule would have been more likely to detect the gorilla and that radiologists who found the nodule would have been more likely to miss the gorilla. Neither of these predictions held true: Of the 7 radiologists who missed the nodule, none detected the gorilla. Furthermore, all of the radiologists who detected the gorilla also detected the nodule on the same slice.

It would be a mistake to regard these results as an indictment of radiologists. As a group, they are highly skilled practitioners of a very demanding class of visual search tasks. The message of the present set of results is that even this high level of expertise does not immunize individuals against inherent limitations of human attention and perception. Researchers should seek better understanding of these limits, so that medical and other man-made search tasks could be designed in ways that reduce the consequences of these limitations.

Author Contributions

T. Drew developed the study concept. All authors contributed to the study design. T. Drew collected and analyzed the data. T. Drew wrote the manuscript in collaboration with J. M. Wolfe and M. L.-H. Võ. All authors approved the final version of the manuscript for submission.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

References

- Aberle, D. R., Adams, A. M., Berg, C. D., Black, W. C., Clapp, J. D., Fagerstrom, R. M., . . . Sicks, J. D. (2011). Reduced lung-cancer mortality with low-dose computed tomographic screening. *New England Journal of Medicine*, *365*, 395–409.
- Armato, S. G., III, McLennan, G., Bidaut, L., McNitt-Gray, M. F., Meyer, C. R., Reeves, A. P., . . . Croft, B. Y. (2011). The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A completed reference database of lung nodules on CT scans. *Medical Physics*, *38*, 915–931.
- Berbaum, K. S., Franken, E. A., Dorfman, D. D., Miller, E. M., Caldwell, R. T., Kuehn, D. M., & Berbaum, M. L. (1998). Role of faulty visual search in the satisfaction of search effect in chest radiography. *Academic Radiology*, *5*, 9–19.
- Lum, T. E., Fairbanks, R. J., Pennington, E. C., & Zwemer, F. L. (2005). Profiles in patient safety: Misplaced femoral line guidewire and multiple failures to detect the foreign body on chest radiography. *Academic Emergency Medicine*, *12*, 658–662.

- Memmert, D. (2006). The effects of eye movements, age, and expertise on inattention blindness. *Consciousness and Cognition, 15*, 620–627.
- Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattention blindness. *Psychological Science, 12*, 9–17.
- Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology, 7*, 480–494.
- Potchen, E. J. (2006). Measuring observer performance in chest radiology: Some experiences. *Journal of the American College of Radiology, 3*, 423–432.
- Richards, A., Hannon, E. M., & Derakshan, N. (2010). Predicting and manipulating the incidence of inattention blindness. *Psychological Research, 74*, 513–523.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception, 28*, 1059–1074.
- Simons, D. J., & Jensen, M. S. (2009). The effects of individual differences and task difficulty on inattention blindness. *Psychonomic Bulletin & Review, 16*, 398–403.
- Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General, 114*, 285–310.