

Accessibility Limits Recall From Visual Working Memory

Jason Rajsic
University of Toronto

Garrett Swan
Pennsylvania State University

Daryl E. Wilson
Queen's University

Jay Pratt
University of Toronto

In this article, we demonstrate limitations of accessibility of information in visual working memory (VWM). Recently, cued-recall has been used to estimate the fidelity of information in VWM, where the feature of a cued object is reproduced from memory (Bays, Catalao, & Husain, 2009; Wilken & Ma, 2004; Zhang & Luck, 2008). Response error in these tasks has been largely studied with respect to failures of encoding and maintenance; however, the retrieval operations used in these tasks remain poorly understood. By varying the number and type of object features provided as a cue in a visual delayed-estimation paradigm, we directly assess the nature of retrieval errors in delayed estimation from VWM. Our results demonstrate that providing additional object features in a single cue reliably improves recall, largely by reducing swap, or misbinding, responses. In addition, performance simulations using the binding pool model (Swan & Wyble, 2014) were able to mimic this pattern of performance across a large span of parameter combinations, demonstrating that the binding pool provides a possible mechanism underlying this pattern of results that is not merely a symptom of one particular parametrization. We conclude that accessing visual working memory is a noisy process, and can lead to errors over and above those of encoding and maintenance limitations.

Keywords: visual working memory, visual short-term memory, memory retrieval, computational models

Although our subjective visual experience is rich with details, our ability to recall visual information from the recent past is surprisingly poor (O'Regan & Noë, 2001). The systems and processes that allow us to retain visual information for brief periods are referred to as visual working memory (VWM; Luck, 2008; Postle, 2006). Although much consideration has been given to the limitations of encoding and maintenance in VWM, few systematic examinations have addressed potential limitations of retrieval in VWM, that is, how information in VWM is accessed. The seminal studies of VWM have largely relied on the one-shot change detection technique, where a one-to-one comparison of all information in a display to all information in memory is all that is theoretically necessary to determine a response (Luck & Vogel, 1997; Wheeler & Treisman, 2002). Indeed, Hyun, Woodman, Vogel, Hollingworth, and Luck (2009) have demonstrated that changes between remembered and test displays “pop out” of the display, and quickly attract spatial attention, suggesting that the comparison of remembered and tested objects in change detection

occurs in parallel. However, even in simple change detection, providing a single object at test instead of the entire studied object array results in a performance cost (Jiang, Olson, & Chun, 2000). Such performance costs cannot be attributed to failures in encoding or maintaining visual information over time, and thus provide evidence that the processes that retrieve information from VWM can lead to failures of memory.

Motivated by the goal of determining the type of resources that limit VWM, vision researchers have adopted a new laboratory task for measuring the quality of information in VWM: the delayed-estimation task (Bays, Catalao, & Husain, 2009; Wilken & Ma, 2004; Zhang & Luck, 2008). In the delayed-estimation task, participants study an array of objects, and at test they are provided with a cue to one of the studied objects (usually a cue to its location) so that they can fill in missing information about that object (e.g., its color). Much of the work using this task has sought to uncover the model that best accounts for changes in the shape of the empirical memory error distribution (for a review, see van den Berg, Awh, & Ma, 2014) in order to settle the debate about the nature of representation in VWM. Although the influence of encoding and maintenance on memory error in delayed estimation has been examined through the manipulation of stimulus exposure duration (Zhang & Luck, 2008), presentation format (simultaneous vs. sequential; Emrich & Ferber, 2012; Gorgoraptis, Catalao, Bays, & Husain, 2011), the retro-cuing technique (Murray, Nobre, Clark, Cravo, & Stokes, 2013), and retention interval duration (Zhang & Luck, 2009), little research has attempted to isolate the contribution of selective retrieval processes to memory error. Because the delayed-estimation paradigm is a cued-recall task, memory failures

This article was published Online First March 2, 2017.

Jason Rajsic, Department of Psychology, University of Toronto; Garrett Swan, Department of Psychology, Pennsylvania State University; Daryl E. Wilson, Department of Psychology, Queen's University; Jay Pratt, Department of Psychology, University of Toronto.

Correspondence concerning this article should be addressed to Jason Rajsic, Department of Psychology, University of Toronto, 100 St. George Street, Toronto, Ontario, Canada, M5S 3G3. E-mail: jason.rajsic@mail.utoronto.ca

may originate from two sources: failures of availability and failures of accessibility (Tulving & Pearlstone, 1966). Whereas availability failures occur when a cued memory was not encoded or stored, an accessibility failure occurs when, despite being encoded and stored, the cued memory is not sufficiently activated by recall cues. It can be difficult to establish that a memory is unavailable rather than inaccessible, as an absence of evidence is not evidence of absence. On the other hand, establishing inaccessibility is possible by demonstrating a reliable memory performance gain with a particular cue. This is the primary concern of the present article: whether manipulating the characteristics of memory probes in a VWM task will reveal accessibility limits in the delayed estimation of visual objects.

While little data exists regarding the possibility of accessibility limits in the delayed estimation of visual objects, the broader working memory (WM) literature includes demonstrations of the importance of retrieval. McElree (2001) has reported that the retrieval efficacy (as assessed by speed and accuracy trade-off functions) of matching judgments decreases as more items are maintained in WM. In addition, Oberauer (2002) has shown that computations performed using items held in WM are slowed when the item being accessed changes from one trial to the next. Both authors have suggested that accessing information in working memory requires bringing a representation into the focus of attention. Investigations of VWM using change detection have shown that spatial rearrangement of stimuli, as well as removal of non-tested items, in probe displays disrupts the recognition of changes (Jiang, Olson, & Chun, 2000), suggesting that spatial correspondence is an important determinant of successful information retrieval. Finally, in the detection of changes to realistic scenes, Hollingworth (2003) has shown that spatial cues directing participants to the location of a possible change improve change detection, thus demonstrating the need to consider how retrieval of information from visual memory determines successful performance.

The tasks used in these cases are, however, notably different from the delayed-estimation task used to assess VWM, limiting their generalizability. In principle, however, the delayed-estimation task requires selective reporting of one of multiple objects, often with multiple features (e.g., Fournie & Alvarez, 2011), which would require selecting among candidate memory representations. Relatedly, Flombaum and colleagues (Bae & Flombaum, 2013; Levillain & Flombaum, 2012) have shown that task-irrelevant featural overlap between objects can lead to correspondence errors; if objects differ on features that are integral to those being reported (e.g., objects of different hues in a context where luminance memory is tested), decrements in memory precision can be eliminated. The authors argued that reducing correspondence problems led to this improvement in performance, although it is not clear what stage, or stages, of memory were affected by their stimulus manipulation (see also Bays, Catalao, & Husain, 2009). Some support for a retrieval-based locus of correspondence problems can be found in Rajsic and Wilson (2014), who showed that the presence of nontarget items at test substantially reduces swap errors, analogous to Jiang, Olson, and Chun's (2000) observation in change detection. Still, the processes by which the selective reports in delayed-estimation tasks are made remains poorly understood and may constitute an additional source

of variability to memory reports that is worth capturing in models of VWM.

In order to uncover memory retrieval processes involved in delayed estimation from VWM, we conducted three experiments wherein we provided identical encoding and maintenance conditions within and across experiments, but adjusted the information provided by the recall cues on each trial. In every experiment, participants saw objects composed of two features—a color and an orientation—that appeared in varying locations. This meant that every to-be-remembered stimulus was defined by values along three dimensions: a location, color, and orientation. In each experiment, participants consistently recalled one of these three features (e.g., color), and the two remaining features (e.g., location and orientation) were used as retrieval cues. A retrieval cue could provide the feature value of an object along the first, second, or both cue dimensions. For example, in Figure 1, the recalled feature is orientation in all trials, but a given trial's retrieval cue might include only color, only location, or both color and location information. We hypothesized that VWM representations are accessed by matching representations in a probe display to representations stored in VWM. This leads to the prediction that the more features contained in a memory probe, the more likely participants would be to report the probed item. In the case when only one feature was presented in the memory probe, multiple representations might be activated by the memory probe, leading to swap errors, in the case that the activation process has a low threshold, or even guess errors, if the activation process has a high threshold, such that one representation must be activated considerably over others before a memory-guided response is made. In summary, we expected that VWM performance would indeed be limited by accessibility, and that performance would be maximized by memory probes with more object features. While intuitive—indeed, such a retrieval process is implicit in studies of VWM using delayed estimation—the question of how retrieval occurs in VWM is empirical, and our study provides insight into how this memory operation functions.




Figure 1. A sample trial, depicting the report-orientation variation (Experiment 1). Stimuli not drawn to scale. On the right, the top row depicts a color cue trial, the middle row depicts a location cue trial, and the bottom row depicts a both-feature cue trial. Report feedback was presented as a dot indicating where on the report-circle a correct click should have occurred. See the online article for the color version of this figure.

Experiment 1

In this experiment, we assessed the contribution of color and location used as cues to recall orientation of simple objects (triangles). Participants reported the orientation of a recently encoded triangle when provided with a color cue, a location cue, or a cue providing both the color and location of the target triangle. If accessibility limits the information that can be retrieved from VWM, then providing both-feature cues should improve performance, increasing the probability of reporting the cued orientation, and reducing the likelihood of reporting a noncued object's orientation.

Methods

Participants. Thirty participants in total were recruited for this experiment. All participants were students in a first-year undergraduate Psychology course at the University of Toronto, participating for course credit. Participants provided informed consent before participating. Fifteen participants completed a version of this experiment where the to-be-remembered stimuli were presented for 100 ms on each trial, and 15 completed a version of the experiment where the to-be-remembered stimuli were presented for 600 ms on each trial. This sample size was maintained for Experiments 2 and 3.

Materials and procedure. Stimuli were constructed and presented using Matlab by Mathworks using the Psychophysics toolbox version 3.0.11 (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997). Stimuli were displayed on 16 in. CRT monitors at a viewing distance of approximately 50 cm. To ensure consistent stimulus exposure, participants viewed stimuli using a chin rest. The experiment was conducted in a dimly lit, sound-attenuated room. Each experimental session consisted of 512 trials, with two distinct stages: the encoding stage and the test stage. The encoding stage was identical for all experiments reported in this article, and so will be described only here for economy.

The encoding stage consisted of a 1.5-s fixation display, consisting of a single white fixation cross on a gray background. The memory sample display occurred next, consisting of either two or five colored triangles, appearing approximately 6.5° from fixation. The triangles were isosceles in shape, with a base of approximately 1.25° and a height of approximately 2.5° . Each triangle was hollowed, to allow for discriminability despite occasional partial overlap, and the thickness of each triangle's contour was approximately 0.25° . Each triangle was pseudorandomly rotated around its center (defined as the point lying halfway between its short side and opposite vertex) by selecting an angular value for each triangle in a given trial's display from between 0 and 358° , in two-degree steps without replacement. Triangle colors and locations were determined using an identical angular sampling approach. For color, angular values were translated into RGB values by converting from the L^*a^*b space, using the angles to select a point in L^*a^*b space on the radius of a circle centered at [70, 0, 0], with a radius of 60. Although the luminance value was chosen to equate color luminance, variation in measured luminance did exist, and so color memory in our experiments may have included some degree of memory for luminance as well. For location, angular values were translated into screen positions by centering a triangle on a point on an imaginary circle of radius 6.6° around the fixation cross. The memory display was removed after either 100 ms (for

15 participants) or 600 ms (for a separate 15 participants). Following the offset of the memory display, a 900-ms retention interval of a blank screen with a fixation cross was presented.

Following the retention interval was the test stage of the trial. In Experiment 1, the test stage was one of three types: color cue, location cue, or both cue. For color cue trials, a single circle outline, with a 1° diameter and a line width of 0.25° , appeared in the center of the screen whose color exactly matched one of the triangles that had appeared in the display earlier.

For location cue trials, a single, white circle outline appeared centered on the exact location of one of the triangles that had appeared in the presentation stage. For both cue trials, a single circle outline appeared whose location and color exactly matched one of the triangles from earlier in the trial. In addition, a large, white circle outline was drawn on-screen, centered on fixation, with a radius of approximately 8.25° and a thickness of 0.35° . This was added in order to visually equate the test display in Experiment 1 with the test display of Experiment 2, where this circle was drawn as a color wheel of identical physical dimensions. In all three conditions, the participant used the mouse to produce an oriented triangle whose orientation matched his or her memory of the cued object. The mouse cursor was always set to the center of the screen at the beginning of the test phase, and when the cursor was moved at least 5° away from fixation, the cue circle was replaced by a triangle whose orientation was calculated using the angle of arc between the mouse cursor's position and the center of the screen. Participants, submitted their matching response by clicking the mouse button. After a response was given, feedback was provided in the form of a small, black, filled circle of radius 0.16° on the larger circle, whose radial angle from fixation matched the correct orientation of the cued triangle.

Across all experiments, both factors (Set Size and Cue Condition) were randomly and equally seeded, leading to an approximately equivalent, with small variation, number of trials per cell of the design. Participants completed 512 trials across 8 blocks in one experimental session. One group of 15 participants were shown the triangles at encoding for 100 ms while another group of 15 was shown the triangles for 600 ms. Two sample durations were used as Rajsic and Wilson (2014) found a retrieval-context effect for a nonspatial feature (color) only when stimuli had been presented for 600 ms, but not 100 ms. Thus, we anticipated a possible interaction between Cue Condition and Sample Duration.

Results

Raw memory error. We first analyzed raw error, calculated as the mean absolute error between the probed item's orientation and its reported orientation, in degrees. Raw memory error in each condition can be seen in Figure 2. A mixed-model ANOVA with Set Size (2, 5) and Cue Condition (Color Cue, Location Cue, or Both Cue) as within-subjects factors and Sample Duration (100 ms, 600 ms) as a between-subjects factor showed that increasing Set Size increased memory error, $F(1, 28) = 961.69$, $p < .001$, $\eta_p^2 = .97$, and that Cue Condition also modulated memory error, $F(2, 56) = 6.60$, $p = .003$, $\eta_p^2 = .19$. Overall, memory error was lower when both features were present in a cue than when either color alone, $F(1, 28) = 17.14$, $p < .001$, $\eta_p^2 = 0.38$, or location alone, $F(1, 28) = 4.95$, $p = .03$, $\eta_p^2 = .15$, was present. Cue Condition did not interact with either Set Size or Sample Duration.




Figure 2. Raw memory error (mean absolute deviation) in Experiment 1. Error bars depict one within-subjects standard error.

The main effect of Cue Condition shows that access to VWM was improved (memory error was lower) when more informative cues were provided.

Three-component model analysis. Given that memory cues did affect the amount of memory error in our experiment, we used Bays's three-component model (Bays, Catalao, & Husain, 2009; also referred to as the "swap" model: Suchow, Brady, Fournie, & Alvarez, 2013) to understand the source of this change in error. This model estimates four performance descriptors (one redundant, hence the term "three-component model") from the trial-wise list of response errors and stimulus values: the precision of memory, the probability of correct access,¹ the probability of a swap response, and the probability of a guess response. The three latter parameters describe the three possible sources of any given response: a distribution of responses from a correctly accessed item, where the reported value is sampled from a circular normal distribution (the von Mises distribution) centered on the cued feature value; a distribution of "swap" responses, where the reported value is sampled from a combination of circular normal distributions centered on the feature values of the nontarget items that had been presented in the memory display; and a distribution of "guess" responses, where the reported value is sampled from a uniform distribution, meaning that every feature value is equally likely to be reported. Importantly, memory precision can be quantified using the standard deviation of the circular normal distributions for both the "correct" distributions and the "swap" distributions. Parameters are estimated using maximum likelihood. In our analyses, we fit parameters separately in each condition for each participant. Although we endeavored to maximize the number of trials in each condition for the purposes of parameter fitting, to keep each experimental session at approximately one hour in length, we were able to collect approximately 85 observations per condition. Lawrence (2010) found relatively modest gains in the reliable recovery of $p(\text{Correct Access})$ between 80 samples and 160 samples per fit, albeit using simulations and fits with a two-component model of memory (correct responses and guesses from Zhang & Luck, 2008). Nevertheless, it is possible that parameter estimation suf-

fered from noise due to a modest number of trials, and so these results—as well as those from Experiments 2 and 3—should be interpreted with some discretion.

Given that our analyses of raw memory error showed only main effects of Set Size and Cue Condition, we ran two-way repeated measures ANOVAs on each set of estimated memory parameters, using only Set Size and Cue Condition as factors, and concentrating exclusively on the source of the main effect of Cue Condition found in raw memory error. The resulting parameter estimates are plotted in Figure 3. Although Set Size affected all memory parameters, $F_s > 19.07$, $p_s < .001$, only the probability of a correct response [or $p(\text{Correct Access})$], $F(2, 58) = 9.75$, $p = .001$, $\eta_p^2 = 0.25$, and the probability of a swap [or $p(\text{Swap})$], $F(2, 58) = 14.94$, $p < .001$, $\eta_p^2 = 0.34$, were affected by memory cues. Compared to both-feature cues, color cues and location cues alone led to a lower probability of correct responses, $F_s(1, 29) > 5.54$, $p_s < .026$, $\eta_p^2 > 0.16$, and a higher probability of swap responses, $F_s(1, 29) > 6.47$, $p_s < .017$, $\eta_p^2 > 0.18$. On the basis of these findings, the benefit of multifeature retrieval cues can be characterized as an improvement in memory disambiguation; some swaps that occurred when only one feature was available in the cue were due to selection of the wrong remembered item, when the correct remembered item was actually available to be reported.

In addition to a main effect of Cue Condition, we also observed interactions between Cue Condition and Set Size for the $p(\text{Swap})$, $p(\text{Correct Access})$, and the circular Standard Deviation of correct responses (SD), indicating that the effect of memory cues differed by Set Size. Given that the purpose of our study was to understand the source of the cue-related main effect found in raw memory error, we do not report these statistics here. However, curious readers can find the details of these interactions in the Appendix.

Discussion

The results of Experiment 1 demonstrate that increasing the amount of information provided by the cue can allow participants to correctly recall an object's orientation more often. Providing two retrieval features allowed participants to access the correct object feature more often, reducing swap errors. This change in performance suggests that the additional information gained with multiple cues allowed participants to better discriminate between activated item representations, as opposed to activating memory representations that had been otherwise not accessible. If the latter were the case, multiple cues should have led to a reduction in guess responses. To determine whether the same findings hold for other object features, we ran two additional experiments, testing recall of color and locations, respectively.

Experiment 2

In Experiment 2, we altered the mapping between which features (location, color, and orientation) were used as cues and which feature was recalled. The results of Experiment 1 revealed that single-feature cues led to poorer performance than cues including both features, characterized primarily by an increase in swap errors at the expense of accessing the cued item. In this experiment,

¹ We thank an anonymous reviewer for suggesting this terminology.

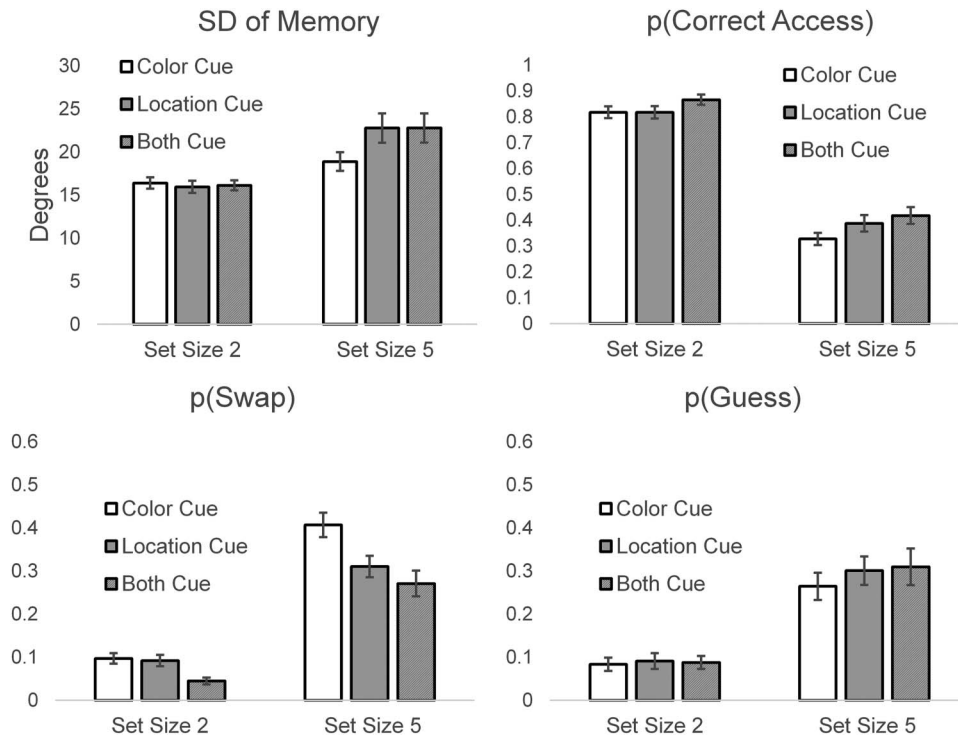


Figure 3. Summaries of memory performance in Experiment 1, recalling orientation. Error bars depict one standard error.

orientation and location were used as cues, and color was the recalled stimulus feature. We expected that providing both features in a cue would again maximize the probability of correctly reporting a target object's color, and reduce the likelihood of swaps.

Methods

Participants. As in Experiment 1, a new sample of 30 participants in total were recruited for this experiment. All participants were students in a first-year Psychology class at the University of Toronto, participating for course credit. All participants provided informed consent before participating. Fifteen participants completed a version of this experiment with a 100-ms exposure duration, and 15 participants completed a 600-ms exposure duration version.

Materials and procedure. With the exception of the test phase of trials, materials and procedure for this experiment were identical to Experiment 1. The test phase of a trial consisted of three types: orientation cues, location cues, or both-feature cues. Regardless of the cue type, the participant's task was to recall the color of the cued object from earlier in the trial using the mouse and a peripherally presented color wheel. All cue displays contained a central fixation cross, and a color wheel, centered on fixation with a radius of 8.25° and a line thickness of 0.35° . This color wheel depicted all of the possible stimulus hues, described in the Experiment 1 methods section. For orientation cues, a central, white triangle appeared on-screen whose orientation and size matched one of the triangles presented earlier in the trial. For location cues, a single, white, line-drawn circle, with a 1° diameter

and a line width of 0.25° appeared 6° from location, centered on the position of one of the triangles that had appeared in the memory display earlier in the trial. Lastly, for both-feature cues, an oriented white triangle appeared 6° from fixation, whose position and orientation matched one of the triangles from earlier in the trial. In all cases, when participants moved the cursor farther than 5° from fixation, the cue shape was filled in with the hue on the color wheel whose angular position relative to the center of the screen matched that of the mouse. After recalling the desired color, the participant submitted his or her response with a mouse click, and received feedback for 1s in the form of a small, black circle of radius 0.16° appearing on the color wheel over the exact color of the cued triangle.

Results

Raw memory error. Overall memory error can be seen in Figure 4. Initial analyses were again conducted on the raw error from memory reports in each Cue Condition (Orientation Cue, Location Cue, both-feature Cue) and Set Size (2 items, 5 items) for participants in both Sample Duration conditions (100 ms, 600 ms). Increasing Set Size increased memory error, as expected, $F(1, 28) = 835.29$, $p < .001$, $\eta_p^2 = 0.97$. In addition, Cue Condition affected memory error, $F(2, 56) = 24.69$, $p < .001$, $\eta_p^2 = 0.47$, such that memory error was lower when both-feature cues were used compared to orientation cues, $F(1, 28) = 41.14$, $p < .001$, $\eta_p^2 = 0.60$, and location cues, $F(1, 28) = 5.93$, $p = .02$, $\eta_p^2 = 0.18$. Although no two-way interactions were observed, $F_s < 0.99$, $p_s > .37$, $\eta_p^2 < 0.03$, a three-way interaction existed between Set Size,



Figure 4. Raw memory error in Experiment 2. Error bars depict one standard error.

Cue Condition, and Sample Duration, $F(2, 56) = 5.07, p = .009, \eta_p^2 = 0.15$. Analyzing performance separately by Set Size and Sample Duration showed that the benefit of both-feature cues over Location cues was limited to Set Size 2 of the 600-ms exposure duration, $F(1, 14) = 11.80, p = .004$. In all other conditions, no benefit was present for both-feature cues over Location only cues, $F_s(1, 14) < 2.79, p_s > .12$. Nonetheless, it is important to emphasize that the overall effect of Cue Condition on memory

error mirrored the results of Experiment 1; memory error was overall reduced with multifeature cues, albeit improvements over location-alone cues were inconsistent.

Three-component model analysis. To uncover the sources of the memory-cue benefit, responses were again transformed into performance parameters using the three-component mixture model (Bays, Catalao, & Husain, 2009) depicted in Figure 5. The main effect of Cue Condition was found for $p(\text{Correct Access})$ and $p(\text{Swap})$, as expected from the memory error analyses, $F_s(2, 56) > 6.90, p_s > .002, \eta_p^2 > 0.19$. However, both-feature cues only increased $p(\text{Correct Access})$ relative to Orientation cues, $F(1, 28) = 36.64, p < .001, \eta_p^2 = 0.57$, and did not boost performance relative to Location cues, $F(1, 28) = 0.47, p = .50, \eta_p^2 = 0.02$. The converse was true of $p(\text{Swap})$; fewer swaps occurred for Both-cue than Orientation cue trials, $F(1, 28) = 12.04, p = .002, \eta_p^2 = 0.30$, but only a marginal difference in swaps occurred between Both-cue and Location cue trials, $F(1, 28) = 0.39, p = .54, \eta_p^2 = 0.01$. This finding parallels the findings in the analyses of raw memory error, showing better recall of color from location cues than from orientation cues, but little improved recall when adding orientation information to a cue containing location information already.

Discussion

When reporting the color of objects at test, manipulating the type of cue once again altered the accessibility of information in VWM. Overall, cues with more visual information about an item led to improved ability to recall that item's color. Correct access was more likely in lieu of swap errors. One additional important caveat is that both-feature cues did not improve the probability of

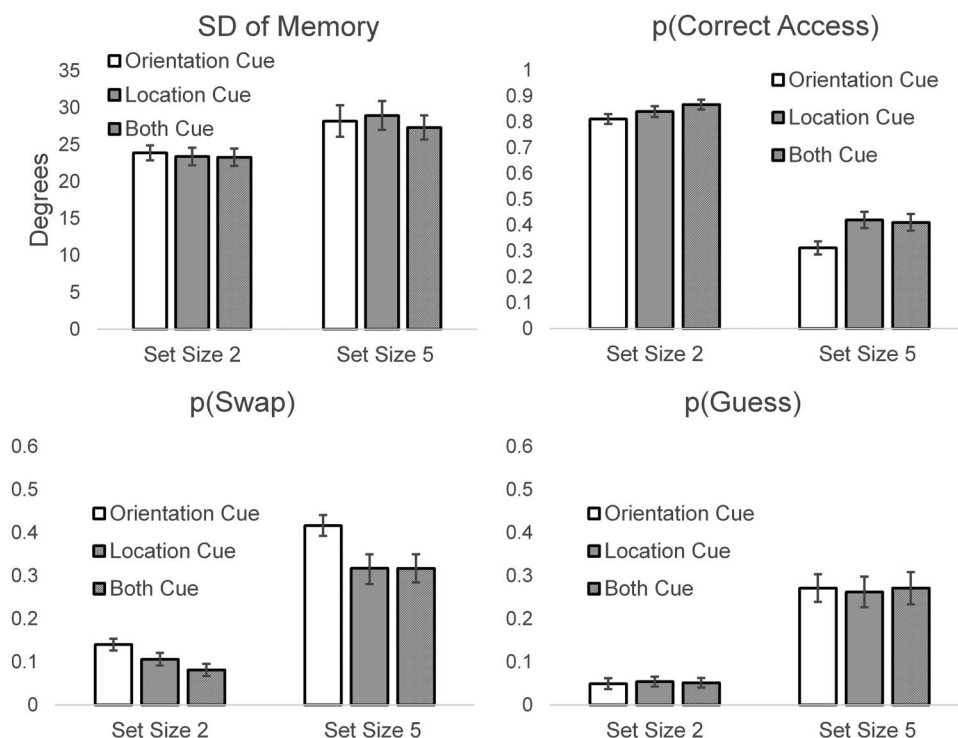


Figure 5. Summaries of memory performance parameters from Experiment 2, reporting color. Error bars depict one standard error.

recalling the correct item's color over a location cue alone. It seems that adding nonspatial features in a memory probe cannot always be counted on to improve upon retrieval over a location cue, unlike what we found with color. While we did not expect this discrepancy, orientation and color are fundamentally different features; orientation is an extrinsic feature of objects (assuming that different two-dimensional orientations do not produce a different perceived three-dimensional object shape, which we highly doubt with our stimuli) and color is an intrinsic feature, reflecting surface properties (leaving aside issues of color constancy). Empirically, it is known that search for a predefined color target in an array of heterogeneous colored dots is efficient (Wolfe et al., 1990), whereas search for a predefined orientation in an array of heterogeneous oriented lines is quite inefficient when orientation targets are not categorical (Wolfe, Friedman-Hill, Stewart, & O'Connell, 1992). Thus, there is the possibility that orientation may be less capable of guiding search through VWM. In our final experiment, we assessed the utility of the nonspatial features (color and orientation) in retrieving the locations of objects.

Experiment 3

The results of Experiments 1 and 2 showed that the type of information provided to access VWM does affect the probability that an object's features will ultimately be recalled. In Experiment 3, we compared the efficacy of color and orientation cues in recalling an object's location. Once again, we were most interested in the comparisons between single-feature and both-feature cues. In particular, Experiment 3 provided an opportunity to see whether the findings of Experiment 2, where orientation information paired with location information did not improve retrieval over location information alone, indicates that orientation information is not used in retrieval when another feature can be used instead.

Methods

Participants. Thirty participants were recruited for Experiment 3, all of whom were students enrolled in a first-year Psychology course, participating for course credit. Participants provided informed consent before participation. Fifteen participants completed a version of the experiment where stimuli were presented for 100 ms, and 15 participants completed a version in which stimuli were presented for 600 ms. None of the participants had participated in either of the preceding experiments.

Materials and procedure. As in Experiment 1, we ran separate sets of participants through a 100-ms sample duration condition and a 600-ms sample duration condition. Once again, with the exception of the test phase of trials, materials and procedure for this experiment were the same as for Experiments 1 and 2.

Three types of cues were provided in the test phase of trials: color cues, orientation cues, or both-feature cues. For all cue types, the participant's task was to move a centrally placed object to its original location in the periphery using the computer mouse. All cue displays contained a central fixation cross, and a white circle whose physical dimensions matched the color wheel from Experiment 2: centered on fixation with a radius of 8.25° and a line thickness of 0.35° . For orientation cues, a central, white triangle appeared in the center screen whose orientation and size matched one of the triangles presented earlier in the trial. For color cues, a

single line-drawn circle, with a 1° diameter and a line width of 0.25° whose color exactly matched one of the stimuli from earlier in the trial, appeared in the center of the screen. Lastly, for both-feature cues, an oriented, colored triangle appeared centrally whose color and orientation matched one of the triangles from earlier in the trials. In all cases, when participants moved the cursor farther than 5° from fixation, the cue shape moved to the periphery to the angular position corresponding to the mouse's deviation from fixation. The object was always constrained to have a radial distance of 6.6° from fixation (the same distance from fixation that triangles appeared at the beginning of the trial). Therefore, position errors could only be angular errors, analogous to the report orientation and report color experiments reported earlier. After placing the object in the desired position, the participant submitted his or her response with a mouse click, and received feedback for 1 s in the form of a small, black circle of radius 0.16° appearing on the white response wheel over the exact angular position of the cued triangle.

Results

Raw memory error. Raw memory error in each condition is depicted in Figure 6. Once again, initial analyses were performed on this raw error of memory reports. Set Size affected memory error, as expected, $F(1, 28) = 610.65, p < .001, \eta_p^2 = 0.96$, as did Cue Condition, $F(2, 54) = 82.89, p < .001, \eta_p^2 = 0.75$. Memory error was reduced when both-feature cues were provided in a cue compared to orientation cues, $F(1, 28) = 154.14, p < .001, \eta_p^2 = 0.85$, and color cues, $F(1, 28) = 57.77, p < .001, \eta_p^2 = 0.68$. Set size and cue condition also interacted, $F(2, 56) = 19.81, p < .001, \eta_p^2 = 0.42$, which we examined in the context of the memory parameters, below.

Three-component model analysis. To determine the source of the memory error gain, responses were once again transformed into performance parameters using the three-component mixture model (Bays, Catalao, & Husain, 2009), depicted in Figure 7. An analysis of these estimates demonstrated expected effects of set size on all parameters, $F_s(1, 28) > 84.09, p_s < .001, \eta_p^2 > 0.75$,

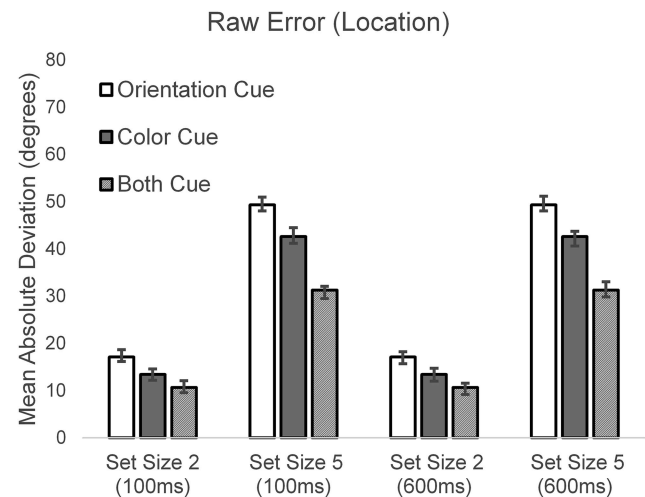


Figure 6. Raw memory error in Experiment 3. Error bars depict one standard error.

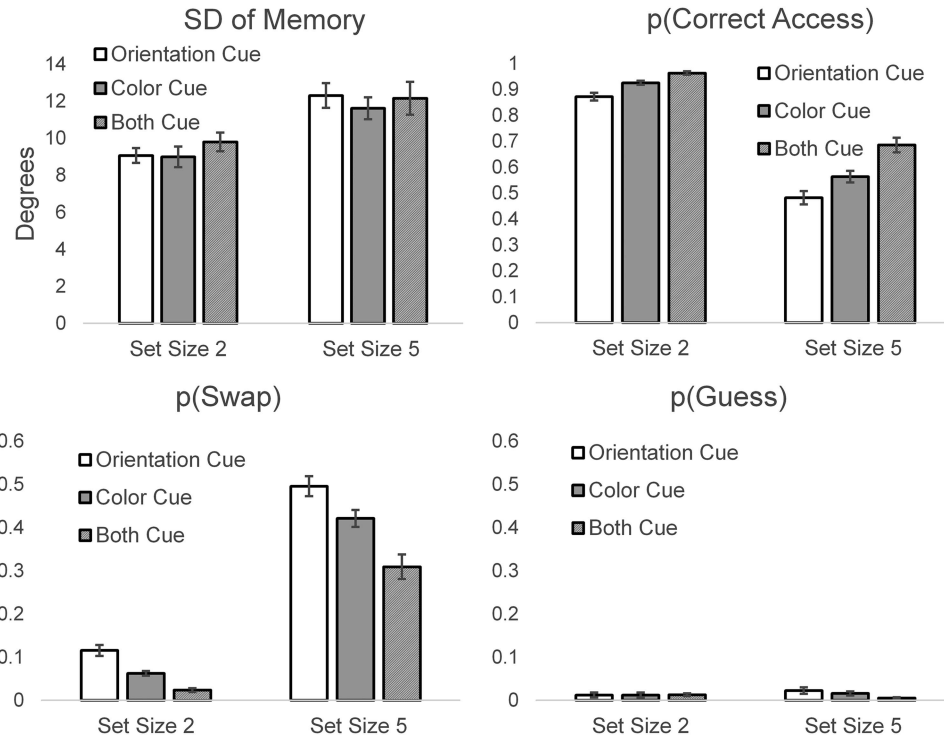


Figure 7. Summaries of memory performance in Experiment 3, reporting location. Error bars depict one standard error.

except for p (Guess). This lack of an effect for p (Guess) was due to the fact that, overall, random guess errors were very rare in our location recall task. In no condition did the average p (Guess) for participants exceed 0.03.

As in Experiments 1 and 2, Cue Condition affected p (Correct Access) and p (Swap), such that both-feature cues led to higher p (Correct Access) than orientation cues, $F(1, 28) = 193.00$, $p < .001$, $\eta_p^2 = 0.87$, and color cues, $F(1, 28) = 54.04$, $p < .001$, $\eta_p^2 > 0.66$, alone. Both-feature cues also led to lower p (Swap) than orientation cues, $F(1, 28) = 214.61$, $p < .001$, $\eta_p^2 > 0.89$, and color cues, $F(1, 28) = 38.77$, $p < .001$, $\eta_p^2 > 0.58$. Finally, cue condition also interacted with set size in determining p (Correct Access) and p (Swap), $F_s(2, 56) > 9.80$, $p_s < .004$, $\eta_p^2 > 0.26$. Importantly, both set sizes exhibited the same effects of cue condition on p (Correct Access), $F_s(2, 56) > 43.84$, $p < .001$, $\eta_p^2 > 0.61$, and p (Swap), $F_s(2, 56) > 41.58$, $p < .001$, $\eta_p^2 > 0.60$, and so this interaction reflects an amplification of the memory cue effect as set size increased. These data very clearly show that memory cues that provide more visual information can improve the likelihood of recalling an item's location.

Discussion

As in Experiments 1 and 2, the likelihood of correctly recalling an item's feature (in this case, location) was improved by cues with more features from the probed item. These correct responses primarily traded off with swap errors. In the context of the present experiment, this trade-off is not surprising given that participants did not opt to randomly guess in any condition. These results also

show that VWM retrieval can benefit from redundant retrieval information: Here, we consistently found benefits for both-feature cues over and above those for the best single feature cue.

Binding Pool Simulations

The results of three experiments showed that a manipulation of retrieval conditions (Cue Type) affected the probability of recalling a feature of an object. This result shows that the p (Correct Access) parameter, often referred to as "probability of memory" cannot be taken as a pure measure of the presence or absence of the representation of an object in VWM (see Bays, Catalao, & Husain, 2009). Given that the vast majority of VWM models are concerned with the quantity of information that is encoded or maintained, and not the processes by which items are recognized or recalled (van den Berg, Shin, Cou, George, & Ma, 2012; Wei, Wang, & Wang, 2012; Zhang & Luck, 2008, but see Johnson, Spencer, Luck, & Schöner, 2009 for a model that outlines a mechanism for same/different judgments and Pearson, Raškevičius, Bays, Pertzov, & Husain, 2014 for a mathematical model relating set size and precision to decision times), few models of VWM can account for our finding that the manipulation of retrieval factors influences performance. One recent exception is the recently developed binding pool model (Swan & Wyble, 2014), which specifies mechanisms used to extract a response given the information in a probe display for both change detection tasks and cued-recall tasks. Given that the binding pool provides a candidate mechanism for accessibility limits, we chose to include an analysis of simulated performance using the binding pool to determine whether it can

exhibit patterns of memory error caused by the retrieval manipulations used in our experiments.

Before describing our simulations, a brief summary of the binding pool is warranted. The binding pool model formalizes memory retrieval as a two-stage process: First, a retrieval cue activates an object-like representation, which then allows the desired features of the object to be retrieved. Noise at both stages may cause failure to retrieve information. The binding pool consists of three kinds of layers: type layers, which code particular features of remembered stimuli (e.g., their location, color, orientation); token node layers, which index particular objects akin to object files; and the binding pool layer, which acts as a hidden layer, associating the features comprising an object with their respective object codes in the token layer (see Figure 8).

Objects are encoded through a serial conjunction operation. For a given object, a node in the token layer is activated, along with the type layer neurons that code for its feature values. Each neuron in the token and type layers are randomly and pseudorandomly connected, respectively, to a subset of neurons in the binding pool. The representation of the object is the set of neurons in the binding pool that are jointly connected to the active token node and type layer neurons. This information is summed across object presentations, leaving a single, distributed code of activity in the binding pool that acts as the stored memory trace for all objects encoded.

For memory retrieval, type layers are used to “reactivate” a token, via the binding pool. If, for example, a dot is used to probe the memory of a stimulus in a particular location, the feature neuron of the location layer would be activated. This would, in turn, activate the neurons in the binding pool that are connected to the active location neuron. The binding pool activity that had been sustained from the encoding phase would be reduced to a subset of neurons that are jointly active for both the original memory code and the activated feature. The resulting pattern of activity in the binding pool then activates nodes in the token layer, with each token layer node’s activity being a function of the activation of binding pool neurons that connect to it. As a result, each token node would have some amount of activity. A particular object is considered “recognized” or “recalled” if its activation exceeds other token nodes’ activation by a particular threshold. Once this

winner-take-all process occurs, the single activated token node then prunes the binding pool activation again, leaving active only the neurons jointly activated by the winning token, and the binding pool activation established earlier in retrieval. Lastly, this resulting binding pool activation is used to activate each type layer to retrieve information about the recalled object’s appearance. Because this activation is noisy, a vector average of each type layer is used to establish each remembered feature value.

Given the large parameter space of the model, we opted to simulate performance in the present experiment over a wide sampling of the parameter space. This allowed us to see whether our main findings—an increase in $p(\text{Correct Access})$ and decrease in $p(\text{Swap})$ —would appear in simulations using different parameters. In other words, we sought to determine whether these results would emerge because of the algorithmic structure of the model, and not simply because of a particular parameter setting. To accomplish this, we produced a set of simulations using a coarse grid-search of the model’s parameter space. In each simulation, the model’s memory performance was simulated in an experiment using two set sizes, and three cue conditions, just as with our previous experiments. The model’s results were then fitted using the three-component model (Bays, Catalao, & Husain, 2009) and averaged, as in our preceding analyses.

In the grid-search, we simulated experimental results under all combinations of the following values of four model parameters for each feature: the degree of connectivity between a feature and the binding pool (type layer connectivity: 0.2, 0.275, 0.35, 0.425, 0.5), the proportion of shared connections between adjacent nodes in a type layer to the binding pool (similarity gradient: 0.05, 0.125, 0.2, 0.275, 0.35), the proportion of nodes in the binding pool connected to each node in the token layer (token connectivity: 0.2, 0.275, 0.35, 0.425, 0.5), and the threshold of activation required to retrieve a bound object representation given a memory probe (token individuation: 0.005, 0.0125, 0.02, 0.0275, 0.035). This resulted in the simulation of 625 simulated experiments.

To interpret these simulations, we opted to compare the change in memory performance when using two retrieval cues over one for the two set sizes. Because there were always two types of single-feature trials, we used the average difference between single- and both-feature performance, calculated as $\frac{\sum_{i=1}^2 M_i - M_{1,2}}{2}$, where M refers to the memory parameter in question, and the subscripts refer to the features used in memory retrieval, to quantify the both-feature advantage. These values were compared to the difference between memory performance for the two single-cue trials, M_1 and M_2 , which was simply calculated as $M_2 - M_1$. The distribution of changes in memory performance between the two single-cue trial types provides a convenient null distribution, as we did not implement any systematic differences between features. The distribution of changes in memory performance for double cues can then be compared against this null distribution to determine the extent to which different implementations of the model can be expected to show the retrieval effects that we found in our experiments. Figure 9 plots these values for each memory parameter as histograms.

As can be seen in Figure 9, only $p(\text{Correct Access})$ and $p(\text{Swap})$ are clearly, reliably affected by increasing the number of memory cues used in retrieval, despite changes in model parameter settings. At Set Size 5, a decrease in memory SD tended to appear with




Figure 8. A schematic illustration of the binding pool model’s architecture. See the online article for the color version of this figure.

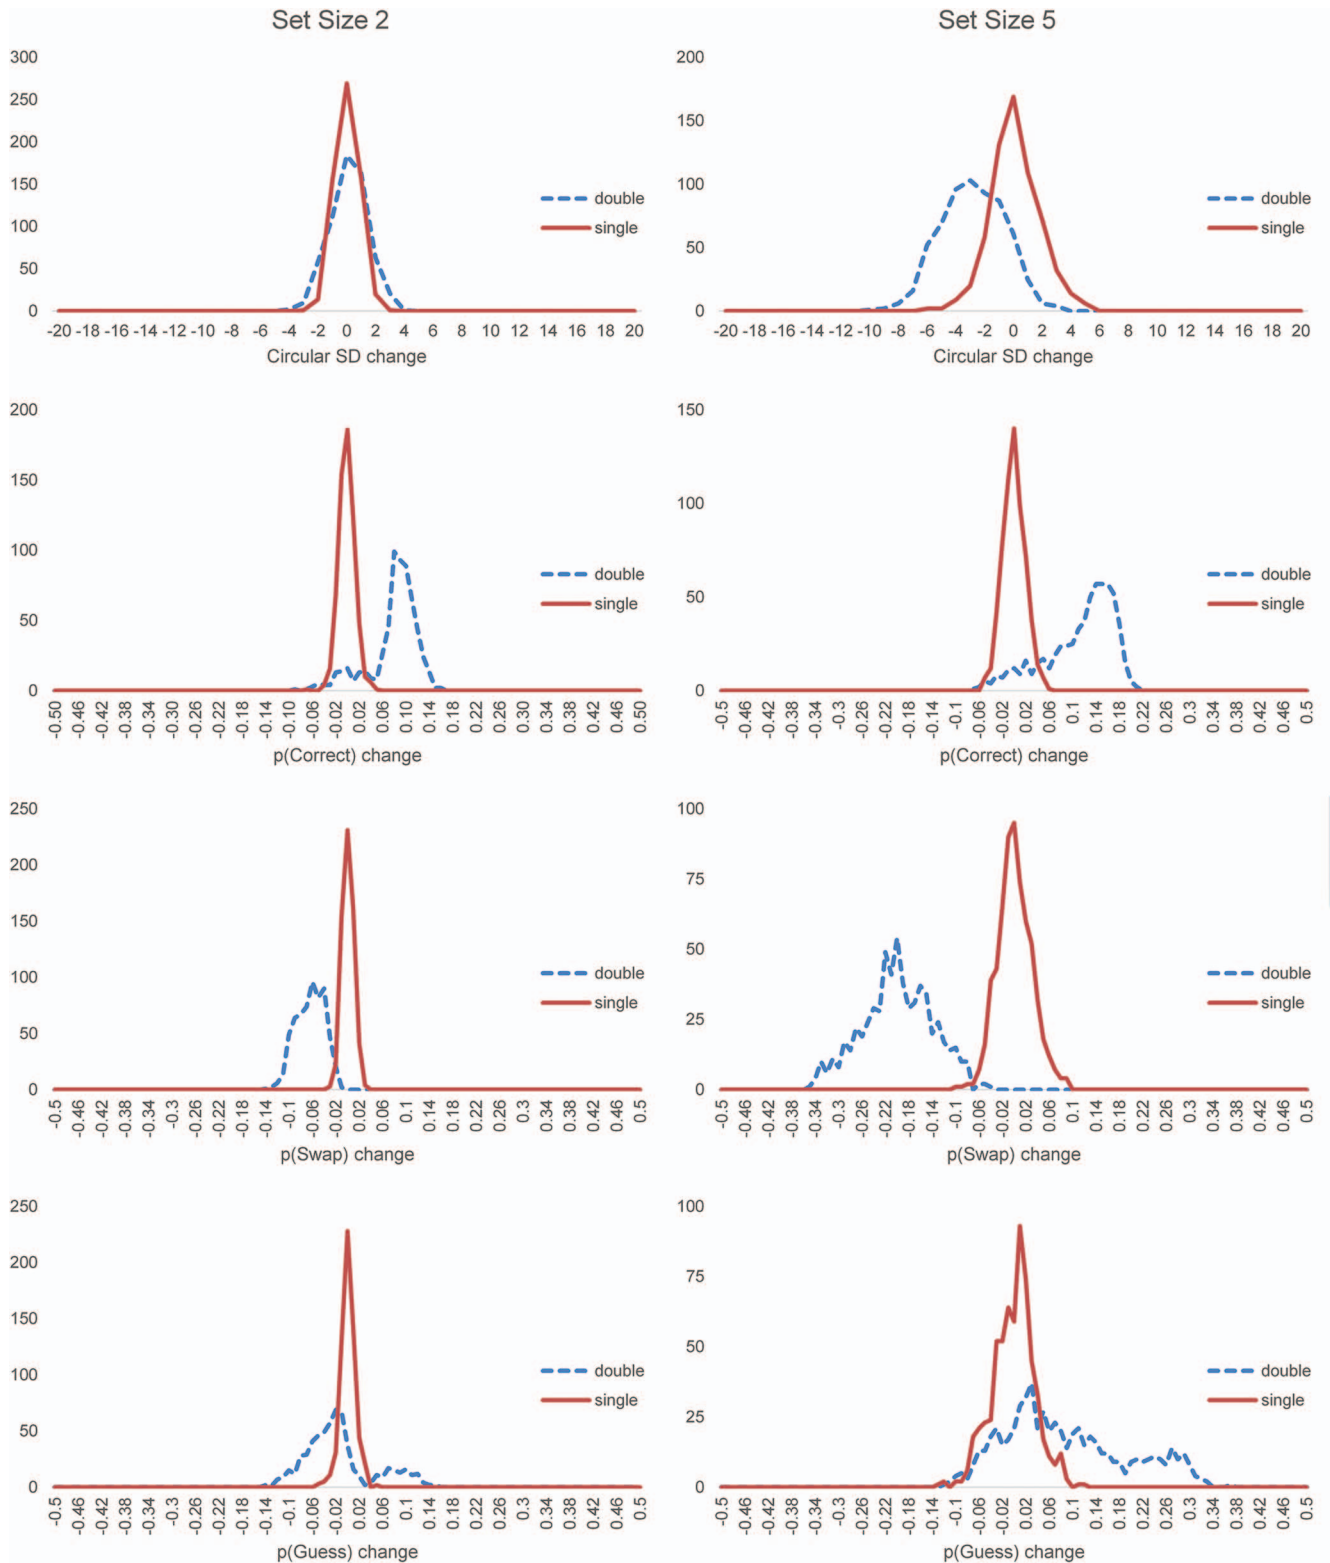


Figure 9. Histograms of the effect of different retrieval cues on memory performance in binding pool simulations. "Single" corresponds to the average difference between the trials where a single-feature cue was used in retrieval, and "double" corresponds to the average difference between both-feature cues were used, compared to a single-feature cue. See the online article for the color version of this figure.

more memory cues, but only 42.7% of simulations showed an increase outside of a 95% confidence interval constructed from the single-cue simulations. For comparison, 85% of simulations showed an increase in $p(\text{Correct Access})$ outside of the 95% confidence interval for single-cue simulations (for both Set Sizes 2 and 5), and 98% (Set Size 2) and 99% (Set Size 5) of simulations showed a reduction in swaps with two-feature cues that was beyond the 95% confidence interval surrounding the single-cue simulations. Guesses, like memory SD , were affected by the use of two features in a memory cue, but only increased beyond the 95% confidence interval on single-cue simulations 19% and 46% of the time for each set size, respectively. Overall, our simulations show the two consistent findings of our experimental results, an increase in $p(\text{Correct Access})$ and decrease in $p(\text{Swap})$ with both-feature cues, occur for the vast majority of parameter settings of binding pool, but that changes in memory precision and guessing depend on how the parameters are set.

To understand how the binding pool leads to these changes in memory performance, we inspected the distribution of average binding pool neuron activations during retrieval. Figure 10 shows the average difference in the number of binding pool neurons activated during retrieval between memory cue conditions at two stages of retrieval. In the first stage, the number of binding pool neurons is determined by the pattern of activity established after encoding and the neurons that are activated by the retrieval cue. In the second stage, after a token has been selected, the selected token further narrows down binding pool activity in order to isolate

information about the retrieved object. As can be seen, an additional feature at retrieval reduces the number of binding pool neurons activated in Stage I, as well as Stage II to a lesser extent. The reduction in Stage I in the number of active binding pool neurons is critical for token node retrieval, as the binding pool activity codes for all items simultaneously. When two cue features are available to constrain the binding pool activity, this reduces the overall number of active binding pool neurons, but importantly leaves a larger proportion that are unique to the binding of the target item's features. This allows the correct object representation, or token, to be uniquely activated in retrieval. That the difference in active binding pool neurons is reduced between both-feature and single-feature conditions in Stage II reflects the contribution of the retrieved token node; regardless of how many cues are presented, once a token node is retrieved, that will provide a further, constant reduction in the binding pool activity in order to solely represent the probed object.

Unlike our empirical data, these simulations occasionally show increases in guessing when more features are provided for memory retrieval. One reason for this may lie in the decision mechanism of token retrieval. The current decision rule is that, once tokens are activated in Stage I, if one token node is activated sufficiently above others (by a threshold amount) it will win the retrieval competition and activate its object's stored features. If token nodes are not sufficiently different in activation, a random response will occur. This suggests that, when uncertainty exists between two or more objects, the model will guess. One issue with this when

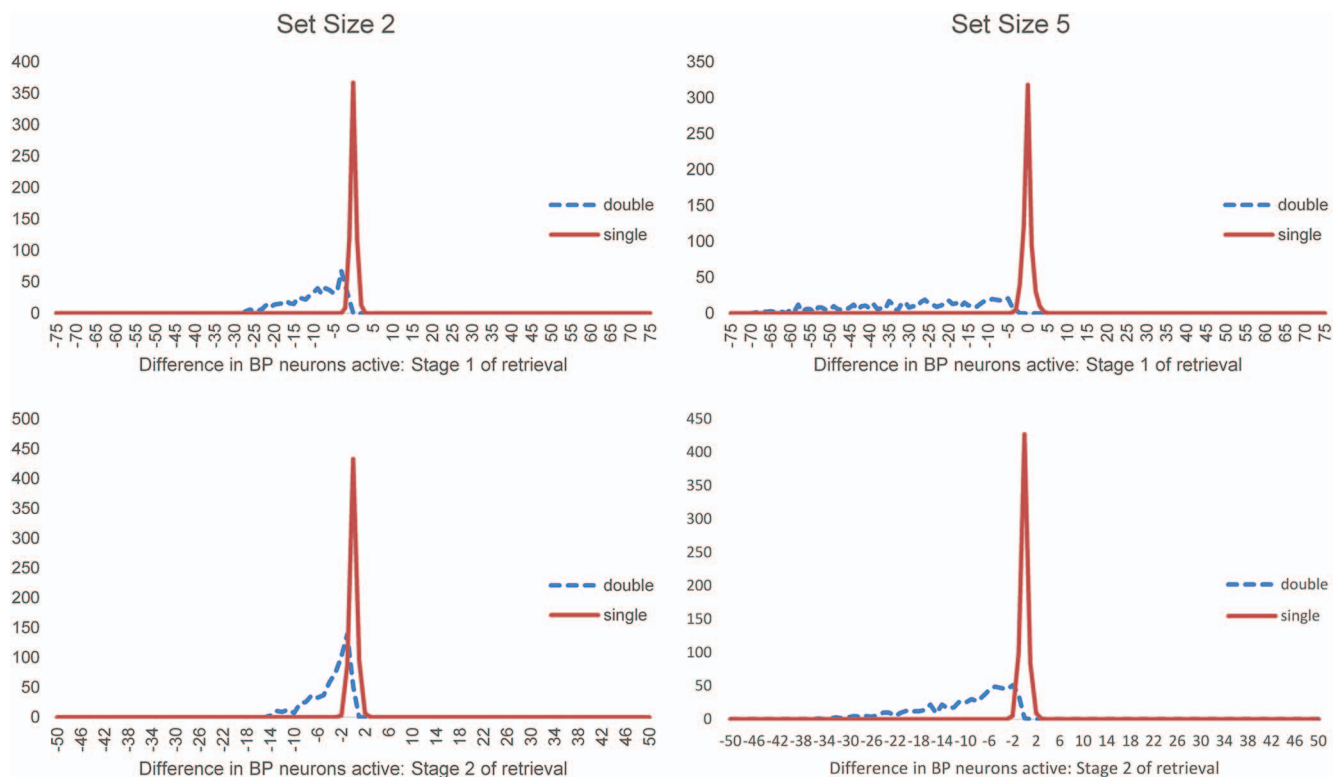


Figure 10. Histograms of BP neuron activation differences when using a single- or two-feature cue for Set Size 2 (left column) and Set Size 5 (right column) and for Stage I (upper row) and Stage II (lower row) of retrieval. See the online article for the color version of this figure.

considering variability in retrieval cues is that, as seen above, more cues leads to fewer active binding pool neurons. Because token activation is determined by summing the activity of the binding pool neurons connected to each token node, this means that the total activity of each token node will be reduced, making it more likely that no token node will be higher than another token node by the threshold amount. If token node selection were based upon the ratio of activity, instead, this could eliminate the increase in guessing that we observed in some simulations.

To summarize, our simulations using the binding pool show that the improvement in correct memory retrieval, and the reduction in incorrect item retrieval with additional retrieval cues, is a robust prediction of the binding pool's architecture. The critical factor in correct retrieval of an item is the reduction of initial binding pool activity, which represents all stored items simultaneously, to the subset of neurons that represent the probed item. The number of features that are used to retrieve an item, then, help to individuate one particular object in memory.

General Discussion

An often overlooked issue in the VWM literature is the nature of access to stored visual information. In three experiments, we assessed the variation in cued-recall performance caused by different types of cues at the test stage of a delayed-estimation task. As we expected, providing memory cues with more features maximized participants' ability to recall a tested object's orientation, color, or location. However, it is not the case that the single-feature cues were consistently inferior to double-feature cues. When reporting color, providing the location information alone was in some cases enough to maximize participants' ability to access the probed item's features, such that adding a nonspatial feature did not provide further improvement in memory performance. That location-based cues were occasionally superior to nonspatial cues is consistent with previous demonstrations of a precedence of spatial information in VWM (Jiang, Olson, & Chun, 2000; Olson & Marshuetz, 2005, but see Logie, Brockmole, & Jaswal, 2011). Indeed, in our experiments, the overall probability of correctly reporting a cued location was greater than reporting a cued color or cued orientation for the same class of objects (see also, Rajsic & Wilson, 2014).

It is possible that the reason that locations were occasionally a superior feature for retrieving nonspatial information may be due to the relatively higher precision with which location is remembered (or can be perceived), and that location does not have any special role in memory representation or retrieval. While the superior precision of location coding may explain its utility in retrieval, we should note that the circular *SD* for correct reports in our data was, on average, better for orientation ($M_{100\text{ms}} = 19.01^\circ$, $SE_{100\text{ms}} = 1.33^\circ$, $M_{600\text{ms}} = 18.62^\circ$, $SE_{600\text{ms}} = 1.27^\circ$) than for color ($M_{100\text{ms}} = 27.11^\circ$, $SE_{100\text{ms}} = 1.61^\circ$, $M_{600\text{ms}} = 24.56^\circ$, $SE_{600\text{ms}} = 1.27^\circ$), but color proved to be the superior feature in retrieving object locations for both set sizes and sample durations compared to orientation, $t_s(14) > 2.69$, $p_s < .02$. It is therefore tempting to speculate that the efficacy of retrieving information from VWM with different features may be related to other known feature-differences in perception, for example, the ability to guide visual attention

using different features (Wolfe & Horowitz, 2004). In fact, visual search provides a nice parallel for our finding that an increased number of features aids in the retrieval of, or search for, a visual memory: Triple conjunction search tasks (where more features are available to disambiguate targets from distractors) show better search efficiency than standard, two-feature conjunction tasks (Wolfe, Cave, & Franzel, 1989). However, the specific task and stimulus conditions likely mediate the relative ability of different features to retrieve information from VWM (see Heuer & Schubö, 2016).

Two salient possibilities for how multifeature memory cues could affect recall from VWM appeared plausible. First, multifeature cues may have been more effective because they resolve conflict regarding correspondence; and, second, multifeature cues may have been more effective because they overcome the problem of partially complete representations. The first suggests that matching visual information across time is a noisy process. Several researchers have argued that memory contains inherent uncertainty (Fougnie, Suchow, & Alvarez, 2012; Ma, Husain, & Bays, 2014) that is measurable when object features are recalled from VWM. However, this uncertainty should also contribute to error in the process of accessing memory. Adding features to a cue may aid in constraining the matching process—activating fewer object representations that match the cue, and preventing swap errors, as we demonstrated with the binding pool model. Although we were not able to show an improvement in memory precision when multiple-feature cues were presented—a situation that should reduce correspondence ambiguity—our results are compatible with the overall conclusion that correspondence is an additional source of memory failures in VWM, alongside limited capacity for information, as we often did observe a reduction in swap errors with more informative memory cues.

In addition to alleviating correspondence problems, single-feature cues could have failed to retrieve information for those representations in VWM that are only partially complete. Fougnie and Alvarez (2011, see also: Bays, Wu, & Husain, 2011) have shown that loss of information in VWM can occur at the feature level, such that a representation may contain, for example, a location and color, but not orientation. Such representations would prove problematic if the cue provided only orientation information. In such a case, it would not be possible for the cue to activate the appropriate object representation for report, even though reportable information would be present. If this is indeed occurring, our data suggest that participants opt to report some known feature in these cases. Given that it is unclear whether swap errors in location-recall tasks reflect lost information about the cued object or a correspondence problem (see Rajsic & Wilson, 2014), this issue is deserving of further investigation. Indeed, if swap errors are simply strategic responses to situations where the cue does not retrieve item-specific details, then our data would be entirely compatible with a partial-representation account of VWM, where some objects have missing information about their nonspatial features. Until a thorough account of response strategy in the delayed-estimation task is available, whether swap errors reflect ignorance of a cued object's features or simply confusion about which known objects' features should be reported will remain unknown. We note that Rajsic and Wilson (2014) completely

eliminated swap responses by presenting all nontested items on the test display of each trial, suggesting that swap responses reflect uncertainty about the specific object being cued, albeit when the cued object's feature is unable to be reported. Thus, random guesses may only occur when participants are confident that they do *not* know the feature of the cued object. As such, partial representation is consistent with our results, as cuing a missing feature (e.g., using “blue” to cue a blue triangle) may still sufficiently activate a similar item (a green triangle) above others (a red and an orange triangle), leading to a swap response.

Throughout our results, we consistently observed that our retrieval manipulations affected the retrieval of discrete features. Providing more information in a memory cue did not reliably increase the precision of retrieved information. Similarly, retro-cues, which provide participants information about which item will be tested after memory encoding has already occurred, appear to affect only the likelihood of retrieval, and not precision (Hollingworth & Hwang, 2013; Murray et al., 2013, but see Gunseli, van Moereselaar, Meeter, & Olivers, 2015). Taken together, these results suggest that the representational precision of memory items is established at encoding. As mentioned previously, Bae and Flombaum (2013) have shown that correspondence failures can affect representational precision. However, their manipulation was perceptual in nature; when features were reported with higher precision, they also appeared within a physically different stimulus. Higher memory precision was observed when simultaneously presented stimuli did not share an irrelevant feature (color, shape, or frequency) compared to when they did share an irrelevant feature, and therefore the difference in precision may have emerged during memory encoding in their study.

While our study was able to show that failures of memory can emerge due to accessibility limits, it is unclear how much these failures may account for performance limits in the many studies that have used the delayed estimation paradigm (Luck & Vogel, 2013; Ma, Husain, & Bays, 2014). One unique feature of our paradigm (but see Emrich & Ferber, 2012) was that our stimuli were not always highly discriminable on the dimension used to cue memory. It is possible, then, that poorer performance on single-cue trials could be simply due to guess and swap responses stemming from trials where the cued object and a noncued object were close on the cue-feature dimension. However, when we reanalyzed mean absolute memory error after excluding all trials where a noncued object appeared within 20° (clockwise or counterclockwise) of the cued object on either cue dimension (e.g., color and location for Experiment 1), we still observed a main effect of Cue Condition in all experiments (Experiment 1: $F(2, 56) = 3.15, p = .05, \eta_p^2 = .10$; Experiment 2: $F(2, 56) = 24.62, p < .001, \eta_p^2 = .47$; Experiment 3: $F(2, 56) = 54.23, p < .001, \eta_p^2 = .66$). Because of the large reduction in trial counts associated with removal of these “near miss” trials, we could not confidently analyze performance on these trials using the mixture-modeling approach. As a way of confirming that a similar trade-off between correct reports and swaps occurred here, however, we combined trials across all observers for a given experiment and condition, and fit a single mixture model to these data. The resulting fits are shown in Figure 11. As can be seen, they mirror the data from Experiments 1–3 qualitatively; $p(\text{Correct Access})$ is greater for Both-Cues than

single cues, and $p(\text{Swap})$ is lower for Both-cues than single cues. Thus, cue ambiguity alone cannot account for our findings. We do note, however, that most existing studies have endeavored to minimize accessibility issues, such as by using highly discriminable locations and marking the locations of nontested items (e.g., Zhang & Luck, 2008). As such, we do not intend to claim that accessibility differences in VWM account for well-established memory performance reductions associated with, for example, set size. Our goal here is simply to provide insight into the mechanisms of cued-recall from VWM, which is an integral component of delayed estimation that remains poorly understood.

In our article, we have used the binding pool model of VWM to account for our data. The binding pool has an explicitly defined retrieval algorithm, making it ideal for understanding our findings. Indeed, the binding pool was able to provide a computational explanation of the results of our experiments—its two-stage retrieval process fits with the finding that manipulations at the recall stage of a delayed-estimation experiment affect the retrieval of discrete objects. In addition, our data provide a confirmation of the most robust prediction of the binding pool's retrieval process: that multiple cues improve retrieval of bound item representations. Our later simulations showed that the binding pool produces this behavior over a wide range of parametrizations; in fact, it was present in the vast majority of them. This lends support to the argument that the binding pool indeed captures important aspects of how information is retrieved from VWM. In future applications of the binding pool, this data will be able to place constraints on plausible parametrizations. For example, a sizable number of binding pool parametrizations showed an increase in guessing with multiple cues, whereas this was not observed in experimental data. We speculate that the critical difference between our data and simulations may lie in the process of deciding whether sufficient evidence exists for a correspondence between a remembered item and a probe. The binding pool's initial decision was a relative threshold rule: If one item's token activation exceeds other items' activation by a particular amount, it “wins” the retrieval competition. However, other rules, such as a ratio-based threshold, could be the key to these differences.

One aspect of the data that we did not capture in our simulations was the “special” status of location in retrieval that occasionally emerged in our data. At this stage of its implementation, the binding pool model treats all features as homogeneous, and so a natural way of accommodating this result would be to introduce inhomogeneities in feature coding, for example, richer representational resources (i.e., more type nodes) for the location layer. Another potential change that may reproduce a special status for location would be to encode object features in a location-based manner, sampling bindings between locations and nonspatial features independently for each object. For example, location-color and location-orientation bindings for each object could be probabilistically sampled. This is consistent with several accounts of encoding (Bundesen, Kyllingsbaek, & Larsen, 2003; Cowan, Blume, & Saults, 2013; Vul & Rich, 2010) that suggest bindings between locations and different nonspatial features are independently sampled. Importantly, this sampling algorithm could produce




Figure 11. VWM performance parameters, fit using all trials across participants that contained no objects within 20° of either feature used as a cue. Bars with white backgrounds depict data from the following single-cue conditions: Experiment 1: Color, Experiment 2: Orientation, Experiment 3: Orientation; gray backgrounds depict data from the following single-cue conditions: Experiment 1: Location, Experiment 2, Location, Experiment 3, Color; and black backgrounds depict data from Both-cue conditions.

the partial object representations that may underlie our measured retrieval effects.

As a final note, our results underscore the difficulty in inferring the properties of VWM directly from measured parameters; given that decisions about testing procedure alter performance in the delayed estimation task, empirically derived memory parameters cannot be considered a complete picture of memory representations without considering the process that produces responses. We have chosen to ground our interpretation of performance in the network structure of the binding pool (Swan & Wyble, 2014). A distinct advantage of the binding pool is that it specifies not only how information is encoded and stored in VWM, but how it is retrieved.

Conclusions

By manipulating the features provided in memory cues at test, we show that access to information in VWM is a source of performance limits. The likelihood of correctly reporting an object's orientation, color, or location was sensitive to the type and amount of information provided by a cue. We suggest that these memory cue effects may stem from two sources: reduction of correspondence errors between cues and representations in VWM, and overcoming problems of partial-information. Our results highlight the limitations inherent in the visual system for dealing with

information over the short term, and extend the issue of information accessibility to visual working memory.

References

- Bae, G. Y., & Flombaum, J. I. (2013). Two items remembered as precisely as one: How integral features can improve visual working memory. *Psychological Science*, 24, 2038–2047. <http://dx.doi.org/10.1177/0956797613484938>
- Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9, 7–11. <http://dx.doi.org/10.1167/9.10.7>
- Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, 49, 1622–1631. <http://dx.doi.org/10.1016/j.neuropsychologia.2010.12.023>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436. <http://dx.doi.org/10.1163/156856897X00357>
- Bundesen, C., Kyllingsbaek, S., & Larsen, A. (2003). Independent encoding of colors and shapes from two stimuli. *Psychonomic Bulletin & Review*, 10, 474–479. <http://dx.doi.org/10.3758/BF03196509>
- Cowan, N. (1993). Activation, attention, and short-term memory. *Memory & Cognition*, 21, 162–167. <http://dx.doi.org/10.3758/BF03202728>
- Cowan, N., Blume, C. L., & Saults, J. S. (2013). Attention to attributes and objects in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 731.
- Emrich, S. M., & Ferber, S. (2012). Competition increases binding errors in visual working memory. *Journal of Vision*, 12, 12–12.

- Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision*, 11, 3.
- Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, 3, 1229. <http://dx.doi.org/10.1038/ncomms2237>
- Gorgoraptis, N., Catalao, R. F., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *The Journal of Neuroscience*, 31, 8502–8511. <http://dx.doi.org/10.1523/JNEUROSCI.0208-11.2011>
- Gunseli, E., van Moorselaar, D., Meeter, M., & Olivers, C. N. L. (2015). The reliability of retro-cues determines the fate of noncued visual working memory representations. *Psychonomic Bulletin & Review*, 22, 1334–1341. <http://dx.doi.org/10.3758/s13423-014-0796-x>
- Heuer, A., & Schubö, A. (2016). Feature-based and spatial attentional selection in visual working memory. *Memory & Cognition*, 44, 621–632. <http://dx.doi.org/10.3758/s13421-015-0584-5>
- Hollingworth, A. (2003). Failures of retrieval and comparison constrain change detection in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 388–403. <http://dx.doi.org/10.1037/0096-1523.29.2.388>
- Hollingworth, A., & Hwang, S. (2013). The relationship between visual working memory and attention: Retention of precise colour information in the absence of effects on perceptual selection. *Philosophical Transactions of the Royal Society of London: Series B, Biological Sciences*, 368, 20130061. <http://dx.doi.org/10.1098/rstb.2013.0061>
- Hyun, J. S., Woodman, G. F., Vogel, E. K., Hollingworth, A., & Luck, S. J. (2009). The comparison of visual working memory representations with perceptual inputs. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 1140–1160. <http://dx.doi.org/10.1037/a0015019>
- Jiang, Y., Olson, I. R., & Chun, M. M. (2000). Organization of visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 683–702. <http://dx.doi.org/10.1037/0278-7393.26.3.683>
- Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science*, 20, 568–577. <http://dx.doi.org/10.1111/j.1467-9280.2009.02329.x>
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception 36 ECVF Abstract Supplement*.
- Lawrence, M. A. (2010). Estimating the probability and fidelity of memory. *Behavior Research Methods*, 42, 957–968. <http://dx.doi.org/10.3758/BRM.42.4.957>
- Levillain, F., & Flombaum, J. I. (2012). Correspondence problems cause repositioning costs in visual working memory. *Visual Cognition*, 20, 669–695. <http://dx.doi.org/10.1080/13506285.2012.683050>
- Logie, R. H., Brockmole, J. R., & Jaswal, S. (2011). Feature binding in visual short-term memory is unaffected by task-irrelevant changes of location, shape, and color. *Memory & Cognition*, 39, 24–36. <http://dx.doi.org/10.3758/s13421-010-0001-z>
- Luck, S. J. (2008). Visual short-term memory. In S. J. Luck & A. Hollingworth (Eds.), *Visual memory* (pp. 43–86). New York, NY: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195305487.003.0003>
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281. <http://dx.doi.org/10.1038/36846>
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17, 391–400. <http://dx.doi.org/10.1016/j.tics.2013.06.006>
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, 17, 347–356. <http://dx.doi.org/10.1038/nn.3655>
- McElree, B. (2001). Working memory and focal attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 817–835. <http://dx.doi.org/10.1037/0278-7393.27.3.817>
- Murray, A. M., Nobre, A. C., Clark, I. A., Cravo, A. M., & Stokes, M. G. (2013). Attention restores discrete items to visual short-term memory. *Psychological Science*, 24, 550–556. <http://dx.doi.org/10.1177/0956797612457782>
- Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 411–421. <http://dx.doi.org/10.1037/0278-7393.28.3.411>
- Olson, I. R., & Marshuetz, C. (2005). Remembering “what” brings along “where” in visual working memory. *Perception & Psychophysics*, 67, 185–194. <http://dx.doi.org/10.3758/BF03206483>
- O'Regan, J. K., & Noë, A. (2001). What it is like to see: A sensorimotor theory of perceptual experience. *Synthese*, 129, 79–103. <http://dx.doi.org/10.1023/A:1012699224677>
- Pearson, B., Raškevičius, J., Bays, P. M., Pertzov, Y., & Husain, M. (2014). Working memory retrieval as a decision process. *Journal of Vision*, 14, 2–15. <http://dx.doi.org/10.1167/14.2.2>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442. <http://dx.doi.org/10.1163/156856897X00366>
- Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, 139, 23–38. <http://dx.doi.org/10.1016/j.neuroscience.2005.06.005>
- Rajsic, J., & Wilson, D. E. (2014). Asymmetrical access to color and location in visual working memory. *Attention, Perception, & Psychophysics*, 76, 1902–1913. <http://dx.doi.org/10.3758/s13414-014-0723-2>
- Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, 13, 9. <http://dx.doi.org/10.1167/13.10.9>
- Swan, G., & Wyble, B. (2014). The binding pool: A model of shared neural resources for distinct items in visual working memory. *Attention, Perception, & Psychophysics*, 76, 2136–2157. <http://dx.doi.org/10.3758/s13414-014-0633-3>
- Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning & Verbal Behavior*, 5, 381–391. [http://dx.doi.org/10.1016/S0022-5371\(66\)80048-8](http://dx.doi.org/10.1016/S0022-5371(66)80048-8)
- van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, 121, 124–149. <http://dx.doi.org/10.1037/a0035234>
- van den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 8780–8785. <http://dx.doi.org/10.1073/pnas.1117465109>
- Vul, E., & Rich, A. N. (2010). Independent sampling of features enables conscious perception of bound objects. *Psychological Science*, 21, 1168–1175. <http://dx.doi.org/10.1177/0956797610377341>
- Wei, Z., Wang, X. J., & Wang, D. H. (2012). From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *The Journal of Neuroscience*, 32, 11228–11240. <http://dx.doi.org/10.1523/JNEUROSCI.0735-12.2012>
- Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131, 48–64. <http://dx.doi.org/10.1037/0096-3445.131.1.48>
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4, 11. <http://dx.doi.org/10.1167/4.12.11>

- 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. <http://dx.doi.org/10.1037/0096-1523.15.3.419>
- Wolfe, J. M., Friedman-Hill, S. R., Stewart, M. I., & O'Connell, K. M. (1992). The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 34–49. <http://dx.doi.org/10.1037/0096-1523.18.1.34>
- Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5, 495–501. <http://dx.doi.org/10.1038/nrn1411>
- Wolfe, J. M., Yu, K. P., Stewart, M. I., Shorter, A. D., Friedman-Hill, S. R., & Cave, K. R. (1990). Limitations on the parallel guidance of visual search: Color x color and orientation x orientation conjunctions. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 879–892. <http://dx.doi.org/10.1037/0096-1523.16.4.879>
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representation in visual working memory. *Nature*, 453, 233–235. <http://dx.doi.org/10.1038/nature06860>
- Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, 20, 423–428. <http://dx.doi.org/10.1111/j.1467-9280.2009.02322.x>

Appendix

Analysis of Statistical Interactions

Experiment 1

In addition to a main effect of Cue Condition, we also observed interactions between Cue Condition and Set Size for the $p(\text{Swap})$, $p(\text{Correct Access})$, and the circular Standard Deviation of correct responses (SD), indicating that the effect of memory cues differed by Set Size. Analyzing Set Sizes separately showed that, at Set Size 2, Cue Condition affected $p(\text{Correct Access})$ and $p(\text{Swap})$ alone, $F_s(2, 58) > 8.41$, $ps \leq .001$, $\eta_p^2 > 0.22$, such that both-cue trials increased $p(\text{Correct Access})$ relative to color cues, $F(1, 29) = 14.75$, $p = .001$, $\eta_p^2 = 0.34$, and location cues, $F(1, 29) = 11.03$, $p = .002$, $\eta_p^2 = 0.28$, and decreased $p(\text{Swap})$ correspondingly, $F_s(1, 29) > 14.91$, $ps \leq .001$, $\eta_p^2 > 0.34$. This fits the pattern noted earlier, with better access to visual memories when both features were used to cue an item than when either feature alone was provided.

At Set Size 5, memory cues affected correct SD , $F(2, 58) = 3.16$, $p = .05$, $\eta_p^2 = .10$, such that color-cued SD was lower (and, therefore, memory precision was higher) compared to both-feature cued SD , $F(1, 29) = 3.82$, $p = .06$, $\eta_p^2 = 0.12$, whereas no difference existed between the SD for location cues and both-feature cues, $F(1, 29) = 0$, $p = .995$, $\eta_p^2 = 0$. This accounts for the interaction between Cue Condition and Set Size for SD , as no effects on SD we observed at Set Size 2; at Set Size 5 only, orientation was more precisely recalled when retrieved using color than location, or location along with color. With regards to access, the differences between both-feature cues and single-feature cues in $p(\text{Correct Access})$ and $p(\text{Swap})$ only occurred when the single-feature cue was a color-cue, $F_s(1, 29) > 12.87$, $ps = .001$, $\eta_p^2 > 0.30$, and no difference existed between both-feature cues and location cues, $F_s(1, 29) < 1.65$, $ps > .20$, $\eta_p^2 < 0.06$. At larger set sizes, then, having two features in a recall cue only improved access over a color cue alone, suggesting that participants may have relied on location primarily at higher set sizes for retrieval. Nonetheless, at no point did either single-feature cue lead to more frequent access than the both-feature cue condition, indicating that more informative cues led to maximal access.

Experiment 2

A three-way interaction existed between Set Size, Cue Condition, and Sample Duration, $F(2, 56) = 5.07$, $p = .009$, $\eta_p^2 = 0.15$, and so our follow-up analyses using the three-component memory model were done separately for each Sample Duration.

Three-component model analysis: 100-ms sample duration.

To uncover the sources of the memory-cue benefit, responses were again transformed into performance parameters using the three-component mixture model (Bays, Catalao, & Husain, 2009) depicted in Figure 5. With a 100-ms memory display duration, we observed the expected main effects of Set Size for all memory parameters, $F_s(1, 14) > 5.84$, $ps < .04$, $\eta_p^2 > 0.28$. More importantly, for the present investigation, Cue Condition produced a reliable change in $p(\text{Swap})$, $F(2, 28) = 5.20$, $p = .01$, $\eta_p^2 = .27$, with no change in the probability of guessing [$p(\text{Guess})$], $F(2, 28) = 1.87$, $p = .17$, $\eta_p^2 = 0.12$, or memory SD , $F(2, 28) = 1.19$, $p = .32$, $\eta_p^2 = .08$. Instead, we observed a marginal effect on $p(\text{Correct Access})$, $F(2, 28) = 3.10$, $p = .06$, $\eta_p^2 = 0.18$, suggesting that the change in $p(\text{Swap})$ was driven by a complementary change in $p(\text{Correct Access})$, as in Experiment 1.

As we observed with the data from Set Size 5 in Experiment 1, cues with both spatial and nonspatial information were superior only to nonspatial only cues for the short sample duration performance in Experiment 2. Both-feature cues improved color recall compared to Orientation cues, such that $p(\text{Correct Access})$ was higher and $p(\text{Swap})$ was lower, $F_s(1, 14) > 11.21$, $p < .005$, $\eta_p^2 > 0.44$, but this was not true for both-feature cues when contrasted with Location Cues, $F_s(1, 14) < 2.98$, $ps > .10$, $\eta_p^2 < 0.18$. Finally, a marginal interaction was observed for $p(\text{Guess})$ only, $F(2, 28) = 3.00$, $p = .07$, $\eta_p^2 = 0.018$, but given that no other interactions were observed, $F_s(2, 28) = 2.01$, $ps > .16$, $\eta_p^2 < 0.13$, any changes in the effect of Cue Condition with Set Size on $p(\text{Guess})$ were subtle enough to not produce a corresponding change in other sources of memory error, and so we did not analyze this potential interaction further.

(Appendix continues)

Three-component model analysis: 600-ms sample duration.

When memory stimuli were presented for 600 ms, Set Size again affected all aspects of memory performance, $F_s(1, 14) = 9.79$, $p_s < .007$, $\eta_p^2 > 0.41$, as expected. Critically, Cue Condition again exhibited main effects on $p(\text{Correct Access})$, $F(1, 28) = 23.35$, $p < .001$, $\eta_p^2 = 0.63$, and $p(\text{Swap})$, $F(1, 28) = 3.50$, $p = .04$, $\eta_p^2 = 0.20$. However, interactions between Cue Condition and Set Size for $p(\text{Correct Access})$, $p(\text{Swap})$, and $p(\text{Guess})$, $F_s(2, 28) > 3.31$, $p_s < .05$, $\eta_p^2 > 0.19$, indicated that memory cueing effects were best examined separately for each Set Size.

At Set Size 2, memory cues affected $p(\text{Correct Access})$, $F(2, 28) = 6.48$, $p = .005$, $\eta_p^2 = 0.32$, and $p(\text{Swap})$, $F(2, 28) = 6.18$, $p = .01$, $\eta_p^2 = 0.31$. Both-feature cues led to higher $p(\text{Correct Access})$ than either Orientation Cues, $F(1, 14) = 11.53$, $p = .004$, $\eta_p^2 = 0.45$, or Location Cues, $F(1, 14) = 5.55$, $p = .03$, $\eta_p^2 = 0.28$. Correspondingly, $p(\text{Swap})$ was lower for both-feature cues relative to Orientation Cues, $F(1, 14) = 11.85$, $p = .004$, $\eta_p^2 = 0.46$, and Location Cues, $F(1, 14) = 5.10$, $p = .04$, $\eta_p^2 = 0.27$. Set Size 2, then, exhibited a straightforward effect of accessibility: cues with more features prevented swap errors and promoted correct item retrieval.

At Set Size 5, Cue Condition again affected $p(\text{Correct Access})$, $F(2, 28) = 15.56$, $p < .001$, $\eta_p^2 = 0.53$, but this was accompanied by an effect on $p(\text{Guess})$, $F(2, 28) = 3.57$, $p = .042$, $\eta_p^2 = 0.20$,

and only a marginal effect on $p(\text{Swap})$, $F(2, 28) = 3.12$, $p = .06$, $\eta_p^2 = 0.18$. As we observed in Experiment 1, at this larger Set Size, both-feature cues increased $p(\text{Correct Access})$ compared to Orientation Cues, $F(1, 14) = 13.67$, $p = .002$, $\eta_p^2 = 0.49$, but not compared to Location cues, $F(1, 14) = 0.10$, $p = .76$, $\eta_p^2 = 0.01$. Importantly, only $p(\text{Guess})$ mirrored this pattern, with Orientation Cue trials leading to more guessing, $F(1, 14) = 6.22$, $p = .026$, $\eta_p^2 = 0.31$, than both-feature cue trials, whereas no such difference was present for $p(\text{Swap})$, $F(1, 14) = 0.15$, $p = .71$, $\eta_p^2 = 0.01$. We did observe, however, that Location-Cue trials had fewer swaps than Both-Cue trials, $F(1, 14) = 4.43$, $p = .05$, $\eta_p^2 = 0.24$, but guessing was higher for Location-Cue trials, $F(1, 14) = 4.60$, $p = .05$, $\eta_p^2 = 0.25$, possibly reflecting a more liberal retrieval threshold for Location-Cue than for both-feature cues. Overall, these results are qualitatively quite similar to Experiment 1, where at the larger Set Size, memory retrieval with a location cue was equal to memory retrieval with a location cue that also contained information about an item's nonspatial features. One notable caveat is that the improvement in $p(\text{Correct Access})$ at Set Size 5 with richer retrieval cues reduced guess responses instead of swap responses.

Received May 21, 2016

Revision received December 15, 2016

Accepted December 16, 2016 ■