Addressing properties and limitations of flexible resource allocation in visual short-term memory.

Holly A. Lockhart, Kevin J. MacDonald, Stephen M. Emrich

Acknowledgments:

This project was funded by an NSERC Discovery Grant awarded to SME [435945], and NSERC Postgraduate Scholarship Doctoral Program awarded to HAL. We would like to thank Tom Nelson and Joseph Capozza for their assistance programming Experiments 1 and 2 respectively.

Abstract

Although recent evidence suggests that visual short-term memory is a continuous resource, little is known about how flexibly this resource can be allocated. The current study used a continuous report procedure and attentional prioritization via probabilistic spatial cues to address three unknown properties of a flexible continuous resource. The first experiment measured multiple responses from each trial to assess whether multiple items could simultaneously be prioritized. Second, since past work has shown that participants could prioritize VSTM representations according to two states of attentional priority, we examined whether participants could maintain three levels of priority. Lastly, we examined whether flexible allocation is possible when the prioritization cues are provided after initial encoding (i.e. a retro-cue). The results demonstrated that when participants had to recall multiple items on each trial, there were clear differences in response precision between cued and uncued items; however, two items of the same category were not always stored with equal precision. When three attentional priority levels were provided, participants' precision was no different between high- and medium-priority, but significantly improved over low priority items, suggesting participants did not assign different memory weights to the two higher-priority conditions. When prioritization was performed via retro-cues, participants could re-allocate memory resources, but not when more than one item was to be prioritized, suggesting limitations in the flexible allocation of resources after initial encoding. Together, the results provide evidence of a VSTM resource that is flexibly, but variably, allocated using up to two attentionally guided priority goals, primarily during encoding.

Over the past two decades, visual short-term memory (VSTM), the system used to hold visual information on-line for short periods of time, has received considerable attention (Luck & Vogel, 2013). In particular, the nature of the representational limitations of VSTM has been the subject of much debate. Using a change-detection task, seminal work by Luck and Vogel (1997), suggested that VSTM was a richly detailed but very limited memory system only able to remember approximately 4 items, consistent with a discrete-resource and fixed-capacity model of VSTM. More recently, studies measuring the fidelity of memory with a continuous report procedure (Wilken & Ma, 2004) and other measures (Bays & Husain, 2008) have demonstrated that memory precision changes with the number of remembered items, and further, that more than 3 – 4 items can be remembered, although with much lower representational quality (Bays, Catalao, & Husain, 2009; van den Berg, Shin, Chou, George, & Ma, 2012). These findings support an alternative to the discrete-resource and fixed-capacity models of VSTM: continuous resource models. This family of models proposes that VSTM does not have a fixed capacity, and that resources could theoretically be infinitely divided at the cost of fuzzy, low-resolution memory representations. Although the debate surrounding the precise nature of the limits of VSTM continues (Fougnie, Cormiea, Kanabar, & Alvarez, 2016; Schurgin, Wixted, & Brady, 2018; Zhang & Luck, 2008), the continuous-resource models of VSTM have gained more support in recent years compared to discrete-resource models (van den Berg, Awh, & Ma, 2014).

Many previous studies examining the nature of VSTM representations have focused on how different models explain changes in performance across memory load (e.g. van den Berg, Awh, & Ma, 2014). However, another key prediction of continuous-resource models is that an individual should be able to flexibly allocate resources to objects according to their goals. Recently, a number of studies have shown that resources can indeed be flexibly allocated

according to the attentional priority of each object (Dube, Emrich, & Al-Aidroos, 2017; Emrich, Lockhart, & Al-Aidroos, 2017; Klyszejko, Rahmati, & Curtis, 2014). For example, Emrich et al., (2017) demonstrated that simultaneously-presented spatial cues that varied in their predictive validity impacted the precision of memory recall, where items that were most likely to be probed were reported with greater precision, while guesses and non-target errors were made less frequently in a similarly predictable fashion. Critically, the effect of cue probability on performance was best fit by a power-law function, as predicted by a flexible continuous resource model of VSTM. Similar effects have been observed using feature-based attentional cues (Dube et al., 2017) as well as through manipulating reward (Klyszejko et al., 2014). Thus, these studies suggest that the reported precision of an item in memory is best characterized by the amount of attentional prioritization received by that object, and not by the number of objects in memory, consistent with a flexible and continuous model of VSTM resources.

Although these studies tested some of the predictions a flexible resource model of VSTM, there are many untested questions surrounding the properties and limitations of this flexible resource; the purpose of the present study is to address some of these open questions. In three experiments, we examined: how high- and low-priority objects are maintained concurrently, relative to each other; whether three attentional goals can guide the allocation of memory resources simultaneously; and whether flexible allocation of resources can occur after encoding. Each experiment is a variant on the delay-estimation task with spatial cues indicating relative priority of objects as shown by Emrich et al. (2017) to address the limitations of flexibility on a continuous memory resource.

How are multiple items of different priority stored?

As mentioned above, previous studies (Dube et al., 2017; Emrich et al., 2017) have demonstrated that the precision of VSTM recall is strongly predicted by the attentional

prioritization that a given item receives. Emrich et al. (2017), used spatial cues presented during memory encoding that predicted the likelihood that a given item would be probed (probe-likelihood). For example, in one condition, six items were presented, but one cued item would be probed on 50% of trials. Consistent with the predictions of flexible allocation, the precision of memory recall for the one cue 50%-valid item was nearly identical to that of trials when two items were probed with 100% validity (i.e., when resources should be split equally across two items). Critically, performance on the uncued trials was similarly predicted by the proportion of resource-allocation, fitted by a power-law function; performance for the 5-uncued items was consistent with each item receiving approximately 10% of the total resources (i.e., $1/5^{\text{th}}$ of the remaining 50%). In other words, probe-likelihood predicts the proportion of resources that should be allocated to each item. The two studies in Emrich et al. (2017) demonstrated this effect across conditions ranging from 9% – 100%, and similar effects were observed using feature-based cues in Dube et al. (2017). Thus, these studies demonstrate that (1) memory resources can be flexibly allocated via attention and (2) individual items can receive a small proportion of resources, consistent with a capacity-less, continuous-resource model of VSTM.

It is possible, however, that the objects held in memory in these previous studies were not actually being prioritized according to the different levels of priority, but rather, were being held or excluded from memory probabilistically, such that when averaged across trials it appeared to follow the priority set by the attentional cues. This explanation would potentially support a capacity-limited discrete resource model. Indeed, previous studies using whole report tasks demonstrate that not every item can be reported with the same accuracy, such that when only obtaining one response on each trial, it is likely that meaningful variability is being averaged out across individual items (Adam, Vogel, & Awh, 2017; Cowan, Saults, Elliott, & Moreno, 2002).

To address this possibility, Experiment 1 used the same delay-estimation task with probabilistic cues as in Emrich et al. (2017), but participants reported two items on each trial. On each trial, we obtained measures of precision for two high-priority items, two low-priority items, or one high- and one low-priority item. If participants were indeed allocating resources according to the validity of the cues, then on a given trial, the responses towards multiple items (whether cued or uncued) should reflect that prioritization.

How flexible is the allocation of memory resources?

In both Emrich et al. (2017) and Dube et al. (2017) there were two categories of objects presented: high- and low-priority objects. We suggest this process of sorting items into high- and low-priority is akin to having two active attentional goal-sets. Some theories from the attentional control literature suggest that there can only be one goal-set in memory at a time due to a limited single focus of attention (see Olivers, Peters, Houtkamp, & Roelfsema, 2011 for review). Despite this, using feature based cues, Dube et al. (2017) previously demonstrated that multiple objects could be prioritized under two attentional goals; however, when features are considered targets in a rapid serial presentation task Houtkamp & Roelfsema, (2009) found target identification was impaired. Houtkamp and Roelfsema (2009) suggest these findings are consistent with switching between target templates rather than the templates sharing a limited resource, potentially indicating that there is a cost to maintaining multiple goals in attention. Similarly, when participants could search for either of two different colored squares in a display, reaction time was longer than when searching for both colored squares again suggesting that subjects were switching between target templates (Menneer, Barrett, Phillips, Donnelly, & Cave, 2007).

However, Adamo, Pun, Pratt, and Ferber (2008) showed that two different attentional goal-sets could be active over separate space in the visual field. In further work, it was shown that multiple attentional goals could be maintained for different features such as shape and color

(Adamo, Wozny, Pratt, & Ferber, 2010). As well, Irons, Folk, & Remington, (2012), thoroughly tested whether multiple target colors were set and maintained independently of each other or if they reflected a large band of target colors (i.e. red and blue, vs. red to blue). Across five experiments they consistently showed that cues matching multiple target colors captured attention while irrelevant colored cues did not, even when the distractor color would be linearly between the target colors. These works provide evidence that individuals can maintain multiple goal-sets over space and features, which may be the process underlying attentional prioritization in VSTM as in Emrich et al. (2017), and Dube et al. (2017).

Although this past research has shown that individuals are able to maintain at least two attentional goal-sets and use these goals to guide prioritization, it is still unclear whether the top-down allocation of VSTM resources is limited to only two active goals. Accordingly, Experiment 2 tested whether it is possible to use three levels of priority in memory. Participants were shown two cues indicating three distinct attentional goal-sets on each trial: a high-priority object tested 50% of the time, a medium-priority object tested 25% of the time, and two low-priority objects with no spatial cues each tested 12.5% of the time. If it is possible for participants to allocate memory resources using three attentional goal-sets, then we expect precision to reflect the probe-likelihoods of all three levels of attentional priority, a proxy for resource allocation, as previously shown with two attentional goal-sets.

Are memory resources flexible after encoding?

In studies demonstrating flexible allocation of memory resources, the cues were presented simultaneously with, or ahead of, the memoranda (Dube et al., 2017; Emrich et al., 2017; Taylor, Thomson, Sutton, & Donkin, 2017). However, a number of recent studies have examined prioritization in VSTM using retro-cueing paradigms (i.e., cues presented after encoding). Evidence from the retro-cue literature is mixed on whether participants effectively

use these cues to benefit performance in the same way that pre-cues have demonstrated. In contrast to pre-cues, a single informative retro-cue did not improve the precision of recalled items; although similar to pre-cues, it did improve the probability of an item to be recalled (lower guess rate) (Murray, Nobre, Clark, Cravo, & Stokes, 2013). Other studies have demonstrated that informative retro-cues improved both the speed of recall and the precision of the memory report for auditory pitch, suggesting retro-cues may benefit auditory memory performance in a similar manner as pre-cues (Lim, Wöstmann, Geweke, & Obleser, 2018). Gunseli, van Moorselaar, Meeter, & Olivers, (2015) manipulated the validity of retro-cues, demonstrating that the precision and recall probability differ between cues with 50% and 80% validity (Gunseli et al., 2015). However, these investigations did not address whether attention is acting to flexibly reallocate memory resources by systematically examining the effect of single probabilistic cues and multiple retro-cues.

Accordingly, Experiment 3 examined whether it was possible to flexibly allocate memory resources with retro-cues or whether flexible resource allocation is limited to pre-cues and simultaneous-cues. To test this question, probe-likelihood information was indicated by a retro-cue, and consequently tested whether memory resource allocation can be optimized after encoding. In this design all items should be initially encoded as equally likely to be probed but must be updated with the prioritization information for optimal performance. If memory resources are only flexible during encoding then participants will not be able to update the memory representations with the priority information from the retro-cue; however, if memory resources remain flexible and participants can update their representations with the priority information, this would provide evidence that memory resources are flexible after encoding.

Summary

Although previous studies have provided evidence that VSTM resources can be flexibly allocated according to attentional prioritization cues, the properties and limits of this flexible resource allocation remain unclear. In the current study, we examined three questions regarding flexible resource allocation: whether multiple items on a given trial are allocated VSTM resources in proportion to cue validity of probe-likelihood, whether three attentional goal-sets can be used to direct flexible allocation, and lastly, whether memory resources are flexible after encoding. All of these studies used variations on the delay-estimation task and seek to expand the understanding of an assumed flexible resource, rather than dissociate between a discrete and flexible model of VSTM resources.

Methods

General procedure

Participants were recruited from Brock University using an online research pool or through posters around campus. Each experiment had independent samples. Participants were offered a choice of research credit for courses or paid at a rate of \$10/hour of participation. All participants had normal or corrected to normal vision and were tested for normal color vision. All procedures were approved by the Brock Research Ethics Board.

All stimuli were presented using PsychoPy (Peirce, 2007) on a 20" LCD display at a distance of approximately 57 cm. All stimuli were colored squares of 1° of visual angle evenly spaced around the central fixation. Color stimuli were randomly selected from 360 unique colors from a continuous color wheel made using CIE L*a*b* color space with a minimum distance of 30°. This color wheel was presented during the color judgment phase of each experiment and had a radius of 7° visual angle. The probe location was indicated by a bolded square. All responses were made using the mouse by clicking on the color that best matched their memory. The color would appear inside the probed square as participants made their selection. Response location

was recorded and the difference between the reported color and the target color was calculated as the raw error. All participants completed a session of practice trials before the experimental session and were given the opportunity to ask question or clarify the instructions. During each experiment participants were given self-paced breaks every 25 trials. Participants also completed a change detection task to assess working memory capacity with Pashler's formula prior to the experimental task (Pashler, 1988). This task had set sizes 2, 4, and 6, 50% change trials, and whole probe test presentation. After each experiment each participant filled out the Media Multitasking Inventory (Ophir, 2016); this data is secondary to the main purpose of the current study and as such are not presented here.

Experiment 1

Participants:

Twenty participants, 4 males, with a mean age of 20.50 (range 18-26). Participants were given the option to listen to non-lyrical music during the task on their own devices. Four participants chose to listen to music.

Procedure and Stimuli:

Six colored squares were presented around a central fixation cross with two spatial cues pointing at two of the locations for 500ms followed by a delay of 900ms. The test display stayed on the screen until both responses were made. On each trial, participants were always probed on two locations: 50% of the time both of the probe locations were the two cued locations; 25% of the time participants were probed on one cued location and one uncued location (mixed-report). 25% of the time two uncued locations were probed. The participants were told these probabilities at the start of the session. This meant that each cued item would be probed 31.25% of the time, and each uncued item would be probed 9.375% of the time participants chose which square to report first by clicking inside the square and choosing their response with the mouse. The color

inside the chosen probed location changed to reflect the mouse position around the colorwheel.

Participants completed 400 intermixed trials: 200 trials of both cued locations probed, and 100 trials each of both uncued location probed, and cued-uncued probes. See Figure 1A.

Figure 1. Example trials from each experiment. All experiments were delay estimation tasks consisting of a study phase, delay, and report. Responses were made with the mouse. The color inside the probed square changed to reflect the mouse position around the wheel. **A.** Example trial of experiment 1. Six colored items were presented on grey background for 500ms. Spatial line cues indicated two of the six items were more likely to be probed after the delay. Each cued item had a 31.25% chance to be probes, each uncued item had a 9.33% chance to be probed. After a 900ms delay two color responses were made. **B.** Example trial of experiment 2. Four colored stimuli were presented on a grey background with one black and one white spatial cue. The cue color meaning was counterbalanced across participants. The white cue may have indicated the high priority items (50% probe-likelihood), and the black cue indicated the medium priority items (25% probe-likelihood). The uncued items each had a 12.5% chance to be probed. Participants made one response at the end of a 900ms delay. **C.** Example trial of Experiment 3. One to four spatial cues were presented in the location of the item 800ms after the item was displayed for 300ms (retro-cue) followed by a second delay period of 800ms. Conditions were given in randomly ordered blocks.

Experiment 2

Participants:

Twenty-one participants, 7 males, with a mean age of 20.38 (range 17-29). Participants

were given the option to listen to non-lyrical music during the task on their own devices. Four participants chose to listen to music.

Procedure and Stimulus:

Four colored squares were presented around a central fixation with two spatial cues, one black and one white. The color stimuli were presented for 500ms followed by a delay of 900ms. The test display stayed on the screen until response. The color probabilities were counter-balanced across subjects: for half of subjects the black cued item was probed on 50% of trials and the white cued item was probed 25% of trials, the two uncued items were probed the remaining 25% of trials (12.5% each). If the white cued item was probed the bold box around the probe location was white, otherwise it was black. Participants completed a total of 400 trials were presented all conditions intermixed. See Figure 1B.

Experiment 3

Participants:

Twenty-six participants, 2 males, with a mean age of 18.87 (range 18-25). Six participants only completed three of four of the conditions due to a computer error. One subject was removed for having error on the single cue condition above 3 standard deviations from the mean.

Procedure and Stimulus:

Four colored squares were presented around a central fixation for 500ms. Following a delay of 800ms, one, two, or four spatial retro-cues appeared for 300ms and were followed by a delay of 800ms before the test probe was displayed. These retro-cues were empty squares in the cued locations. In the valid conditions, one, two, or four squares appeared with 100% validity; in a critical condition one cue appeared with only 50% validity. Conditions were blocked, and participants were told both the number and validity of the cues. Participants completed 500 trials:

100 trials of each of the three 100% valid cueing condition and 200 trials of the 50% valid cuing condition. See Figure 1C.

General Analysis

Responses were measured as the angular distance between the target and reported color value in degrees. The standard deviation of the error was used as a measure of (inverse) precision; when not explicitly testing between different models (as is the case here), model free summary statistics such as standard deviation of the circular error have been demonstrated to be the most reliable measure for testing differences between conditions (Ma, 2018). Tests between conditions were performed using repeated-measures ANOVA. Post-hoc tests used Bayesian pairwise t-test to quantify the evidence for the null-hypothesis (BF_{01}) or the alternative hypothesis (BF_{10}). A Bayes factor larger than 1 indicates the degree of evidence in the given direction, values larger than 3 are considered substantial evidence (Wetzels et al., 2011). Repeated measures ANOVA and Bayesian pairwise t-tests were conducted using JASP (JASP Team, 2018).

A complementary mixed-effect model with random effects of participant (random intercept) and fixed effects for the test variables (fixed-slope) was applied using the lme4 package in R (Bates, Maechler, Bolker, & Walker, 2015). This technique is similar to repeated-measures ANOVA as this removes the variance related to participant identity while testing the effects of the variables of interest. These tests model all trials the participants completed rather than modeling the summary statistics for each participant. Unless otherwise stated the models are tested against an intercept model (random effect of participant only) for goodness-of-fit. These tests can verify the main effects and interactions of the repeated-measures ANOVA with the added benefit of better accounting for each subject's behavior across all trials since all trials are used in the model. Secondly, the mixed-effect models are able to test linear (and log-linear)

relationships across factors and thus are able to suggest if the precision of responses is best fit by a log-linear fit predicted by the continuous resource model. This log-linear test was only performed on Experiment 3 where there are enough data points to test the fit.

Experiment 1: How are multiple items of different priority stored?

Results

In the first experiment, participants were prompted to report the colors of two out of the six sample items presented. Importantly, the two selected items for report were controlled such that the probed items had 1) both been cued during encoding (with a cuing probability of 31.25 % per item), 2) both uncued (9.3 % per item), or 3) probed on one item from each category (mixed trials). Responses were self-ordered, and thus were recorded by order of response as well as whether targets were cued or uncued. The mixed report condition was split into whether participants reported the cued item first (cued-first) or the uncued item first (uncued-first). There was no instruction to report the better remembered item first but it was anticipated that participants would report the better remembered item first (Adam et al., 2017). The cued item was reported first on approximately 75% of total trials but this varied by participant. The range of trials where participants reported the cued item first was between 48 - 99% of the mixed report trials (96 - 198 of 200 trials).

Figure 2. A. Distribution of all participants error in degrees by response condition. Boxplots overlaid on all data points. *B.* Graph of the standard deviation of circular error by response condition and response order. Individual subjects shown in transparent lines, subjects with less than 4 trials in the Mixed, Uncued First responses were omitted from graphing. Standard error bars shown.

The distributions of the data by condition and order can be seen in Figure 2A (bar graphs can be seen in Figure 2B). A 4x2 Repeat-measures ANOVA was used to test the condition by order interaction to understand how two items are reported under differing cue conditions. There was a significant effect of condition, $F(3, 54) = 48.15, p < 0.001$, and of order, $F(1, 18) = 35.88, p < 0.001$, confirming that participants remembered the cued items better than the uncued items. The main effect of order suggests that, in general, there is a difference between the first and second responses; this effect remained significant when excluding mixed-response trials from the analysis. Finally, there was a significant condition by order interaction, $F(3, 54) = 9.73, p < 0.001$. To investigate the interaction, we looked at the simple main effects of condition by order (Figure 2B). This showed that there was a significant effect of order when reporting both cued items, $F(1, 18) = 54.65, p < 0.001$, when reporting the cued item first in the mixed-response condition, $F(1, 18) = 228.30, p < 0.001$, and when reporting both uncued items, $F(1, 18) = 7.04, p = 0.016$, but not when reporting the uncued item first in the mixed-response condition, $F(1, 18) = 0.463, p = 0.50$. The effect of order was stronger in the both-cued condition than in the both-uncued condition, suggesting that the effect of order may not entirely driven by interference from the first report and that there are other factors affecting the difference in precision between responses.

To further examine the precision of the cued items relative to each other, we compared the first response in the both-cued condition against the cued response in the mixed-report-cued-first condition. Both of these responses were the first ones made in a trial and as such should have been reported with a similar amount of time since the study phase and should not be affected by interference from a previous response. The only difference between these responses is whether participants were assigned one of the two cued items to report (mixed response trials)

or whether they could choose which of the two cued items to report first (both-cued response trials). There was substantial evidence suggesting that the first reported item from the both-cued trials were reported with greater precision than the cued item in the mixed response condition, $BF_{10} = 5.839$ ($Cued_1 = 44.40$ [95CI: 37.29, 51.51]; $Mixed-Cued_1 = 53.51$ [95CI: 47.14, 59.87]). This suggests that there are differences in the precision of the memories between cued items.

We similarly examined whether there were differences in the precision of the uncued items by comparing whether the precision for the uncued item (reported second) in the mixed-response cued-first condition was the same as in the both-uncued response condition (second response). Both of these responses have the similar time decay and previous report interference, and only differed in whether the previous response was a cued or uncued item. There was moderate evidence for the null hypothesis, suggesting that the second responses from the mixed-response condition (uncued item) were similar to the second responses from the both-uncued response condition $BF_{01} = 3.205$ ($Mixed-Uncued_2 = 101.73$ [95CI: 97.80, 105.66]; $Uncued_2 = 102.93$ [95CI: 99.34, 106.52]). This analysis suggests that the priority of the previously reported item did not affect the precision of an uncued item reported second in a trial.

A complementary mixed-effect model was run on the absolute error of each response and input in a hierarchical fashion in the order intercept (random effects of participant only), response order effect, and probe-likelihood. This three-step process tests if the effect of probe-likelihood is above the effect of response order. Response order was a significant predictor of absolute error $\chi^2(1) = 548.46, p < 0.0001$. Probe-likelihood was a significant predictor of error above response order, $\chi^2(1) = 3312.12, p < 0.0001$. This confirms that probe-likelihood affects the precision of the output when controlling for order effects and between subject variance, consistent with past

studies (Dube et al., 2017; Emrich et al., 2017). As well, the benefit to cued items is maintained when controlling for response order suggesting that two items were prioritized as expected.

Discussion

The aim of Experiment 1 was to examine whether evidence for flexible prioritization in VSTM was evident when subjects had to provide a response for multiple items on each trial. Accordingly, two items were reported from an array of six colored items, with two levels of priority. When objects are assigned priority through probabilistic spatial cues it is apparent in the data that participants do hold the high- and low-priority items with different precision. Specifically, items with a high probability of being probed are reported with lower error than lower priority items. This finding supports a model of a flexible resource that can be assigned to objects through top-down control (Dube et al., 2017; Emrich et al., 2017; Klyszejko et al., 2014).

The mixed-report, uncued-first trials stand out from the expected pattern of results: cued-item should be reported first because they would have a better memory since it was high-priority (Adam et al., 2017). There are fewer trials of this report order overall, and 10 of the participants had fewer than 40 trials in this order. One possible explanation for this unexpected response order is that it reflects trials in which the cued item was not prioritized over the uncued items, or potentially forgotten entirely. Previous studies have reported that guessing may occur due to lapses in attention (Fougnie, Suchow, & Alvarez, 2012), and it is known that memory performance is affected by location binding errors (Bays et al., 2009; Emrich & Ferber, 2012), as well as differences in strategy (Linke, Vicente-Grabovetsky, Mitchell, & Cusack, 2011). Consequently, the low proportion of these trials may suggest this behavior may be due to lapses of attention, misremembered cued locations, inadequate prioritizing, or a suboptimal strategy. Thus, in general, while participants were able to use the cues to flexibly allocate attentional

resources, the order of multiple responses may reveal information about imperfect flexible resource allocation on some trials (Adam et al., 2017).

Our data also suggest that there is variability between items stored in memory. It was expected that there would be some loss of precision on the second report due to the increased time before the response, as well as possible interference from the previous report (Cowan et al., 2002). Although there is a significant difference between the two responses in the two-cued-report condition, it is unknown how much of this difference could be the result of time and interference. However, since the cued report in the mixed-report-cued-first trials was not as precise as the first-cued report in the two-cued-report trials it was apparent that the memories of both cued items are not exactly equal at the time of the first report. This suggests there are factors other than time (decay) and report interference leading to variation in precision across the cued items.

What might cause the difference in precision between the two cued items? Since participants chose which item to respond to first, they likely chose the item that was better remembered. Previous studies have demonstrated that participants will typically report the most accurately remembered items first (Adam et al., 2017). Thus, on two-cued-report trials, subjects likely reported the most precisely stored item of the two cued items first. However, on mixed-report trials, subjects could only report one of two cued items; thus, even though the cued item was remembered better than the uncued item, this may not have been the more precisely represented of the cued items. This discrepancy in precision between equally-prioritized items appears to support a variable precision continuous resource model (Fougnie et al., 2012; van den Berg et al., 2012). In other words, although the multi-response procedure confirmed that multiple

items were indeed prioritized, it did reveal that not all items are prioritized equally, and that there is some variability in the precision between prioritized items.

There was less variability between the uncued responses. Given that these reports are averaged over the sample of 4 low-priority objects, it may be that variability is lost in the average across these items, and a full-report method would better capture variability in the low-priority objects. Moreover, there may be less room for variability while maintaining such low-priority objects in memory without losing a meaningful representation or reaching floor effects.

Interestingly, a recent study analyzed the data of Emrich et al. (2017) and found evidence for guessing for some very low-priority items (van den Berg & Ma, 2018). Thus, it is possible that variability for low priority items could include a significant number of trials in which those items were not encoded or were removed from memory. Exactly how variable these low priority items are, including the factors that contribute to guessing, should be the subject of future research.

Experiment 2: How flexible is the allocation of memory resources?

Results

Experiment 1 presented evidence, consistent with previous findings (e.g., Emrich et al., 2017), that a set of items can be assigned two different levels of attentional priority in VSTM. The goal of Experiment 2 was to examine whether more than two priority levels (or goal sets) can be established. This was achieved by having three levels of probe-likelihood (low: 12.5% each, medium: 25%, and high: 50%) (See Figure 1B). Across the three cue conditions, the standard deviation of circular error was analyzed using a repeated-measures ANOVA (Figure 3). There was a significant effect of cue condition $F(2, 38) = 26.09, p < 0.001$, showing participants differentially prioritized items given the cue information. Planned comparisons showed that the uncued (low-priority) item was reported with significantly worse precision than the medium-priority item ($t(19) = 4.94, p < 0.001$) as well as the high-priority item ($t(19) = 5.69, p < 0.001$),

demonstrating that participants were able to prioritize cued items over uncued items. However, there was no significant difference between the medium-priority and high-priority items ($t(19) = 0.12, p = 0.91$). Bayesian pairwise t-tests found moderate evidence in favour of the null hypothesis that there is no difference in response precision between medium- and high-priority conditions $B_{01} = 4.276$ (Medium = 52.80 [95CI: 42.86, 62.73], High = 53.02 [95CI: 44.20, 61.84]). No difference in the response precision suggests that there was no difference in the memory representations for the two higher levels of priority.

Figure 3. A. Distribution of all subjects' error in degrees by probe-likelihood conditions. Boxplots overlaid on all data points. *B.* Graph of the standard deviation of circular error by probe-likelihood. Individual subjects shown in transparent lines. Standard error bars shown.

A mixed-effects model confirmed these results. Task condition was a significant predictor of response error, $\chi^2(2) = 529.32, p < 0.0001$, supporting previous results that participants could prioritize items given cues. As in the ANOVA, there was no difference between the high-priority and medium-priority conditions, $B = 0.31 (0.58), t(7996) = 0.053, p = 0.958$, suggesting these items were not prioritized according to their known probe-likelihood. There was a significant difference between the average of both cued conditions and the uncued condition $B = 17.34 (0.74), t(7980) = 23.30, p < 0.001$, this demonstrates prioritization and supports a flexible allocation of a continuous memory resource.

Although these results so far demonstrate that participants could not maintain three levels of attentional priority, it may have been the case that participants initially treated these levels differently, but over time they stopped differentiating the conditions (e.g., due to fatigue); alternatively, participants could have learned to use the priority information as the task went on. Either explanation would be evident if there was a change in precision from the start of the experiment to the end that differed between the cued conditions, with either response precision becoming more similar (fatigue) or becoming differentiated (learning). To test this, an exploratory analysis was performed to determine if there was a change in the response precision over time, and if that change differed between the cueing conditions. A mixed-effect model was run on absolute error for each trial with the interaction of trial and probe-likelihood using random slopes, which allows the effect of trial to vary between participants. This analysis showed there was a significant main effect of trial. Over 100 trials, participants reduced their absolute precision by 1.90° , meaning participants on average were 7.60° less precise at the end of the 400 trials $B = 1.90 (0.42), t(7961) = 4.537, p < 0.001$; however, this effect was not different between the cuing conditions $B = 0.79 (0.63), t(7961) = 1.24, p = 0.215$, see Figure 4. It

is evident in Figure 4 that the response precision for higher priority items are not different throughout the entire duration of the task while the low priority items are reported with worse precision overall, $B = 17.28 (3.25)$, $t(20) = 5.31$, $p < 0.001$, and that absolute error increases from the start to end of the task. Thus, participants did not learn to use the priority cues better over time nor did they get worse at using the cues over time: they simply responded with worse precision over time.

Figure 4. Absolute error in degrees of all trials (conditions shown in colored shapes), averaged over all participants. 95% confidence intervals shown.

Discussion

How flexible is the allocation of memory resources? This experiment used three levels of attentional priority to determine if participants could allocate memory resources using a number

of attentional goals. The data support the evidence that memory resources can be flexibly allocated, as there was increased precision for the high-priority cued items compared to the low-priority uncued items. Contrary to the idea that resources are infinitely flexible, participants appeared to treat the two higher priority objects the same; that is, there was no difference in the precision of responses between the high- and medium-priority items. Closer analysis revealed that the lack of distinction in precision between the medium- and high-priority items was consistent across the experiment, and that participants neither learned to allocate to three levels, nor became fatigued of allocating to three levels of priority (see Figure 4). This finding suggests that it was not possible for participants to allocate memory resources using three attentional goals and instead treated all cued items as high priority. Taken together, the results suggest that prioritization of VSTM resources via top-down attention is limited to only two levels, consistent with past work that demonstrating that at least two separate attentional control sets are possible (Adamo et al., 2008, 2010; Liu, Becker, & Jigo, 2013; Liu & Jigo, 2017).

The absence of a significant difference between the two cued conditions is particularly surprising in light of a recent study by Yoo, Klyszejko, Curtis, & Ma, (2018). In that study, participants were cued using colored wedges with proportional radial size to indicate the relative priority of each quadrant, with cuing probabilities of 0.6, 0.3, 0.1, and 0. Their results showed three distinctive levels of response error which were proportional to the validity of the distinct cues. It is possible that our task design was not adequate for revealing more than two levels of attentional priority. For example, the two types of spatial cues may not have been sufficiently distinct to readily establish different goal sets during encoding. Given that flexible encoding can be mediated by either spatial attention (Emrich et al., 2017), feature-based attention (Dube et al., 2017) or reward (Klyszejko et al., 2014), it's possible that multiple levels of attention could be

established using different forms of attentional prioritization. Alternatively, it is possible that the two levels of high- and moderate-priority were not different enough to motivate flexible allocation to three attentional goals sets (which may be mentally costly). Although the priority levels chosen for this study (50% and 25%) were showed different precision when resources were allocated using only two attentional goal sets (Emrich et al., 2017). Consequently, while further studies should address the differences between the results observed by Yoo et al. (2018) and those observed here, our results reveal that there are potentially limits on the ability to flexibly prioritize information in VSTM via attention.

Interestingly, we also observed that memory performance decreased over the course of the experiment. A recent finding showed the likelihood of low-precision responses increased with time on task when participants did not nap before testing, but when participants did nap before testing, performance was maintained throughout the task (MacDonald, Lockhart, Storace, Emrich, & Cote, 2018). This finding provides evidence that it is an effect of fatigue leading to increased variability in VSTM precision. However, this is in contrast to a previous study, which examined session-by-session analysis of mixture model precision, and which failed to show significant variability in precision across blocks of 160 trials (Rademaker, Tredway, & Tong, 2012), although there are a number of methodological differences between this study and the one by Rademaker and colleagues which could account for these differences. Nevertheless, the finding that memory precision changes over the course of the experiment is consistent with the general class of models which suggest that there is variability in VSTM precision.

Experiment 3: Are memory resources flexible after encoding?

Results

In the previous two experiments, we confirmed that even when reporting multiple items, participants can flexibly allocate VSTM resources according to probabilistic cues, but that

participants cannot flexibly prioritize three discrete levels of priority. Experiment 3 tested the question of whether flexible (i.e., proportional) prioritization can occur after initial encoding. This was achieved through retro-cues presented during the delay period rather than during the sample display (Figure 1C). There were four task conditions resulting in five types of responses that could be compared. Three conditions had 100% valid cues, but varied between one, two and four items. Since there were always four items presented, the 100% valid, four-cue trials represent a “neutral” cue. Comparisons between these conditions will address the efficiency of resource allocation when using a retro-cue, as varying the number of 100% valid retro-cues represents a typical load manipulation. The final condition was a single 50% valid retro-cue, which resulted in responses for both cued and uncued items (on different trials). If near-optimal flexible allocation is possible (i.e., if resource allocation in a retro-cue tracks with cue probability), this cued item should be allocated approximately 50% of memory resources, similar to each of the items in the 100% valid, two-cue condition, as we have observed in previous studies (Dube et al., 2017; Emrich et al., 2017). Thus, if retro-cue probability efficiently allocates resources in a similar fashion to simultaneous cues (Emrich et al., 2017), we would predict differences between the 100% valid, one-cue and 50% valid, one-cue conditions, while the 100% valid, two-cues and 50% valid, one-cue conditions should demonstrate nearly identical precision. By contrast, the uncued items should each be allocated approximately 16.67% of the memory resources ($0.50/3$) and should therefore be less precise than the neutral cue condition (25%). In addition, by comparing across all cued and uncued conditions, we can test whether the precision of responses follow an expected log-linear pattern according to the probe-likelihood of each item, as predicted by our past results and continuous resource models of VSTM.

To test if flexible prioritization occurred, a repeated-measure ANOVA confirmed that there were differences in the precision of responses across conditions, $F(4, 72) = 25.34, p < 0.001$ (Figure 5). As described above, we followed up this test with paired comparisons to determine whether prioritization reflected the likelihood of each item being probed (i.e., the number and validity of the cues). The first comparison of interest was whether the 100% valid, one-cue condition and the 50% valid, cued condition differed. These conditions are visually identical and only differ in how the cue is treated by the participant. Consistent with past retro-cueing studies (Gunseli et al., 2015), the 100% valid one-cue condition was significantly more precise than the 50% valid cued condition, $t(18) = 3.44, p = 0.018$; $BF_{10} = 14.698$, $M = -0.961$, $[95CI = -1.659, -0.269]$. Moreover, the uncued items in the 50% valid condition showed evidence of being less precise than the neutral, 100% valid four-cue condition, $t(18) = 2.57, p = 0.063$; $BF_{10} = 3.019$, $M = 0.692$, $[95CI = 0.066, 1.343]$, in line with the predictions of flexible prioritization. This suggests that the uncued items were deprioritized compared to the 100% valid, four-cue (neutral) condition, in line with the predictions of flexible prioritization. Together these tests reveal that participants can control the degree to which they maintain a single item at the cost of the remaining items, as expected by flexible allocation of a continuous memory resource.

To further establish the effectiveness of flexible prioritization, we examined the effect prioritizing multiple items using a retro cue, as the 100% valid conditions (one, two, and four) reflect a typical load manipulation. The one-cue condition was significantly more precise than the four-cue condition, $t = 6.29, p < 0.001$; $BF_{10} = 6.194 \times 10^5$, $M = -2.203$, $[95CI = -2.638, -1.508]$, however, the two-cue and four-cue conditions did not differ in precision ($t = 0.103, p =$

0.92; $BF_{01} = 4.10$, $M = -0.124$, $95CI = -0.638, 0.364$). This finding suggests that unlike simultaneous cues, multiple items cannot be prioritized using retro-cues.

An important feature of flexible allocation is that items that have a similar probe likelihood should be allocated a similar proportion of resources. This characteristic was demonstrated through rigorous comparison with simultaneous cues in two studies done by Emrich et al. (2017). For example, a singled item cued with 50% validity showed identical precision to trials in which two items were probed with 100% validity, as in both cases each item had a 50% chance of being probed. By contrast, in the current study, the 50% valid, cued condition was statistically different from the 100% valid, two-cue condition, $t = 2.88$, $p < 0.05$; $BF_{10} = 5.190$, $M = 0.789$, $[95CI = 0.149, 1.456]$. This suggests that flexible allocation is not as efficiently allocated with a retro-cue as has been observed with simultaneous cues.

We also examined whether the precision of responses followed the expected log-linear pattern predicted by the probe likelihood of each item, consistent with previous studies that used simultaneous spatial and feature-based simultaneous-cues (Dube et al., 2017; Emrich et al., 2017). Two mixed-effect models compared the fit of a linear and a log-linear trend by modeling the absolute error with the probe-likelihood (linear fit) and the log of absolute error and the log of probe-likelihood (log-linear fit). Both the linear and log-linear models are a significant improvement over the intercept model ($\chi^2(4) = 474.47$, $p < 0.0001$; and $\chi^2(1) = 334.04$, $p < 0.0001$); however, the log-linear relationship is a better fit evident by a much smaller AIC value ($37752 < 117773$). Thus, although these results are inconclusive to determine if flexible prioritization using a retro-cue is done optimally, there is sufficient evidence that some amount of flexible allocation did occur across the conditions.

Discussion

The primary question of Experiment 3 was whether memory resources could be flexibly allocated after encoding. Three aspects of the data suggest that a retro-cue can be used to flexibly allocate resources: first, the uncued items showed some evidence of being less precise than the neutral 100% valid, four-cue condition; second, a single cue with 50% validity showed lower precision than a single cue with 100% validity, suggesting that a single item can be prioritized strategically; and third, predicting memory precision by probe likelihood according to a log-linear function was a good predictor of the data, consistent with previous simultaneous and feature-based cue results (Dube et al., 2017; Emrich et al., 2017).

Despite this evidence, the data do not support a resource that is fully flexible after encoding. Notably, we failed to observe two effects that have been observed using other cue types. First, we did not observe flexibility over multiple loads, such that the precision of responses for two cues was no different than that of neutral (four cue) trials. This finding replicates previous results finding there is limited ability to prioritize more than one item using a retro-cue (Makovski & Jiang, 2007). Second, the precision of responses did not closely follow the likelihood that an item would be probed. Most notably, unlike in previous studies (e.g., Emrich et al., 2017), we observed a significant difference in the precision of responses between 100% valid, two-cue items and the 50% validly-cued items. Thus, although in both conditions each item had a 50% likelihood of being probed on each trial, participants did not optimally allocate resources according to this information. However, since the precision of the 100% valid two-cue condition did not differ from that of the 100% valid four-cue condition, it is unclear if the two items were effectively prioritized at all. Consequently, this calls into question whether the comparison between the validly-cued, 50% one-item condition and the 100% valid, two-cue condition reflects the intended comparison. Nevertheless, the finding that the precision of

responses in the 100% valid, two-cue condition did not differ from that of the 100% valid, four-cue (“neutral”) condition is consistent with the interpretation that individuals were ineffective at prioritizing more than a single item.

Why did we not observe flexibility in the prioritization of multiple items in VSTM using retro cues? Bays and Taylor (2018) argued that when an item is retro-cued there is no way to gain any perceptual information; in a population coding model, this results in attention driving the neural tuning functions to increase in amplitude (gain) without changing the tuning width (memory precision). Thus, if signal can be increased in a continuous manner, then it’s possible that the 100% valid cue condition receives a larger increase in gain than the 50% valid cue condition, resulting in a proportional change in error. However, increasing gain at multiple cued locations without a change in the tuning width may not be sufficient to prevent interference or drift between those representations. This may suggest that the mechanism of flexible allocation in VSTM observed in previous studies (Emrich et al., 2017; Dube et al., 2017) is achieved by decreasing the tuning width of high priority items during encoding by acquiring more perceptual information when it is available. Further work would need to directly investigate the neural tuning functions of multiple retro-cues as well as simultaneous probabilistic cues to assess the direct effect on signal gain and tuning width.

The obvious alternative explanation that a multi-item benefit is not seen is that memory resources are simply not flexible after initial encoding and cannot be re-allocated; however, our current data, as well as that of Gunseli et al., (2015), show that flexible re-allocation is possible for a single item. Similarly, Shimi and colleagues found that a 50% valid retro-cue was less effective than a 100% valid cue (Shimi, Nobre, Astle, & Scerif, 2014). Thus, it is apparent that there is at least some flexibility in the allocation of resources using retro-cues. Another

explanation may be that utilizing retro-cues is very effortful, which may interfere with maintaining multiple items in memory. Lim et al., (2018) reported that a retro-cuing benefit to auditory precision was more reliable in high-capacity subjects, suggesting that processing a retro-cue itself was cognitively demanding. Thus, it may be that multiple items do not experience a retro-cue benefit because it is too cognitively demanding.

Recent evidence from a study by Dube, Lumsden, and Al-Aidroos, (2018) found that a single probabilistic retro-cue valid at 70% could improve accuracy of change detection to a similar extent as the 100% valid “deterministic” cued item; critically, however, they showed evidence that the item cued with 70% validity did not bias attention in a subsequent visual search task, whereas the 100% validly cued item did. This evidence may suggest that retro-cueing multiple items is ineffective because multiple (retro-cued) items are not maintained in the focus of attention. Thus, while the present data are consistent with the interpretation that memory resources can only be flexibly allocated to a single item post-encoding, future studies may be needed to delineate the precise mechanisms and boundary conditions of these limits.

General Discussion

Using probabilistic cueing paradigms, previous research has demonstrated that VSTM is a continuous resource that can be flexibly allocated according to the attentional priority of a given item (Dube et al., 2017; Emrich et al., 2017; Klyszejko et al., 2014; Yoo et al., 2018). While this past work has characterized memory resource allocation as flexible and continuous, it did not address the limitations of this flexibility. The current study examined three questions related to the properties and limitations of a flexible and continuous VSTM resource. Experiment 1 confirmed that high-priority items are reported with better precision than low priority items, even when examining multiple responses within the same trial. In addition, items of equal low-

priority are reported with similar precision while high priority items have more variability in precision that is not entirely explained by the order of response or time since encoding. This suggests that even when subjects are prioritizing a small subset of items within the memory array, items with equal priority are not necessarily encoded and maintained with equal resolution. This finding is consistent with variable precision model variants of VSTM (Fougnie et al., 2012; van den Berg et al., 2012). Indeed, when compared side-by-side, variable precision models have typically outperformed equal precision models in VSTM (Dube et al., 2017; Emrich et al., 2017; Klyszejko et al., 2014; van den Berg et al., 2014). Thus, this study provides further support for these models.

One alternative explanation for these results is that participants only encoded or maintained a subset of the six items in memory. Indeed, van den Berg and Ma (2018) analyzed data from Emrich et al. (2017) and observed that some subjects appeared to be guessing at very low cueing probabilities. Moreover, Adam et al. (2017) used a whole-report procedure and observed that when the order of responses was randomized, a number of early responses looked like guesses, consistent with their participant's report that they were guessing on those responses by indicating low confidence. Given that in the current study no confidence measures were acquired, it is difficult to assess on which trials participants might have been guessing. However, if participants only held three or four items in memory, we would expect that when participants had to report two uncued items in any order, their first response should be much more precise than the second one, given that the first reported item would have had a higher likelihood of having been one of the three or four items stored in memory (Adam et al., 2017). Thus, the evidence that there was no difference between the two uncued items suggests that all uncued items may have been encoded with similar precision, rather than participants only encoding one

or two of those items. Whether this was true for all subjects remains a question to be further explored (van Den Berg & Ma, 2018).

We showed in Experiment 2 that there is a limitation on the number of attentional goal-sets that can be engaged to allocate memory resources. Specifically, there was no evidence that participants were able to allocate memory resources according to the three different priority levels tested, as the high- and medium-priority conditions had the same level of response precision. This finding held when looking at the trend of response precision over the full experiment, indicating that there was no evidence participants either learned to differentiate the conditions or stopped using the cues over the course of the experiment. It may be that only two goals can be engaged via top-down processes; previous work has demonstrated at least two attentional templates could be maintained but had not decisively tested if participants could effectively use more than two attentional templates (Adamo et al., 2008, 2010; Irons et al., 2012). Alternatively, it may be that only one goal can act on the spotlight of attention but participants may switch through multiple goals serially (Houtkamp & Roelfsema, 2009). Our data suggest there is a limitation on the number of attentional goals that can be engaged for top-down memory allocation. It is not evident in this study, however, if the limitation is specific to spatial cues. It is also unknown whether other factors, like increased motivation, could encourage three levels of priority to be maintained by participants. Recent work by Yoo et al., (2018) has shown multiple priority levels being effectively engaged within a trial, suggesting that there may be conditions that allow for more flexible allocation of memory resources than shown in the current study. Regardless, VSTM resources, while flexible, may not be infinitely flexible due to specific limitations in top-down control mechanisms. Further studies could determine the conditions in which it is possible to find evidence of three concurrently held attentional goals, as

well as investigate the source of this limitation (e.g. bottleneck in switching through multiple goals).

Lastly, Experiment 3 showed evidence of flexible allocation after encoding, but it was not as effective as when cues are presented at the time of encoding. No stepwise load effect was found on the response precision between the 100% valid cued trials at one, two, and four cues as expected from previous work using simultaneous cues to control memory load (Dube et al., 2017; Emrich et al., 2017; Klyszejko et al., 2014). This finding is consistent with past work in by Makovski & Jiang, (2007) which showed no load effect using retro-cues in change detection task. Although other work has appeared to demonstrate an effect of load (Lepsien, Griffin, Devlin, & Nobre, 2005), the changes in load occurred during encoding, not with retro-cuing. The lack of a load effect in the current data suggests there is a limited ability to use retro-cues to improve precision as compared to simultaneous or pre-cues. However, since there was some evidence for proportional prioritization, despite the limitation to prioritize multiple items, it is likely that resources can be reallocated to some extent after encoding.

The results from Experiment 3 are potentially supported by a recent population coding model which suggests that memory signal can only increase in gain (and is thus protected from noise) and not decrease in tuning width (increase in memory fidelity) following a retro-cue (Bays & Taylor, 2018). A review of the retro-cue literature found support for some of the popular hypothesis of the retro-cue benefit such as 1) increased “strength” and 2) protection of the cued item with 3) removal of the uncued items from memory (Souza & Oberauer, 2016). Each of these supported hypotheses can be accounted for in the neural model put forward by Bays and Taylor (2018). Similarly, the principle of flexible reallocation may also unify these concepts. That is, the current study would suggest that each of these hypotheses are the effect of

reallocating memory resources away from the uncued items to the cued item to provide “strength” and protection from perceptual interference.

Although these experiments begin to address the limitations in a flexible continuous resource, there are several limitations. Most notably, although Ma (2018) has argued that raw error is the best measure to use when comparing across manipulations, it is possible that this measure does not capture all aspects of VSTM performance. Indeed, models have been devised to capture several theoretical aspects of memory performance, including binding errors (Bays et al., 2009; Swan & Wyble, 2014), categorical responses (Hardman, Vergauwe, & Ricker, 2017), resource rationality (van den Berg & Ma, 2018) and perceptual similarity (Schurgin et al., 2018). While our analysis may not capture all nuances of resource allocation, most or all of those theorized responses are captured in the measure of raw error. Moreover, although the model that best characterizes memory performance may change over time, the model that best accounts for behavior clearly needs to capture both the flexible control of resources (Dube et al., 2017; Emrich et al., 2017; Klyszejko et al., 2014), as well as the limitations therein.

Finally, though it is clear that memory resources are flexibly allocated, it remains unclear what the neural mechanisms are that mediate this process. It has been well established that the inter-parietal sulcus (IPS) is involved in the maintenance and storage of items in visual memory (Todd & Marois, 2004); it has also been suggested that frontal structures control access to VSTM (Edin et al., 2009; McNab & Klingberg, 2008). However looking at structures alone can only explain part of the mechanism. Emrich and Busseri (2015), tested the implied path from frontal regions controlling access (filtering) to storage activity in the parietal cortex that would predict working memory capacity using the data from McNab & Klingberg (2008). When they tested a path model, the direct path of from filtering activity to working memory capacity was a stronger

predictor than the indirect path of frontal activity controlling access to memory storage (unnecessary storage) predicting VSTM performance. Thus, although previous work had focused on the role of frontal activity as a filter for items the authors suggested that the direct path is more consistent with general executive attention having a central role in VSTM. It is unknown how flexible allocation would be reflected in this activity, but we suggest it is accomplished through those networks that mediate top-down attention and executive control. This proposal is consistent with the findings of Edin et al. (2009) finding functional connectivity between frontal activity in the dorsolateral prefrontal cortex and the IPS, as well as the direct path found by Emrich and Busseri (2015) suggesting attention related activity plays a central role in VSTM performance.

Conclusion

Together, these results add to the body of literature demonstrating that VSTM resources can be flexibly allocated via attentional manipulations. Importantly, they begin to form a better picture of the properties and limitations of that flexibility. Specifically, these data reveal that there are limitations in the efficiency of allocation after initial encoding, and that it is ineffectual at more than two attentional goals. Moreover, they also demonstrate that while resource allocation is flexible, it remains variable. While flexibility is more typically a characteristic of continuous resource models, some of these results may also be consistent with discrete resource models, as long as the model allows for the degree of flexibility demonstrated. Whether or not VSTM resources are continuous or discrete, it is apparent that future models and mechanisms of VSTM should consider this flexible allocation.

References

- Adam, K. C. S., Vogel, E. K., & Awh, E. (2017). Clear evidence for item limits in visual working memory. *Cognitive Psychology*, 97, 79–97.
<https://doi.org/10.1016/j.cogpsych.2017.07.001>
- Adamo, M., Wozny, S., Pratt, J., & Ferber, S. (2010). Parallel, independent attentional control settings for colors and shapes. *Attention, Perception & Psychophysics*, 72(7), 1730–1735.
<https://doi.org/10.3758/APP.72.7.1730>
- Adamo, M., Pun, C., Pratt, J., & Ferber, S. (2008). Your divided attention, please! The maintenance of multiple attentional control sets over distinct regions in space. *Cognition*, 107(1), 295–303. <https://doi.org/10.1016/j.cognition.2007.07.003>
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48.
<https://doi.org/10.18637/jss.v067.i01>
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10), 7–7.
<https://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. <https://doi.org/10.1126/science.1158023>
- Bays, P. M., & Taylor, R. (2018). A neural model of retrospective attention in visual working memory. *Cognitive Psychology*, 100, 43–52.
<https://doi.org/10.1016/j.cogpsych.2017.12.001>

- Cowan, N., Saults, J. S., Elliott, E. M., & Moreno, M. V. (2002). Deconfounding Serial Recall. *Journal of Memory and Language*, 46(1), 153–177.
<https://doi.org/10.1006/jmla.2001.2805>
- Dube, B., Emrich, S. M., & Al-Aidroos, N. (2017). More than a filter: Feature based attention regulates the distribution of visual working memory resources. *Journal of Experimental Psychology: Human Perception and Performance*, 43(10), 1843–1854.
- Dube, B., Lumsden, A., & Al-Aidroos, N. (2018). Probabilistic retro-cues do not determine state in visual working memory. *Psychonomic Bulletin & Review*.
<https://doi.org/10.3758/s13423-018-1533-7>
- Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegnér, J., & Compte, A. (2009). Mechanism for top-down control of working memory capacity. *Proceedings of the National Academy of Sciences*, 106(16), 6802–6807. <https://doi.org/10.1073/pnas.0901894106>
- 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(3), 589–597. <https://doi.org/10.3758/s13415-015-0341-z>
- Emrich, S. M., & Ferber, S. (2012). Competition increases binding errors in visual working memory. *Journal of Vision*, 12(4), 12–12. <https://doi.org/10.1167/12.4.12>
- Emrich, S. M., Lockhart, H. A., & Al-Aidroos, N. (2017). Attention mediates the flexible allocation of visual working memory resources. *Journal of Experimental Psychology: Human Perception and Performance*, 43(7), 1454–1465.
<https://doi.org/10.1037/xhp0000398>

- Fougnie, D., Cormiea, S. M., Kanabar, A., & Alvarez, G. A. (2016). Strategic Trade-Offs Between Quantity and Quality in Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*. <https://doi.org/10.1037/xhp0000211>
- Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, 3, 1229. <https://doi.org/10.1038/ncomms2237>
- 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(5), 1334–1341. <https://doi.org/10.3758/s13423-014-0796-x>
- Hardman, K. O., Vergauwe, E., & Ricker, T. J. (2017). Categorical working memory representations are used in delayed estimation of continuous colors. *Journal of Experimental Psychology: Human Perception and Performance*, 43(1), 30–54. <https://doi.org/10.1037/xhp0000290>
- Houtkamp, R., & Roelfsema, P. R. (2009). Matching of visual input to only one item at any one time. *Psychological Research Psychologische Forschung*, 73(3), 317–326. <https://doi.org/10.1007/s00426-008-0157-3>
- Irons, J. L., Folk, C. L., & Remington, R. W. (2012). All set! Evidence of simultaneous attentional control settings for multiple target colors. *Journal of Experimental Psychology: Human Perception and Performance*, 38(3), 758–775. <https://doi.org/10.1037/a0026578>
- JASP Team. (2018). JASP (Version 0.9). Retrieved from <http://jasp-stats.org>

- Klyszejko, Z., Rahmati, M., & Curtis, C. E. (2014). Attentional priority determines working memory precision. *Vision Research*, *105*, 70–76.
<https://doi.org/10.1016/j.visres.2014.09.002>
- Lepsien, J., Griffin, I. C., Devlin, J. T., & Nobre, A. C. (2005). Directing spatial attention in mental representations: Interactions between attentional orienting and working-memory load. *NeuroImage*, *26*(3), 733–743. <https://doi.org/10.1016/j.neuroimage.2005.02.026>
- Lim, S.-J., Wöstmann, M., Geweke, F., & Obleser, J. (2018). The benefit of attention-to-memory depends on the interplay of memory capacity and memory load. *Frontiers in Psychology*, *9*. <https://doi.org/10.3389/fpsyg.2018.00184>
- Linke, A. C., Vicente-Grabovetsky, A., Mitchell, D. J., & Cusack, R. (2011). Encoding strategy accounts for individual differences in change detection measures of VSTM. *Neuropsychologia*, *49*(6), 1476–1486.
<https://doi.org/10.1016/j.neuropsychologia.2010.11.034>
- Liu, T., Becker, M. W., & Jigo, M. (2013). Limited featured-based attention to multiple features. *Vision Research*, *85*, 36–44. <https://doi.org/10.1016/j.visres.2012.09.001>
- Liu, T., & Jigo, M. (2017). Limits in feature-based attention to multiple colors. *Attention, Perception, & Psychophysics*, *79*(8), 2327–2337.
<https://doi.org/10.3758/s13414-017-1390-x>
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. <https://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*(8), 391–400.
<https://doi.org/10.1016/j.tics.2013.06.006>

Ma, W. J. (2018). Problematic usage of the Zhang and Luck mixture model.

<https://doi.org/10.1101/268961>

MacDonald, K. J., Lockhart, H. A., Storace, A. C., Emrich, S. M., & Cote, K. A. (2018). A daytime nap enhances visual working memory performance and alters event-related delay activity. *Cognitive, Affective, & Behavioral Neuroscience*, 18(6), 1105–1120.

<https://doi.org/10.3758/s13415-018-0625-1>

Makovski, T., & Jiang, Y. V. (2007). Distributing versus focusing attention in visual short-term memory. *Psychonomic Bulletin & Review*, 14(6), 1072–1078.

<https://doi.org/10.3758/BF03193093>

Menneer, T., Barrett, D. J. K., Phillips, L., Donnelly, N., & Cave, K. R. (2007). Costs in searching for two targets: dividing search across target types could improve airport security screening. *Applied Cognitive Psychology*, 21(7), 915–932.

<https://doi.org/10.1002/acp.1305>

McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, 11(1), 103–107. <https://doi.org/10.1038/nn2024>

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(4),

550–556. <https://doi.org/10.1177/0956797612457782>

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. <https://doi.org/10.1016/j.tics.2011.05.004>

- Ophir, E., Nass, C., & Wagner, A. D. (2009). Cognitive control in media multitaskers. *Proceedings of the National Academy of Sciences*, 106(37), 15583–15587.
<https://doi.org/10.1073/pnas.0903620106>
- Pashler, H. (1988). Familiarity and visual change detection. *Attention, Perception, & Psychophysics*, 44(4), 369–378. <https://doi.org/10.3758/BF03210419>
- Peirce, J. W. (2008). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2. <https://doi.org/10.3389/neuro.11.010.2008>
- Rademaker, R. L., Tredway, C. H., & Tong, F. (2012). Introspective judgments predict the precision and likelihood of successful maintenance of visual working memory. *Journal of Vision*, 12(13), 21–21. <https://doi.org/10.1167/12.13.21>
- Shimi, A., Nobre, A. C., Astle, D., & Scerif, G. (2014). Orienting Attention Within Visual Short-Term Memory: Development and Mechanisms. *Child Development*, 85(2), 578–592.
<https://doi.org/10.1111/cdev.12150>
- Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2018). Psychological Scaling Reveals a Single Parameter Framework For Visual Working Memory. <https://doi.org/10.1101/325472>
- Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, 78(7), 1839–1860.
<https://doi.org/10.3758/s13414-016-1108-5>
- 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(7), 2136–2157. <https://doi.org/10.3758/s13414-014-0633-3>

- Taylor, R., Thomson, H., Sutton, D., & Donkin, C. (2017). Does working memory have a single capacity limit? *Journal of Memory and Language*, 93, 67–81.
<https://doi.org/10.1016/j.jml.2016.09.004>
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, 428, 751–753. <https://doi.org/10.1038/nature02466>
- van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, 121(1), 124–149. <https://doi.org/10.1037/a0035234>
- van den Berg, R., & Ma, W. J. (2018). A resource-rational theory of set size effects in human visual working memory. *ELife*, 7. <https://doi.org/10.7554/eLife.34963>
- 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*, 109(22), 8780–8785.
<https://doi.org/10.1073/pnas.1117465109>
- Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E.-J. (2011). Statistical Evidence in Experimental Psychology: An Empirical Comparison Using 855 *t* Tests. *Perspectives on Psychological Science*, 6(3), 291–298.
<https://doi.org/10.1177/1745691611406923>
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(12), 11. <https://doi.org/10.1167/4.12.11>
- Yoo, A. H., Klyszejko, Z., Curtis, C. E., & Ma, W. J. (2018). Strategic allocation of working memory resource. *Scientific Reports*, 8(1). <https://doi.org/10.1038/s41598-018-34282-1>
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233–235. <https://doi.org/10.1038/nature06860>