

More Than a Filter: Feature-Based Attention Regulates the Distribution of Visual Working Memory Resources

Blaire Dube
University of Guelph

Stephen M. Emrich
Brock University

Naseem Al-Aidroos
University of Guelph

Across 2 experiments we revisited the filter account of how feature-based attention regulates visual working memory (VWM). Originally drawing from discrete-capacity (“slot”) models, the filter account proposes that attention operates like the “bouncer in the brain,” preventing distracting information from being encoded so that VWM resources are reserved for relevant information. Given recent challenges to the assumptions of discrete-capacity models, we investigated whether feature-based attention plays a broader role in regulating memory. Both experiments used partial report tasks in which participants memorized the colors of circle and square stimuli, and we provided a feature-based goal by manipulating the likelihood that 1 shape would be probed over the other across a range of probabilities. By decomposing participants’ responses using mixture and variable-precision models, we estimated the contributions of guesses, nontarget responses, and imprecise memory representations to their errors. Consistent with the filter account, participants were less likely to guess when the probed memory item matched the feature-based goal. Interestingly, this effect varied with goal strength, even across high probabilities where goal-matching information should always be prioritized, demonstrating strategic control over filter strength. Beyond this effect of attention on *which* stimuli were encoded, we also observed effects on *how* they were encoded: Estimates of both memory precision and nontarget errors varied continuously with feature-based attention. The results offer support for an extension to the filter account, where feature-based attention dynamically regulates the distribution of resources within working memory so that the most relevant items are encoded with the greatest precision.

Public Significance Statement

We have the ability to briefly remember small amounts of visual information using a memory system known as visual working memory (VWM). This system is foundational to many human behaviors; however, VWM capacity is severely limited. How is it that we are able to make the most of this important but limited resource? A longstanding theory states that this is accomplished using attention: When you have a behavioral goal (e.g., remembering the cars in your blind spot while driving), attention determines which information is granted access to memory by filtering out distraction. Here we show that attention does more than just determine what information is remembered or forgotten: Attention also strategically regulates how well (i.e., how precisely) that information is represented in memory. These findings extend the function of attention and provide an important step in understanding the connections between attention and memory.

Keywords: filtering, feature-based attention, visual working memory

Visual working memory (VWM) is foundational to many human behaviors. For example, VWM plays an important role in supporting everyday tasks such as visual search (Wolfe, 1994), and

individual differences in VWM ability consistently correlate with fluid intelligence (for a review, see Brady, Konkle, & Alvarez, 2011). However, the capacity of VWM is severely limited (Cowan, 2001; Ma, Husain, & Bays, 2014). How can this memory system usefully support behavior when it can only precisely represent an amount of information equivalent to approximately three or four visual objects? One long-standing theory attributes this capability to the role that feature-based attention plays in filtering access to memory (McNab & Klingberg, 2008), an account made popular by analogy to nightclubs: Feature-based attention is the bouncer in the brain that ensures that only the best information gains access to VWM (Awh & Vogel, 2008), in turn optimizing the use of this limited resource. However, nearly a decade later, there has been a shift in the way we conceptualize VWM architecture, with much of

Blaire Dube, Department of Psychology, University of Guelph; Stephen M. Emrich, Department of Psychology, Brock University; Naseem Al-Aidroos, Department of Psychology, University of Guelph.

This work was supported by a Natural Sciences and Engineering Research Council (NSERC) Discovery Grant to Naseem Al-Aidroos and a NSERC Canada Graduate Scholarship to Blaire Dube.

Correspondence concerning this article should be addressed to Blaire Dube, Department of Psychology, University of Guelph, 50 Stone Road East, Guelph, Ontario, Canada, N1G 2W1. E-mail: bdube@uoguelph.ca

the conversation now focusing on the flexible distribution of resources within memory—a trait not readily attributable to the filter mechanism as it currently stands. Thus, in the present study we revisit the filter account of how feature-based attention regulates VWM performance.

Origin of the Filter Account

The foundation to the filter account was derived from work by Vogel and Machizawa (2004), who recorded event-related potentials to the onset of lateralized visual memory arrays. Shortly after the onset of the array, and persisting through the memory retention interval, a negative voltage was recorded over the hemisphere contralateral to the memorized array. The amplitude of this component was modulated by the number of remembered items and reached an asymptote at the same array set size for which behavioral measures of VWM performance also plateau. Therefore, the authors suggested that this lateralized activity, which they referred to as the contralateral delay activity (CDA), subserves the encoding and maintenance of objects within VWM.

Vogel and colleagues capitalized on this indirect measure of memory activity by using it to examine the role of distraction in VWM performance (Vogel, McCollough, & Machizawa, 2005). Here, participants were asked to memorize the orientations of a set of target lines presented in one color (e.g., red) while ignoring any distractor lines presented in another color (e.g., blue), inducing participants to prioritize items using feature-based attention. To assess whether participants were able to successfully prevent distractors from being encoded in memory, Vogel and colleagues tested whether varying the number of distractors affected the amplitude of the CDA. It is interesting to note that they found that not only did some participants fail to ignore distractor items, but the ability to ignore distractor items correlated with individual differences in overall performance on the memory task. The authors reasoned that individual differences in VWM performance may not reflect how much information an individual can store; indeed, under some circumstances low-performing individuals may store more information (i.e., both targets and distractors) than high-performing individuals. Instead, differences in performance may be primarily determined by an individuals' ability to filter distraction.

When relevant information is presented among distractors, how exactly is its selection accomplished? McNab and Klingberg (2008) addressed this question with a functional magnetic resonance imaging (fMRI) study designed to identify the neural activity associated with the adoption of a feature-based goal and the accompanying preparation to filter out nonmatching distractors presented at encoding. Some trials of this study, referred to as distraction trials, paralleled the design used by Vogel et al. (2005), although adapted for fMRI. On these trials, participants were instructed to memorize the locations of a set of red circles that were either presented alone or with distracting yellow circles. By measuring the effect of the presence of distracting stimuli on the amplitude of blood-oxygen-level dependent (BOLD) activity in the parietal cortex—the fMRI index of VWM encoding and maintenance (Todd & Marois, 2004)—McNab and Klingberg were able to assess the extent to which participants unnecessarily stored the distracting yellow circles in memory. The authors also extended this paradigm by including no-distraction trials for which partici-

pants were instructed to memorize the locations of all stimuli, both red and yellow. By testing for differences in BOLD activity after the instruction to either filter (distraction trials) or not (no distraction trials), McNab and Klingberg identified areas of prefrontal cortex and the basal ganglia that play a role in the preparation to filter distracting information. Of note, activity in the globus pallidus in particular correlated inversely with the parietal measure of unnecessary storage, suggesting that this region plays a critical role in filtering information to determine whether or not access to VWM is granted. As Awh and Vogel eloquently suggested in 2008, when a feature-based goal is adopted, the prefrontal cortex writes the guest list for the nightclub and passes it along to the globus pallidus—the bouncer in charge of turning away unwanted guests. Together with the findings of Vogel et al. (2005), this study provides a clear picture of the cognitive and neural mechanisms through which feature-based attention regulates VWM performance by filtering distraction, and this filtering account continues to actively influence investigations of VWM performance (Jost & Mayr, 2016; Liesefeld, Liesefeld, & Zimmer, 2014; Vissers, van Driel, & Slagter, 2016).

Reasons to Question the Filter Account

When the filtering account was first proposed, the leading framework defining VWM architecture was the longstanding discrete capacity “slot” model (Luck & Vogel, 1997), the fundamental claim of which is that VWM is divided into a limited number of slots capable of representing an item with fixed precision. According to this model, each individual's upper limit on VWM storage is defined by the number of equal-precision items that he or she can maintain (i.e., the number of slots), and items that exceed this fixed capacity are forgotten. Under this assumption, ensuring that these slots are filled with the correct information—a goal that can be accomplished entirely with filtering—is indeed an effective way to control performance. However, the assumptions of the discrete capacity model have since been challenged, primarily based on the results of a series of experiments using partial report working memory tasks (Bays, Catalao, & Husain, 2009; Wilken & Ma, 2004; Zhang & Luck, 2008a). In these tasks, participants report one aspect of the probed memory item on a continuous scale (e.g., by choosing a color on a color wheel) and, by examining the distribution of a participant's responses, it is possible to identify how different sources of error contributed to their response biases. For example, it is possible to estimate the frequency with which the participant failed to encode the probed item and responded with a guess (guess rate); how often they reported the incorrect item from the memory array (nontarget errors); and, for items they did correctly remember, how precisely those items were represented within memory (precision; Bays et al., 2009; Zhang & Luck, 2008a; see also van den Berg, Awh, & Ma, 2014 for additional parameters that can be estimated from partial report tasks). One of the major discoveries emerging from such tasks is that memory precision is not fixed: When participants are told that one item in the memory array is more likely to be probed than others, that item is represented with greater precision (Zhang & Luck, 2008a), and precision also varies with the number of items that participants are asked to remember (Bays & Husain, 2008). On the basis of these findings, variants of discrete capacity models have been proposed that allow for changes in precision by,

for example, using multiple memory slots to encode a single item. In addition, a competing class of models, termed *continuous resource models*, has been proposed. According to this class of models, memory is not restricted to a fixed number of slots but can be divided across any number of items, with each division leading to a reduction of precision (Alvarez & Cavanagh, 2004; Wilken & Ma, 2004). Thus, beyond influencing whether an item is encoded in memory or not (i.e., filtering), feature-based attention may also regulate the distribution of resources across items within memory.

Recent work has offered further support for the notion that VWM resources can be distributed flexibly among items based on their task relevance (Emrich, Lockhart, & Al-Aidroos, in press; Klyszejko, Rahmati, & Curtis, 2014; Zokaei, Gorgoraptis, Bahrami, Bays, & Husain, 2011). Using a varied number of spatial cues across multiple set sizes, Emrich and colleagues manipulated the likelihood that certain items in the memory array would be probed, with cue validity ranging from 8% to 100%. They observed that higher cue validity was associated with higher subsequent memory precision. It is interesting to note that the relationship between cue validity and precision followed a power law, which is the expected relationship if memory is accomplished using a neural population code that is distributed across items based on cue validity (Bays, 2014). Equivalent patterns were also observed for other measures of memory performance, including guess rates and nontarget errors. In a similar study, Klyszejko and colleagues found that varying the probability of feature-based cues affected the standard deviation of participants' response errors, likely reflecting a change in precision, although a full investigation of response errors was not conducted (Klyszejko et al., 2014). Likewise, when participants perceive two sets of dots simultaneously moving in two distinct directions, memory responses are more precise when one direction of motion is cued (Zokaei et al., 2011). Together, these studies reinforce the flexibility of VWM representations and point to a broader role for attention (generally speaking) in regulating VWM performance.

The assumptions that were widely accepted at the time the filter theory was proposed have since been challenged, and examinations of VWM architecture have since centered more on the flexible, dynamic nature of encoding and resource allocation—a process that is not easily accomplished through filtering alone. In fact, when we revisit the data on which the original filter account was based, there is evidence that activity in the globus pallidus of the basal ganglia associated with the preparation to filter may actually serve a broader function than simply blocking the encoding of distractors. Specifically, Emrich and Busseri (2015) reanalyzed the original dataset that linked decreased activity in the globus pallidus to the inefficient encoding of irrelevant information (McNab & Klingberg, 2008). Their analysis revealed that, although this neural activity is a strong predictor of the encoding of distracting information, the encoding of distracting information is not an important or sufficient predictor of memory performance once other factors are accounted for. Rather, performance was strongly predicted by the level of activity in the globus pallidus directly, suggesting that this activity (previously associated with the preparation to filter information out) may actually be involved in more processes than once thought. Along with the numerous developments about VWM architecture, this development high-

lights the importance of revisiting the filter account. Specifically, there seems to be a need to broaden our understanding of the ways in which attention—particularly feature-based attention—contributes to VWM performance.

The Current Study

In the present study we reexamine the role of feature-based attention in regulating VWM performance: Does feature-based attention affect aspects of performance unrelated to filtering distractors; can evidence of filtering be found after controlling for any such other effects; and, if so, how flexible is control over feature-based attention? Given that these questions serve to extend our understanding of the filter account, we have adapted the paradigms used in the original papers. Namely, in McNab and Klingberg (2008) and Vogel and colleagues (2005), the memory arrays included distractors that were entirely task irrelevant. In contrast to this, in the present study each individual item in a given memory array is always at least minimally relevant. This allows us to assess participants' memory for these items relative to those that are more highly prioritized. In Experiment 1, participants completed a partial report procedure in which they memorized the colors of two squares and two circles (Experiment 1a) or three squares and three circles (Experiment 1b), and we manipulated feature-based attention by instructing participants that one shape was always more likely to be probed for recall than the other. To assess the flexibility of attention, including filtering, we manipulated the strength of the feature-based goal so that, across blocks, the likelihood of probing a high-probability shape was 60%, 70%, 80%, or 90%. According to the original description of the filtering account, this probability manipulation should have minimal impact on memory for high-probability items. Specifically, high-probability items are always more likely to be probed than low-probability items; thus, if the only function of feature-based attention is to serve as a simple filter, granting or denying access to VWM, these items should always receive priority for the limited slots in memory regardless of the probability condition. Consequently, Experiment 1 was designed to test for differences only in high-probability items. We adjusted the design in Experiment 2 to include all memory items, investigating performance across a wider range of probabilities (both high and low). For both experiments, we also estimated the contribution of guesses, nontarget responses, and memory precision to participants' errors. To the extent that feature-based attention acts as a filter, we should see effects of attention on guess rate. Beyond determining which items are, and are not, represented in memory, attention may also affect how items are represented, and we assessed such changes through nontarget errors and precision. To preview our results, the value of the attentional goal affected all measured aspects of memory performance, suggesting that feature-based attention flexibly regulates multiple aspects of VWM performance both by filtering and by redistributing resources within memory.

Experiment 1

Method

Experiments. Experiments 1a and 1b differed only in the size of the memory array; four versus six colored shapes, respectively.

Experiment 1b was conducted to address a potential ceiling effect noted in Experiment 1a. We analyze and report both experiments together.

Participants. Sixty-six undergraduate students (33 in each of Experiment 1a and Experiment 1b) from the University of Guelph between the ages of 17 and 27 years ($M = 18.40$) participated for partial course credit. All participants reported having normal or corrected-to-normal vision and no color blindness.

Stimuli and apparatus. The stimuli and apparatus were identical in Experiments 1a and 1b aside from noted exceptions. Participants reported the colors of memory items using a 360° color wheel. The color wheel used in all reported experiments was created based on the CIELAB color space with the intention that colors would differ only in hue and not in luminance, and the perceptual difference between two adjacent colors anywhere on the wheel would be comparable. In Experiment 1a, this color wheel was not calibrated to the specific monitors used; however, this limitation was remedied in Experiments 1b and 2, in which experimental monitors were assessed using a colorimeter (Konica Minolta CS-100A; Konica Minolta Sensing Americas, Inc., Ramsey, New Jersey) and the color wheel was adjusted for each monitor. The colors of memory items on each trial were randomly selected from the color wheel with the constraint that all colors were separated by a minimum of 40° on the wheel. All stimuli were viewed on a $1,280 \times 1,024$ cathode ray tube (CRT) monitor using a 75-Hz refresh rate with a gray background, and viewing distance was fixed at 57 cm using a head and chin rest.

Procedure. On each trial participants were shown a memory array for 300 msec (see Figure 1). The memory array in Experiment 1a comprised two squares and two circles (each 1.2° in width and height) placed randomly in four of eight possible positions

spaced equally around an invisible circle with a 5.5° diameter that was centered on a central fixation point with a 0.07° radius. The memory array in Experiment 1b comprised three squares and three circles, and we increased the eccentricity of stimuli to 8° to reduce crowding. After a 900-msec delay, a probe screen was displayed, which presented the black outline of all four shapes in their original spatial locations surrounded by the color wheel, the outer edge of which was 12° from fixation. Participants were asked to report the color of a probed (outlined in a thick black line) shape as it had appeared in the preceding memory array by clicking the corresponding color on the color wheel. The orientation of the color wheel was randomized on each trial. This probe screen remained visible until participants made a response and was followed by a 500-msec intertrial interval.

Participants were instructed that for the duration of the experiment, memory items that were one shape (i.e., either squares or circles) were more likely to be probed than the other; shape was pseudo-counterbalanced across participants. Participants completed four blocks of 150 trials, and the likelihood that a high-probability shape would be probed was manipulated across blocks to be 60%, 70%, 80%, or 90%. Participants were informed of this probability at the outset of each block, and the order of the blocks was randomized for each participant. Of note, this procedure results in a large number of trials in which a high-probability shape is probed and only a small number of trials in which the low-probability shape is probed, especially for the 90% block. For Experiments 1a and 1b, we were primarily interested in how the probability manipulation affected memory for high-probability items; thus, our analyses focus on only trials in which a high-probability item was probed, resulting in a fully crossed, mixed design with one within-subjects factor (Proba-

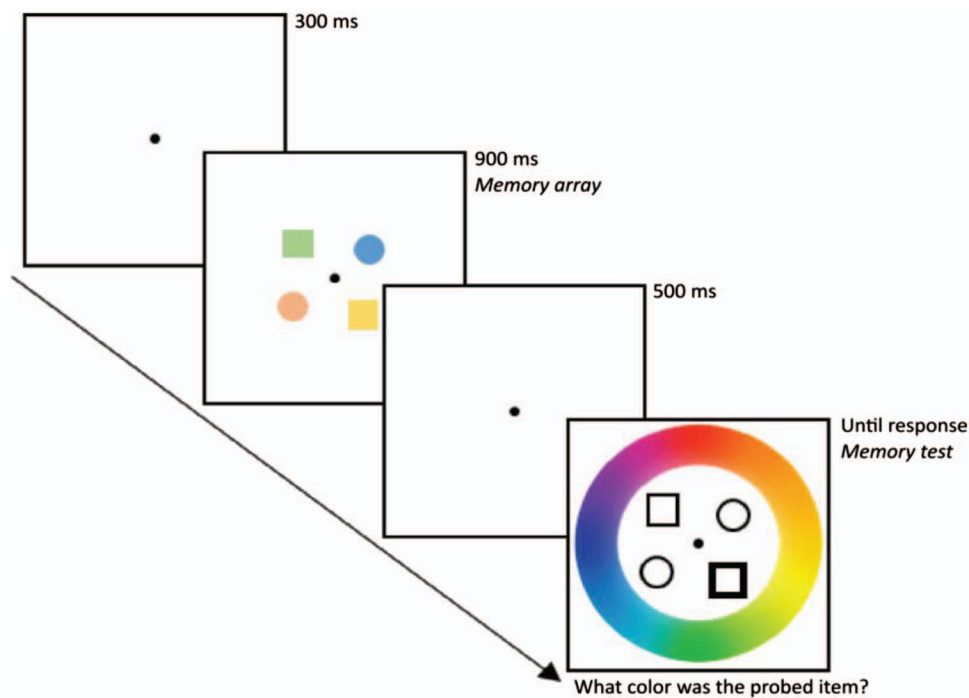


Figure 1. Example trial sequence for Experiment 1a. Actual background color was gray. See the online article for the color version of this figure.

bility: 60%, 70%, 80%, 90%) and one between-subjects factor (Experiment: 1a vs. 1b).

Analysis. In all reported experiments, we assess VWM performance using several measures all based on raw response error (i.e., the angular degree of distance between the color of the probed item and the participant's response). We first calculated the variability of responses for each probability condition by taking the standard deviation of participants' response errors (SD_{response}). Here, lower values reflect responses distributed more closely around the target color, which we interpret as generally reflecting more precise representations in VWM.

Using a three-parameter mixture model (Bays et al., 2009), we further broke raw response errors down into three separate components that each reflect different aspects of VWM performance. This model asserts that three different distributions contribute to the likelihood of a given response: a uniform distribution reflecting the proportion of random guesses; a von Mises distribution reflecting the proportion of responses centered around the actual color of the probed item, the standard deviation of which is inversely related to VWM fidelity; and the proportion of reports that are centered around nonprobed items (i.e., "swap" reports of nontargets, likely a consequence of a spatial binding error). To account for these nontarget errors, Bays and colleagues developed this three-parameter mixture model from the standard two-parameter mixture model proposed by Zhang and Luck (2008a), which included only a uniform distribution and a von Mises distribution. For the precise mathematical definition of this model, see Equation 2 in Bays et al. (2009). We used maximum likelihood estimation (MLE) to decompose raw errors into these three parameters via MATLAB and the MemToolBox library (Suchow, Brady, Fougny, & Alvarez, 2013).

It has been argued that some variability in memory performance is a consequence of random fluctuations in memory fidelity across items within a trial, and this source of variance is captured in a class of models termed *variable precision models* (Fougny, Suchow, & Alvarez, 2012; van den Berg, Shin, Chou, George, & Ma, 2012). To assess whether any observed effects of our probability manipulation are influenced by also accounting for random fluctuations in fidelity, we modeled SD_{response} using a variable precision model that has been previously shown to provide a good account of memory performance—the VP-A-NT model reported by van den Berg et al. (2014). This model assumes that memory is a continuous resource distributed across all memory items and, as such, it does not include a parameter reflecting random guesses. However, this model does estimate how an item's precision (\bar{J}) changes as a power-law function (α) of the number of memorized items. Furthermore, along with the contribution of random fluctuations in memory fidelity, this model estimates how the likelihood of making nontarget errors changes with set size (NT_{slope}). Model fitting was accomplished via MATLAB and a collection of functions provided by van den Berg et al. (2014). Of note, the code was adjusted to allow for our probability manipulation to replace the traditional set size manipulation; to accomplish this we input "set size" as the inverse of probability (Emrich et al., in press). Our primary interest was whether this variable precision model would reveal comparable effects of our probability manipulation on precision (α) and nontarget errors (NT_{slope}) to those revealed by the three-parameter mixture model.

Results and Discussion

Raw error. As an initial assessment of memory performance, we examined the variability of participants' errors (SD_{response}) on trials in which a high-probability item was probed (see Figure 2). A 2 (Experiment: 1a vs. 1b) \times 4 (Probability: 60%, 70%, 80%, 90%) mixed analysis of variance (ANOVA) on SD_{response} revealed statistically significant main effects of experiment, $F(1, 64) = 18.57, p < .001, \eta_p^2 = .225$, and probability, $F(3, 192) = 30.89, p < .001, \eta_p^2 = .326$; the interaction was not statistically significant, $F(3, 192) = 1.76, p = .157$. As would be expected, the increase in memory set size from Experiment 1a to 1b (four items vs. six) was associated with an overall increase in SD_{response} (i.e., a decrement in memory performance). Interestingly, this measure of performance was also affected by the probability manipulation. Looking at Figure 2, in both experiments participants' memory for high-probability items (i.e., participants' SD_{response} values) appears to consistently increase as the value of the feature-based attention goal increases, an interpretation supported by a statistically significant linear within-subject contrast, $F(1, 64) = 67.93, p < .001, \eta_p^2 = .515$, which did not interact with experiment, $F(1, 64) = 1.23, p = .273$. This pattern is difficult to reconcile with the role of feature-based attention as described in the original filter account. Here, high-probability items were always more likely to be probed than low-probability items and should always have had priority for the limited slots in memory regardless of the strength of the probability manipulation. However, the finding that it matters how high a high-priority target is is consistent with feature-based attention causing goal-relevant items to be stored with proportionately greater fidelity as their value increases.

Mixture model. One way to reconcile the observed effect of probability on SD_{response} with a pure filtering account is to assume that participants suboptimally filtered. Perhaps participants can




Figure 2. The standard deviation of participants' errors (SD_{response}) for high-probability trials in Experiments 1a and 1b. Values on the x-axis represent the likelihood of probing any high-probability item—this value divided by the set size for the relevant shape (i.e., two or three in Experiments 1a and 1b, respectively) is the probability of probing any individual high-probability item. Error bars in this figure, as well as in all subsequent figures, are adjusted within-subject standard error (Morey, 2008).

control the effectiveness of the filter to sometimes allow low-probability items to be encoded in memory at the expense of high-probability items, and they varied this effectiveness to match the value of the feature-based goal (i.e., our probability manipulation). To investigate this interpretation, and, more generally, to better understand how the probability manipulation affected memory performance, we next decomposed participants' errors into distributions reflecting guesses, nontarget responses, and memory precision using a three-component mixture model (Bays et al., 2009). According to a pure filtering account, the value of the feature-based attentional goal should primarily affect the rate of guesses. Alternatively, if feature-based attention affects not only which items are encoded in memory but also how those items are encoded, then we should also observe effects on swap rate and precision. The parameter estimates of these mixture model analyses are presented in Figure 3 and reveal that all three parameters were somewhat affected by the probability manipulation.

Before conducting the analyses reported next, we removed any condition data for each participant with model fits characterized as outliers. Specifically, we removed model estimates if guess rate equaled zero (these tended to be associated with an inflated corresponding standard deviation parameter) or if standard deviation fell below 5° or above 105° (the value associated with random selection of colors). If more than 25% of a given participant's data were removed, then we removed the entire participant. This resulted in the removal of one complete participant from Experiment 1b as well as the removal of a single condition for four different participants in Experiment 1b. All Experiment 1a data were retained. The pattern of statistically significant results reported next does not change when outlying model fits are included.

Guess rates. First, we subjected guess rates to a 2 (Experiment: 1a vs. 1b) \times 4 (Probability: 60%, 70%, 80%, and 90%) mixed ANOVA, revealing significant main effects of experiment, $F(1, 55) = 16.33, p < .001, \eta_p^2 = .229$, and probability, $F(3, 165) = 17.78, p < .001, \eta_p^2 = .244$; the interaction did not reach statistical significance. Similar to the results for SD_{response} , guess rates also changed as a function of the probability manipulation.

Nontarget responses. A similar pattern was observed for the proportion of nontarget responses using the same 2×4 ANOVA (corrected degrees of freedom used where assumption of sphericity was violated). This analysis revealed statistically significant main effects of experiment, $F(1, 55) = 9.21, p = .004, \eta_p^2 = .143$, and

probability, $F(2.46, 135.22) = 3.27, p = .032, \eta_p^2 = .056$, and the interaction did not reach statistical significance, $F(3, 165) = 2.05, p = .108, \eta_p^2 = .036$.

Standard deviation. The same 2×4 ANOVA on the standard deviation parameter of the mixture model yielded a slightly different result. Although there was a statistically significant main effect of experiment, $F(1, 55) = 16.55, p < .001, \eta_p^2 = .231$, the main effect of probability was marginally nonsignificant, $F(2.38, 130.67) = 2.28, p = .082, \eta_p^2 = .231$. The interaction term once again did not reach statistical significance ($F < 1$).

Contrasts. To better understand the nature of the observed effects of probability, including the marginally nonsignificant effect on standard deviation, we assessed the within-subject linear contrasts associated with the ANOVAs earlier reported. Of note, our intention with these analyses was not to determine what function best describes the changes in each parameter with probability; indeed, there is reason to avoid overinterpreting such group-level patterns in mixture model parameters because they can distort participant-level effects (van den Berg & Ma, 2014). Rather, our goal was to confirm that any effects of probability were related to the level of the probability manipulation, which itself changed linearly across conditions. All three parameters were associated with statistically significant linear contrasts: Guess rate, $F(1, 55) = 43.22, p < .001, \eta_p^2 = .44$; nontarget errors, $F(1, 55) = 2.57, p = .007, \eta_p^2 = .123$; and standard deviation, $F(1, 55) = 4.71, p = .034, \eta_p^2 = .079$, all decreased linearly with increases in probability. No interactions with experiment reached statistical significance; nontarget error rate, $F(1, 55) = 2.57, p = .115, \eta_p^2 = .045$ and both other F values < 1 . Taken together, these results provide support for two conclusions. Consistent with filtering, changes in the value of a feature-based attentional goal affected the likelihood that high-probability targets were successfully encoded and recalled, as indicated by the effect of probability on guess rate. Probability also affected the likelihood of nontarget errors and precision, suggesting that, beyond filtering, feature-based attention affects how items are encoded in memory.

Variable precision model. We next assessed the generalizability of the reported effects of probability by decomposing participants' errors using a competing model of memory performance, the VP-A-NT variable precision model described in van den Berg et al. (2014). Our primary interest for this analysis was whether probability affected estimates of memory precision and

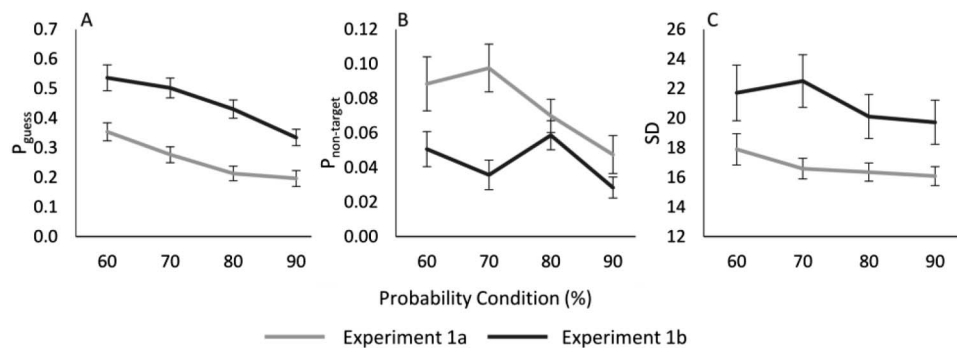


Figure 3. Results of the three-parameter mixture model for Experiments 1a and 1b. (A) Probability of a random guess (P_{guess}) across all probability conditions. (B) Probability of a nontarget report ($P_{\text{nontarget}}$) across all probability conditions. (C) Standard deviations across all probability conditions.

nontarget errors. For this model, the relationship between memory precision (\bar{J}) and probability is assumed to follow a power-law function; as such, an effect of probability on precision will cause a nonzero value in the estimated exponent parameter of the power-law function (α). An effect of probability on nontarget responses is similarly assessed by testing for a nonzero value of the NT_{slope} parameter. For a full list of the VP-A-NT model output parameter values, including those not analyzed here, see Table 1. The observed changes in memory precision with probability are plotted in Figure 4. As can be seen in this figure, precision tends to increase (better performance) with increases in probability, and this interpretation is supported by α parameters that are statistically significantly different from zero for both Experiments 1a, $t(32) = -5.66, p < .001$, and 1b, $t(32) = -7.48, p < .001$. Furthermore, a comparison of the α parameters from each experiment yielded a nonsignificant result, $t(64) = -1.53, p = .130$, indicating minimal differences in how precision changes with probability between the experiments. Analysis of the NT_{slope} parameter revealed the same pattern of effects. In both Experiments 1a and 1b, NT_{slope} was significantly greater than zero, $t(32) = 5.53, p < .001$ and $t(32) = 4.72, p < .001$, respectively, and did not statistically significantly differ across experiments, $t(48.674) = -1.597, p = .121$. Thus, as was observed in the mixture model analysis, the results of this variable precision model analysis support the conclusion that feature-based attention impacts aspects of working memory performance other than filtering, in particular memory precision and the likelihood of nontarget errors.

Summary

Together, the reported analyses of $SD_{response}$, mixture model parameters, and variable precision model parameters reveal a consistent interpretation of how feature-based attention influences VWM performance. On the one hand, in line with filtering accounts, the observed effect of probability on guess rate indicates that one role of attention is to influence whether or not an item will be encoded in memory. It is interesting to note that we advance our understanding of this attentional filter by demonstrating that it can be used flexibly so that the likelihood of encoding an item in memory varies with the probability with which that item will be probed. On the other hand, feature-based attention also regulates other aspects of VWM performance. As the value of a feature-based attentional goal increases, items matching that goal are represented with greater precision and are less likely to be misreported as a nontarget. Thus, similar to the effect on guess rate, here attention regulates performance flexibly based on the value of the attentional goal.

Table 1
Average VP-A-NT Parameter Values

Experiment	VP-A-NT Parameters				
	\bar{J}_1	α	T	κ_r	NT_{slope}
1a	13,291.52	-1.46	9,049.21	27.65	.00011
1b	3,973.79	-2.03	6,089.24	24.65	.00017
2	702.50	-1.30	2,097.76	23.66	.00019

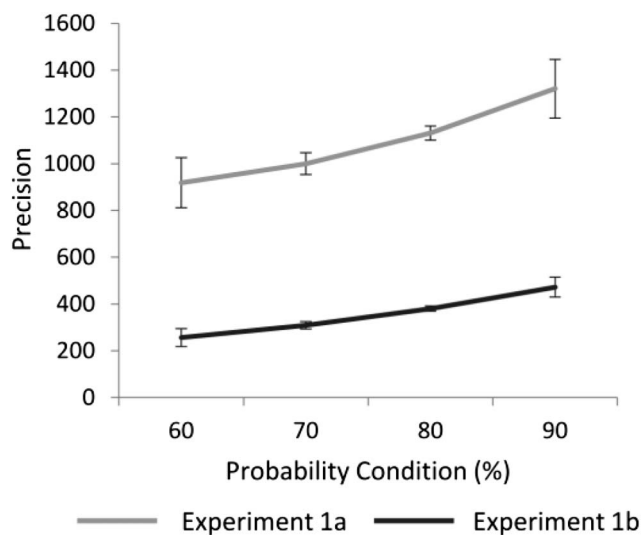


Figure 4. Precision (\bar{J}) across each probability condition for Experiments 1a and 1b. Larger values indicate greater precision.

Is this flexibility achieved by strengthening or weakening attention to match the value of the behavioral goal on every individual trial, or might this flexibility be achieved by switching strategies across trials? For example, did participants prioritize high-probability items on the majority of trials and prioritize low-probability items on a minority of trials? Although such a probability-matching strategy could explain the observed effects on guess rate and nontarget response rate, it cannot readily account for the effects on precision because this parameter is calculated as the variance of response error on only those trials in which the target item was successfully recalled. Thus, the present results tentatively favor the flexible deployment of attention tailored to an individual trial; however, our conclusions regarding filtering accounts of attention ultimately are not dependent on this conclusion.

Experiment 2

In Experiments 1a and 1b, we focused our experimental design and analyses on participants' ability to recall high-probability items because these items provided the strongest test of the role of feature-based attention in regulating VWM performance as described by the filter account. The primary goal of Experiment 2 was to extend these findings through a more comprehensive examination of memory performance across a wider range of probabilities. This goal was principally motivated by the observation that memory precision (i.e., mixture model standard deviation) varied as a function of the feature-based attentional goal, suggesting that VWM resources were strategically distributed among the items in memory based on the likelihood that each item would be probed. We recently reported a similar effect from the manipulation of spatial attention (Emrich et al., in press) and further demonstrated that precision changed as a power-law function of probe probability, which is the predicted relationship if memory is distributed by dividing a neural population code among multiple VWM representations (Bays, 2014). Accordingly, in Experiment 2, we assessed the effects of feature-based attention across a wider

range of probabilities by fitting each measure of memory performance to probability using a power law and testing for a nonzero exponent parameter. This examination of performance on low-probability trials was not possible in Experiments 1a and 1b because of an insufficient number of trials to adequately model performance. In Experiment 2 we overcame this limitation by separately manipulating the number of high- and low-probability items in the memory array. In particular, by including conditions with large numbers of low-probability items, we could probe these items on a large number of trials while still minimizing the probability that any one particular low-probability item would be probed.

Method

Participants. Twenty-seven undergraduate students from the University of Guelph (ages 18–23 years, $M = 18.89$) participated for partial course credit and \$10. Two participants did not complete the experiment, resulting in a final sample of 25.

Stimuli, apparatus, and procedure. The experimental stimuli, apparatus, and procedure were similar to Experiments 1a and 1b. Memory array stimuli were six $1^\circ \times 1^\circ$ colored circles and squares, equally spaced around a central fixation point, visible on screen for 500 msec. Participants were asked to memorize the items and, after a 900-msec delay, a black outline of each stimulus in the original memory array reappeared on the screen surrounded by a color wheel and one item was probed (outlined in bold). Participants indicated the color of the probed shape as it appeared in the initial array by clicking the corresponding color on the color wheel surrounding the test array. Trials were blocked by condition, and at the beginning of each block participants were provided with instructions indicating the high-probability shape, the proportion of the memory array that would match the high-probability shape (one, two, three, or all six items), and how predictive this cue would be of the identity of the probe (ranging from 33% to 100%). There were seven conditions in total, across which we manipulated both the proportion of high- and low-probability items and the predictive validity of the featural cue. Table 2 presents the number of high- and low-probability items in each condition, along with the probability that any individual item would be probed. Participants completed a total of 1,200 trials, with a division of trials among blocks that allowed for 100 high-probability trials per condition.

Analysis. Analyses for Experiment 2 examined the same metrics of VWM performance as in Experiments 1a and 1b. Namely, we analyzed variability in raw response error (SD_{response}) before turning to mixture modeling and the same version of a variable precision model. Rather than comparing performance across the conditions of Experiment 2, we collapsed across condition and treated probability as the predictor of performance. We fit metrics of VWM performance to probability using a power-law function and tested for relationships by comparing the exponent parameter against zero (i.e., no relation). We also report the proportion of variance explained (R^2) for all metrics as a measure of goodness of fit. There were two exceptions to this analysis. We do not report an R^2 value when fitting \bar{J} to probability given that \bar{J} is, itself, calculated using a power-law function. We also tested for linear changes in nontarget errors (NT_{slope}) in the variable precision

Table 2
Probe Likelihoods for High- and Low-Priority Items Across Trial Types

Condition	High-probability items		Low-probability items	
	Number of items	Probability (%)	Number of items	Probability (%)
1	1	100	0	N/A
2	1	50	5	10
3	1	33	5	13.4
4	2	50	0	N/A
5	2	33	4	8.25
6	3	33.3	0	N/A
7	6	16.6	0	N/A

Note. N/A = not applicable.

analysis, rather than fitting to a power-law function, to match the assumptions of the VP-A-NT model.

Results and Discussion

The relationships between metrics of VWM performance and probability are depicted in Figure 5 and Table 1; the latter also reports additional parameters of the variable precision model. For mixture model analyses, we removed outlying model fits using the same criteria as for Experiments 1a and 1b, resulting in the complete removal of one participant, and the removal of a single condition for nine other participants. The mixture model analyses also identified a minimal contribution of nontarget errors (<1% of trials), and we did not examine the relationship between this parameter and probability. Some variance was attributed to nontarget errors in the variable precision model; thus, we did examine this relationship through the NT_{slope} parameter. Similar to the results of Experiments 1a and 1b, all measures of working memory performance varied statistically significantly with probability, with the power-law function providing a good description of the relationship where appropriate (all R^2 values > .7): SD_{response} , $t(24) = -13.12$, $p < .001$, $R^2 = .708$; mixture model guess rate, $t(23) = -5.38$, $p < .001$, $R^2 = .807$; mixture model standard deviation, $t(23) = -3.56$, $p = .002$, $R^2 = .915$; variable precision α , $t(24) = -10.08$, $p < .001$; and variable precision NT_{slope} , $t(24) = 5.38$, $p < .001$.

These results extend the conclusions of Experiments 1a and 1b in three ways. First, we replicate the patterns observed in Experiments 1a and 1b that numerous measures of VWM performance vary with probability, supporting the conclusion that feature-based attention regulates more than just the likelihood of each item being successfully maintained in VWM. Second, changes in performance do not appear to be limited to just high-probability items but extend across the range of measured probabilities (8.25–100%). Third, changes in performance appear to vary with probability as a function of a power law, and this fit appears to be particularly strong for measures of precision (both in the mixture model and variable precision model analyses). Although a formal comparison of competing models is necessary to assess if a power-law function provides the best fit of these data, these results are nevertheless consistent with the flexible division of a continuous VWM resource that is distributed across all memory items based on the probability that each will be probed. These clear effects on

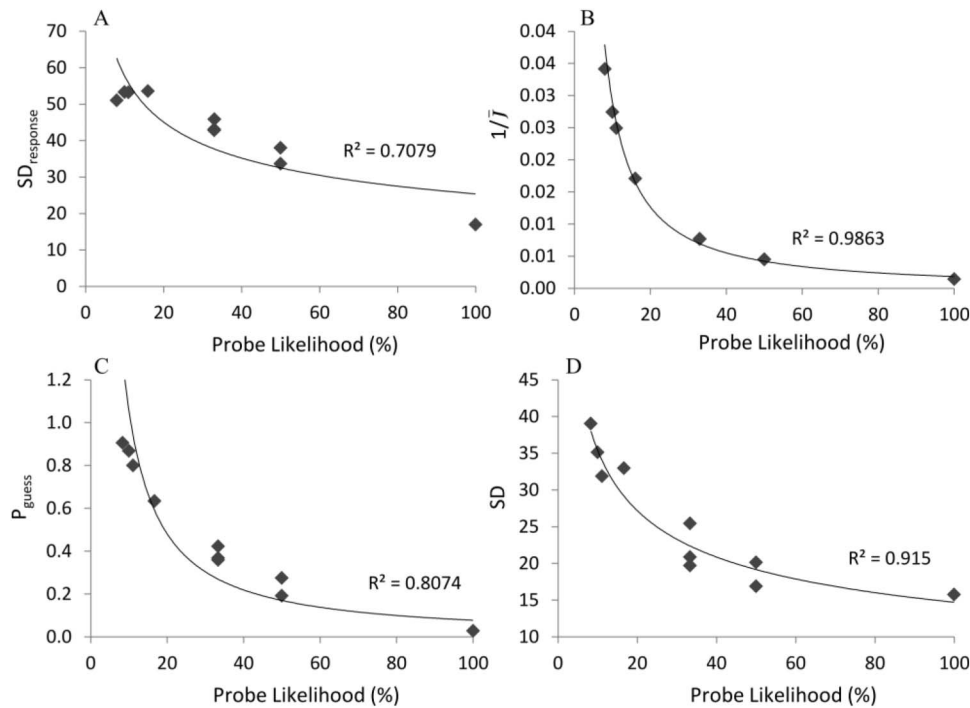


Figure 5. Effects of the probe likelihood manipulation on various measures of VWM performance. Probe likelihood reflects the probability that the probed item would be selected as the probe at the beginning of the trial. (A) SD_{response} at each probe likelihood. (B) For comparability with the measures of precision visualized in this figure, \bar{J} (J bar) is plotted as $1/\bar{J}$ so that smaller values reflect better VWM performance. (C) Estimates of guess rates and (D) standard deviations derived from the mixture model.

precision provide a strong demonstration that, beyond filtering, feature-based attention regulates how precisely an item is represented in memory.

General Discussion

Taken together, the data reported here support the conclusion that feature-based attention does more to regulate VWM performance than simply restrict the encoding of distractors. Beyond filtering, feature-based attention also redistributes resources in VWM across items according to their relative likelihood of being probed in a subsequent memory test. Here, we demonstrated that this distribution can be based on priorities assigned to specific shapes in an array. We focused specifically on feature-based attention given our intention to address the original filter account, support for which has come from tasks in which participants filter items based on their features. However, our larger body of work extends beyond feature-based attention, providing evidence that manipulations of spatial attention have the same consequences for behavior in VWM tasks as those reported here (Emrich et al., in press). Thus, attention, in general, seems to regulate VWM performance by filtering and distributing resources to items in accordance with the goal of the observer.

The suggestion that attention plays a broader role during encoding—and that the breadth of this role is actually better captured by a filter and redistribute account rather than filtering alone—fills an important gap in the existing literature. As we discussed previously, in revisiting the data reported by McNab and Klingberg in

2008 (i.e., the data that served as one of the fundamental bases for the original filter account), our colleagues demonstrated that VWM performance is better predicted directly by activity in the globus pallidus, previously termed *filter set activity*, than indirectly through neural measures of unnecessary storage (Emrich & Busseri, 2015). Emrich and Busseri concluded that filter set activity is associated with more than just filtering, and that perhaps we should “give the bouncer a raise.” We can now speculate with some confidence that variance in this activity also reflects the strategic distribution of resources.

Interpreting Guess Rates

There is ongoing debate regarding the architecture of VWM (e.g., is it a fixed number of discrete slots or a continuous resource that can be flexibility divided?) and, at times, a major point of contention within this debate has been the existence of guessing (Bays et al., 2009; Luck & Vogel, 2013; Zhang & Luck, 2008a). Whereas discrete capacity models of memory predict that guessing should occur any time the memory array exceeds the number of slots, continuous resource models initially argued that, because resources can be divided evenly among all items in an array, observers never actually guess. Here we show that, when you allow for the redistribution of VWM resources by attention, both the discrete and continuous resource models can reasonably predict guessing behavior. That is, by strategically prioritizing memory items, there will at times be circumstances under which guess-

ing on a subset of trials is a consequence of the strategy adopted by the observer (i.e., when some items are assigned zero priority).

In Experiment 2, we assessed how the various parameters of VWM performance (including guess rate) change across a wide range of probabilities, allowing us to examine the nature of these changes. It is interesting to note that all tested measures—VWM error (i.e., SD_{response}), guess rates, and both reported measures of precision (i.e., the SD parameter of the mixture model and \bar{J} values derived from the VP-A-NT model)—were reasonably explained by a power-law relationship with probe probability. Similar power-law relationships for these parameters were also observed by Emrich et al. (in press) after the manipulation of spatial attention. The effect on measures of precision is, perhaps, not surprising given that Bays and Husain (2008) showed a similar relationship between precision and memory load; the authors argued that if VWM items are represented by a neural population code, then the predicted effect of dividing this code among items is a decrease in precision that follows a power law. Why does guess rate exhibit the same relationship? One possibility is that guesses are not a direct consequence of filtering but rather an indirect consequence of the redistribution of resources within VWM. That is, if dividing VWM resources also causes an increased likelihood of failing to encode, maintain, or recall items, then this would explain why guess rate followed the same power-law relationship as precision. Although interesting in theory, this idea will require more investigation to formally assess, including a comprehensive examination of whether a power-law function provides the best explanation of changes in guess rate.

When Does an Attentional Goal Influence Memory Performance?

The effects of attention on VWM performance could reasonably be arising at any stage of information processing from initial perception to VWM maintenance. Indeed, many documented effects of feature-based attention are measured during purely perceptual tasks (i.e., where visual information is accessible to the retina; e.g., Maunsell & Treue, 2006; Moore & Egeth, 1998; Sàenz, Buracas, & Boynton, 2003). Moreover, for arrays of information presented simultaneously, such as the memory arrays in the present study, activity associated with color-based attention is measureable in humans as early as 100 msec after stimulus onset (Zhang & Luck, 2008b), suggesting an early perceptual locus.

Although there are well-established effects of feature-based attention on initial perception, there is reason to question whether these perceptual effects are the source of the biases to VWM performance observed in the present study, in particular those on precision. Specifically, there is little evidence that feature-based attention influences the quality of perceptual representations—an effect that would translate well to our observed effects on VWM precision; rather, the data point to a different effect. Namely, feature-based attention appears to alter the priority of stimuli in the visual field such that stimuli possessing relevant features are processed first (Moore & Egeth, 1998). For example, feature-based attention has large effects in tasks in which multiple stimuli compete for representation, such as visual search (Bacon & Egeth, 1997; Kaptein, Theeuwes, & van der Heijden, 1995), but only minimal effects when stimuli are presented one at a time, such as rapid serial visual presentation tasks (Farell & Pelli, 1993). Thus,

the effects of feature-based attention on initial perception are only measureable when stimuli are in direct competition with one another for representation, which is consistent with behavioral data (Moore & Egeth, 1998) suggesting that feature-based attention alters priority rather than resolution. Consequently, the difference in the nature of the effects of feature-based attention on initial perception and VWM points to the possibility that some of the effects observed in the present study are not simply a byproduct of biases to initial perception.

Furthermore, there is evidence that when information pertinent to an attentional goal is provided after VWM encoding (i.e., in the form of a retro-cue displayed after the removal of a memory array), it can reliably affect VWM performance. For instance, the presentation of a spatial retro cue enhances change detection (Griffin & Nobre, 2003; Sligte, Scholte, & Lamme, 2008) and can affect the estimates of both guess rate as well as precision derived from a partial-report task (Gunseli, van Moorselaar, Meeter, & Olivers, 2015a, 2015b). It is likely that the effects observed in the present study reflect biases of feature-based attention to multiple stages of visual representation. In any case, attention is undoubtedly having a lasting influence on VWM performance in a way that is not accounted for by the existing filter account.

Reciprocal Interactions Between VWM and Feature-Based Attention: Evidence From Visual Search

Beyond understanding the limitations on VWM capacity, assessing the role that attention plays in regulating VWM performance is important for understanding other behaviors. For instance, visual search relies heavily upon both feature-based attention as well as VWM: Attention prioritizes information for access to VWM based on its belonging to a relevant category or possessing a relevant feature and, when selected, the items in a search array are assessed relative to a search target maintained actively in VWM (Bundesen, 1990; Desimone & Duncan, 1995; Wolfe, 1994). Thus, feature-based attention guides visual search and, when additional information (such as which targets are most likely to be relevant) is available, an observer can likely capitalize on this information to improve search. Moreover, as is evident in visual search, the relationship between VWM and attention is not unidirectional: As shown in the present study, feature-based attention biases the contents of VWM; however, there is also a body of work demonstrating that the contents of VWM bias feature-based attention (Olivers, Peters, Houtkamp, & Roelfsema, 2011). Although further investigation is required to assess whether the effects of attention on memory can reciprocally alter the effects of memory on attention (Hollingworth & Hwang, 2013), understanding this potential interplay between attention and memory will be an important step for fully understanding how feature-based attention regulates memory in everyday tasks such as visual search.

The Leaky Filter Account

Thus far, we have described our data as evidence that, beyond acting as a filter that gates access to memory, feature-based attention allows the strategic distribution of resources within memory. However, there is an alternative interpretation. We have long known that the filter mechanism is imperfect (McNab & Kling-

berg, 2008; Vogel et al., 2005), occasionally allowing distracting information to leak into memory. Perhaps participants in our study treated low-probability items on a given trial as distractors to be ignored, and our observed changes in memory precision are simply a consequence of strategic control over the leakiness of the filter used to exclude the distractors. For example, in Experiment 1, as the likelihood of probing high-probability items decreased, the associated precision decreases may simply have been a consequence of more low-probability items being encoded in memory, creating increased competition for memory resources. By this leaky filter account, the decreases in precision are a consequence of the encoding of more items in VWM (with equal precision) rather than the strategic reallocation of some resources from high- to low-probability items.

One way to contrast these two accounts is to compare the precision of high- and low-probability items from individual conditions within our study. Whereas our original strategic distribution account predicts high- and low-probability items within a condition should be encoded with higher and lower precision, respectively, the leaky filter account predicts no difference. This comparison is most easily accomplished using our Experiment 2 precision data (\bar{J}); specifically, Condition 2, in which one high-probability item was cued (50%), leaving five low-probability items (10% each). Consistent with the strategic distribution account, high-probability items ($M = 190.64$, $SEM = 39.33$) were remembered with statistically significantly greater precision than low-probability items ($M = 37.66$, $SEM = 8.34$) in this condition, $t(23) = -4.87$, $p < .001$. Although these data are most consistent with our original interpretation, more work is necessary to completely disentangle these two accounts and isolate the mechanism by which the strategic control over VWM performance is possible. For instance, you might consider that a filter in the traditional sense does not exist at all and that the way to keep items out of VWM is simply to assign them zero resources. On the other hand, perhaps it is best to think of our effects as being accomplished through a filter alone that can both regulate the likelihood that items are encoded and their precision. Regardless, we believe that the message of the present data is the same: A simple interpretation of the filter account in which feature-based attention gates access to memory is insufficient; feature-based attention also allows the strategic distribution of resources within memory.

Summary and Conclusion

Here we discuss both existing and new data to offer an extended function of the filter account proposed nearly a decade ago. Rather than granting access to VWM resources to individual items in an all-or-none fashion, attention plays a more sophisticated role, determining the relative priority of the items competing for VWM representation and assigning VWM resources to each that are proportional to their value. This broader role for attention suggests that the original framework describing how attention controls VWM performance requires updating: The filter mechanism, described as the “bouncer in the brain,” is, indeed, doing more than just bouncing irrelevant information or granting or denying guests access to a nightclub. We suggest that, in addition to gating access to VWM resources, this bouncer also acts as the hostess of the nightclub, determining VIP status and, consequently, how much of

the club’s valuable resources should be allocated to a given guest based on his or her relative importance.

References

- Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, *15*, 106–111. <http://dx.doi.org/10.1111/j.0963-7214.2004.01502006.x>
- Awh, E., & Vogel, E. K. (2008). The bouncer in the brain. *Nature Neuroscience*, *11*, 5–6.
- Bacon, W. J., & Egeth, H. E. (1997). Goal-directed guidance of attention: Evidence from conjunctive visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *23*, 948–961. <http://dx.doi.org/10.1037/0096-1523.23.4.948>
- Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *The Journal of Neuroscience*, *34*, 3632–3645. <http://dx.doi.org/10.1523/JNEUROSCI.3204-13.2014>
- 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., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854. <http://dx.doi.org/10.1126/science.1158023>
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, *11*, 1–34.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*, 523–547. <http://dx.doi.org/10.1037/0033-295X.97.4.523>
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–114. <http://dx.doi.org/10.1017/S0140525X01003922>
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222. <http://dx.doi.org/10.1146/annurev.ne.18.030195.001205>
- Emrich, S. M., & Busseri, M. A. (2015). Re-evaluating the relationships among filtering activity, unnecessary storage, and visual working memory capacity. *Cognitive, Affective & Behavioral Neuroscience*, *15*, 589–597. <http://dx.doi.org/10.3758/s13415-015-0341-z>
- Emrich, S. M., Lockhart, H. A., & Al-Aidroos, N. (in press). Attention mediates the flexible allocation of visual working memory resources. *Journal of Experimental Psychology: Human Perception and Performance*.
- Farell, B., & Pelli, D. G. (1993). Can we attend to large and small at the same time? *Vision Research*, *33*, 2757–2772. [http://dx.doi.org/10.1016/0042-6989\(93\)90234-N](http://dx.doi.org/10.1016/0042-6989(93)90234-N)
- 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>
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, *15*, 1176–1194. <http://dx.doi.org/10.1162/089892903322598139>
- Gunseli, E., van Moorselaar, D., Meeter, M., & Olivers, C. N. L. (2015a). 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>
- Gunseli, E., van Moorselaar, D., Meeter, M., & Olivers, C. N. L. (2015b). Erratum to: The reliability of retro-cues determines the fate of noncued visual working memory representations. *Psychonomic Bulletin & Review*, *22*, 1474. <http://dx.doi.org/10.3758/s13423-015-0914-4>
- 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>

- Jost, K., & Mayr, U. (2016). Switching between filter settings reduces the efficient utilization of visual working memory. *Cognitive, Affective & Behavioral Neuroscience*, *16*, 207–218. <http://dx.doi.org/10.3758/s13415-015-0380-5>
- Kaptein, N. A., Theeuwes, J., & van der Heijden, A. H. C. (1995). Search for a conjunctively defined target can be selectively limited to a color-defined subset of elements. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1053–1069. <http://dx.doi.org/10.1037/0096-1523.21.5.1053>
- Klyszejko, Z., Rahmati, M., & Curtis, C. E. (2014). Attentional priority determines working memory precision. *Vision Research*, *105*, 70–76. <http://dx.doi.org/10.1016/j.visres.2014.09.002>
- Liesefeld, A. M., Liesefeld, H. R., & Zimmer, H. D. (2014). Intercommunication between prefrontal and posterior brain regions for protecting visual working memory from distractor interference. *Psychological Science*, *25*, 325–333. <http://dx.doi.org/10.1177/0956797613501170>
- 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>
- Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*, 317–322. <http://dx.doi.org/10.1016/j.tins.2006.04.001>
- McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*, 103–107. <http://dx.doi.org/10.1038/nn2024>
- Moore, C. M., & Egeth, H. (1998). How does feature-based attention affect visual processing? *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1296–1310. <http://dx.doi.org/10.1037/0096-1523.24.4.1296>
- Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology*, *4*, 61–64.
- Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*, 327–334.
- Sàenz, M., Buraças, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, *43*, 629–637. [http://dx.doi.org/10.1016/S0042-6989\(02\)00595-3](http://dx.doi.org/10.1016/S0042-6989(02)00595-3)
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are there multiple visual short-term memory stores? *PLoS ONE*, *3*, e1699. <http://dx.doi.org/10.1371/journal.pone.0001699>
- Suchow, J. W., Brady, T. F., Fougine, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, *13*, 1–8. <http://dx.doi.org/10.1167/13.10.9>
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, *428*, 751–754. <http://dx.doi.org/10.1038/nature02466>
- 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., & Ma, W. J. (2014). “Plateau”-related summary statistics are uninformative for comparing working memory models. *Attention, Perception, & Psychophysics*, *76*, 2117–2135. <http://dx.doi.org/10.3758/s13414-013-0618-7>
- 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. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *109*, 8780–8785. <http://dx.doi.org/10.1073/pnas.1117465109>
- Vissers, M. E., van Driel, J., & Slagter, H. A. (2016). Proactive, but not reactive, distractor filtering relies on local modulation of alpha oscillatory activity. *Journal of Cognitive Neuroscience*, *28*, 1964–1979. http://dx.doi.org/10.1162/jocn_a_01017
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*, 748–751. <http://dx.doi.org/10.1038/nature02447>
- Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, *438*, 500–503. <http://dx.doi.org/10.1038/nature04171>
- 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. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, *1*, 202–238. <http://dx.doi.org/10.3758/BF03200774>
- Zhang, W., & Luck, S. J. (2008a). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*, 233–235. <http://dx.doi.org/10.1038/nature06860>
- Zhang, W., & Luck, S. J. (2008b). Feature-based attention modulates feedforward visual processing. *Nature Neuroscience*, *12*, 24–25. <http://dx.doi.org/10.1038/nn.2223>
- Zokaei, N., Gorgoraptis, N., Bahrami, B., Bays, P. M., & Husain, M. (2011). Precision of working memory for visual motion sequences and transparent motion surfaces. *Journal of Vision*, *11*, 2. <http://dx.doi.org/10.1167/11.14.2>

Received November 4, 2016

Revision received March 6, 2017

Accepted March 15, 2017 ■