Published in final edited form as:

J Exp Psychol Hum Percept Perform. 2010 December ; 36(6): 1419–1429. doi:10.1037/a0019324.

Statistical Learning Induces Discrete Shifts in the Allocation of Working Memory Resources

Akina Umemoto, Miranda Scolari, Edward K. Vogel, and Edward Awh

University of Oregon

Abstract

Observers can voluntarily select which items are encoded into working memory, and the efficiency of this process strongly predicts memory capacity. Nevertheless, the present work suggests that voluntary intentions do not exclusively determine what is encoded into this online workspace. Observers indicated whether any items from a briefly stored sample display had changed. Unbeknownst to observers, these changes were most likely to occur in a specific quadrant of the display (the *dominant quadrant*). Across 84 subjects and 5 groups of observers, change detection accuracy was significantly higher for items in the dominant quadrant, suggesting that memory encoding was biased towards the dominant quadrant. Only 9 of the 84 subjects were able to correctly specify the dominant quadrant when asked whether any location was more likely to contain the changed item, but more sensitive forced-choice procedures did reveal above-chance discrimination of the dominant quadrant. Nevertheless, because forced choice performance was unrelated to the size of the bias and no observer reported a biased encoding strategy, the bias was unlikely to depend on voluntary encoding strategies. The encoding bias was not due to a reduction in the response threshold for indicating changes in the dominant quadrant (Experiment 2). Finally, separate measures of the number and resolution of the representations in memory suggested that encoding was biased in a discrete slot-based fashion (Experiment 3). That is, although items in the dominant quadrant were more likely to be encoded into memory, mnemonic resolution for the favored items was not affected.

Working memory allows the maintenance of information in an easily updated and accessible state. However, this system is subject to strong capacity limits, such that the average observer can maintain only about 3-4 items at a time (Luck and Vogel, 1997; Pashler, 1988; Sperling, 1960). These severe limits place priority on understanding the processes that guide the encoding of information into this online mental workspace. It has long been known that observers can exert voluntary control over what is encoded into working memory, and previous research has examined basic characteristics of this process such as its time-course (e.g., Reeves and Sperling, 1986; Vogel, Woodman and Luck, 2006) and the efficiency with which relevant items are selected for encoding (McNab and Klingberg, 2008; Vogel, McCollough and Machizawa; 2005). Indeed, a growing body of research has suggested that the voluntary selection of relevant over irrelevant information is intertwined with working memory capacity (e.g., Awh, Vogel and Oh, 2006; Engle, 2002; Kane, Bleckley, Conway and Engle, 2001; Vogel and Awh, 2008). For example, Vogel et al. (2005) examined observers' ability to selectively encode only the relevant items in a display into visual working memory by providing trial-by-trial cues indicating the to-be-remembered items. During some trials, irrelevant items of a different color were also presented in the sample display. The key finding was that individual differences in capacity estimates strongly predicted the ability to block irrelevant objects from entering working memory, as assessed

by a neural measure of the number of distractor items that were encoded into memory (Vogel, McCollough and Machizawa, 2005; see also McNab and Klingberg, 2008). Thus, robust links have been identified between the voluntary gating of information into working memory and the total number of items that can be effectively represented. Indeed, Vogel et al. (2005) proposed that the key source of individual variability in such memory tasks may be the observers' ability to keep irrelevant stimuli from occupying the limited space in working memory, rather than the total amount of space that is available.

Given the clear links between selection and working memory capacity, there is motivation to document the full range of processes that determine which items are encoded into working memory. The present work shows that manipulating the probability of target locations also exerts an influence on which items will be encoded into working memory, even though observers were uninformed about the probability manipulation and only rarely were able to specify the probable target location when directly queried. This works follows from a number of past studies that have shown that the deployment of attention is influenced by acquiring, consciously or unconsciously, the probability with which targets appeared in certain locations, leading to faster and more accurate processing of the probable target locations than the improbable target locations (Shaw & Shaw, 1977; Shaw, 1978; Miller, 1988; Hoffmann & Kunde, 1999; Geng & Behrmann, 2005). For example, Shaw and Shaw (1977) presented a target letter briefly in one of eight possible target locations on an imaginary circle, and manipulated the probability with which target appeared in different locations. Subjects were informed about the location probability in advance, and the authors showed that subjects could use that information to identify a target more accurately (Shaw and Shaw, 1977) and quickly (Shaw, 1978) at the probable location, suggesting that visual attention was biased towards the probable target locations.

The aforementioned studies employed tasks in which the key limiting factor was the speed or quality with which single targets could be visually encoded. By contrast, the present research examined whether manipulating the probability of target locations would also bias encoding when all items in the sample display are potential targets and performance was limited by the storage of information in visual working memory. Although encoding into working memory is traditionally thought to depend on participants' voluntary selection of the stored items, the present work tested whether statistical regularities in prior experience would also influence the deployment of limited storage resources in working memory independent of the voluntary strategies of the observer. To anticipate the results, we found that when targets were more probable in one quadrant of the display, observers were biased to encode items from that quadrant (the *dominant quadrant*) into visual working memory, though observers failed to report any explicit encoding bias. Our work follows from other studies that have examined how statistical regularities can influence performance in memory-limited tasks. For example, Olson, Jiang and Moore (2005; see also Beck, Angelone, Levin, Peterson and Varakin, 2008) measured observers' ability to report the position of a missing stimulus from an array of stimuli held in spatial working memory. Olson et al. (2005) found that when the same item was consistently removed from repeated spatial arrays, observers were better at reporting the missing location than with novel arrays. Importantly, repeated arrays only generated this benefit when the same item was consistently removed from the display, suggesting that likely targets had been prioritized during encoding into working memory while overall capacity remained constant for repeated arrays. Finally, although Olson et al. found that observers could recognize the repeated displays with above-chance accuracy, they also found that recognition performance did not correlate with the magnitude of the benefit for repeated displays. Thus, their data suggest that visual learning of repeated sample arrays biased encoding into spatial working memory without affecting observers' explicit encoding strategies.

In the present work, a key extension of the previous findings is that we offer new evidence regarding how resource allocation in visual working memory is biased when the probable location of targets is manipulated. In particular, we examined whether the bias influenced the probability of encoding into working memory, or the precision with which the encoded items were represented. This question is relevant to the ongoing debate regarding the nature of capacity limits in visual working memory. According to flexible resource models, capacity in working memory is determined by competition for a central pool of resources that can be flexibly allocated across the stored items, such that items of greater complexity receive a larger proportion of the overall resources (e.g., Alvarez and Cavanagh, 2004; Bays and Hussein, 2008; Eng, Cheng and Jiang, 2005). Thus, when items acquire a higher priority for storage in visual working memory, flexible resource models suggest that the optimal encoding strategy would be to allocate a greater proportion of resources to the prioritized items, thereby leading to better mnemonic precision for those items. By contrast, so-called discrete slot models (e.g., Barton, Ester and Awh, 2009; Rouder, Morey, Cowan, Zwilling, Morey & Pratte, 2008; Zhang and Luck, 2008) suggest that resources in visual working memory are allocated in a quantized fashion such that observers can choose which items are stored in working memory (i.e., which items are assigned a “slot”) but without the possibility of asymmetrically dividing resources between the actively represented items. Zhang and Luck (2008) provided evidence in favor of discrete resource allocation by showing that cueing observers to place high priority on a specific item in a sample array increased the probability that the item was encoded into memory, but had no influence on the precision with which that item was represented. Likewise, Barton et al. (2009) found that the precision with which a given item was represented was unaffected by the complexity of the other items in the display, inconsistent with the claim that more complex items are assigned a higher proportion of resources from a shared central pool. The present work provides converging evidence relevant to this debate by demonstrating that increasing the probability of targets in a given region influences only the probability that items from that region will be encoded, without affecting the resolution or clarity with which those items are represented. Thus, encoding biases elicited by statistical learning appear to bias a discrete or quantized process of resource allocation during encoding into visual working memory.

Experiment 1a, 1b and 1c

Experiments 1a, 1b, and 1c tested whether change detection performance would be significantly better at locations where changes were more probable. In addition, post-experiment questionnaires examined whether subjects were aware of this probability manipulation. The key difference between the three experiments was the strength of the probability manipulation. The dominant quadrant, where changes were more likely, contained 80%, 60%, and 50% of the changes in Experiments 1a, 1b, and 1c, respectively.

Method

Subjects

Three different groups of students from the University of Oregon participated in Experiments 1a (N=12), 1b (N=24), and 1c (N=12) for course credit. The procedure lasted one hour. All subjects reported normal or corrected-to-normal visual acuity.

Stimuli

The stimuli were colored squares (white, black, red, blue, green, yellow, and purple). Each object subtended approximately $1.5^{\circ} \times 1.5^{\circ}$ of visual angle. Four or eight randomly selected objects were presented in different locations within a rectangular region 17° tall and 20° wide. The horizontal and vertical coordinates of each object's position were randomly

selected (with continuous variation in both dimensions), with the constraint that equal numbers of objects must occupy each quadrant, and no object could appear within 1.7° of another object. Subjects were seated 50 cm from the computer screen. For each subject, a single quadrant was selected as the *dominant quadrant* (counterbalanced across subjects). In Experiment 1a, 80% of the changes occurred in the dominant quadrant and the other 20% of the changes were randomly divided among the remaining three quadrants (nondominant quadrants). Therefore, approximately 6.7% of the changes occurred in each of the nondominant quadrants. The proportion of trials in the dominant quadrant changed to 60% in Experiment 1b, and 50% in Experiment 1c. For Experiments 1b and 1c, each non-dominant quadrant contained an equal number of changes such that 13.3% and 16.7% of changes fell in each non-dominant quadrant in Experiments 1b and 1c, respectively. Subjects were not told about the probability manipulation.

Procedure

Each trial (illustrated in Figure 1) began with the onset of a light grey region that encompassed all the possible stimulus positions. 1541 ms later, a sample array of either 4 (half of all trials) or 8 (half of all trials) colored squares (randomly selected with replacement, with the constraint that no color appeared more than twice) was presented for 100 ms, followed by a 900 ms delay period. Finally, a test array appeared that was either identical to the sample array (with probability .5), or contained one item whose color had changed (with probability .5). The test array remained visible until subjects pressed the “z” key to indicate that the test array was the same as the sample array, or the “/” key to indicate that it was different. Subjects were instructed to place the highest priority on accuracy, yet to respond quickly without sacrificing accuracy. Subjects completed 6 blocks of 80 trials in Experiment 1a, 8 blocks of 60 trials in 1b, and 10 blocks of 48 trials in 1c. At the end of the experiment, subjects were given a questionnaire to assess their awareness of the probability manipulation. They were asked whether they had noticed a specific location that was more likely to contain the changed items. A paper diagram of a plain square (with no demarcation of quadrants) depicting the square region where stimuli were presented was drawn below the question. If subjects answered yes, they checked the corresponding location on the diagram. Since the quadrants were not specified, they marked a location of their choice anywhere within the square diagram. If more than one location was indicated, they were asked to choose only one¹. Subsequently, all subjects were also asked if they had noticed anything else to provide them opportunities to describe their experience with our task. After the completion of the questionnaire, the experimenter sat with each participant individually and went over the questionnaire to make sure the questions were understood correctly and to verify the intended answers if there was an initial misunderstanding.

Results and Discussion

Experiment 1a The 4 item condition produced ceiling effects (mean accuracy = 93%), so all analyses were performed on the 8 item condition. Memory capacity was estimated at 4.74 on average, based on the formula developed by Pashler (1988) and refined by Cowan (2000)². Accuracy in the 8 item condition in the dominant quadrant (74%) was significantly higher than accuracy in the nondominant quadrants (57%), $t(11)=2.58$, $p=0.03$, resulting in an average bias effect of 17% (see Figure 2). A large bias effect (18%) was present from the 1st block of the trials. A 2 by 6 analysis of variance with probability manipulation (dominant

¹This happened in 5 out of 84 subjects.

²This formula requires the inclusion of accuracy during no-change trials. Thus, in all experiments except for Experiment 2 (in which no-change trials were specifically associated with either the dominant or non-dominant quadrants), the same pool of no-change trials was used to derive k in both the dominant and non-dominant quadrants. Capacity (k) is estimated by the formula: $k = \text{set size} * (\text{hit rate} - \text{false alarm rate})$.

and nondominant) and block (6 blocks) as factors to examine the timecourse of the bias effect revealed a significant main effect of probability manipulation, $F(1,11)=7.03$, $p=.02$, Partial $\eta^2=.39$, no effect of block, and no significant interaction, suggesting that a similar bias was present from the first block onward. Indeed, we failed to see evidence of such an interaction throughout all of the studies reported here. This null result may seem surprising at first glance. Given that observers were not exposed to the probability manipulation prior to the first block, one might presume that the effect should grow larger across blocks as the evidence of location bias accumulates. A likely explanation, nevertheless, is that statistical learning occurred at a relatively rapid rate in these experiments, such that the effect had reached its maximal amplitude during the first block of trials. This hypothesis is consistent with past studies that observed measurable effects of visual statistical learning after only several trials of exposure (e.g., Turk-Browne, School, Chun, and Johnson, 2008). The post-experiment questionnaire revealed that observers only identified the dominant quadrant correctly in 2 of 12 cases (17%). Five subjects of the subjects that did not point out the dominant quadrant reported that they did not notice any biases in target position during the experiment, and the rest selected either one of the nondominant quadrants, or a broad region (e.g. “the edges of the display”) that was not centered on the dominant quadrant.

Experiment 1b

The average capacity estimate was 4.1 objects. Similar to the findings in Experiment 1a, Experiment 1b revealed significantly higher accuracy in the dominant quadrant (73%) than in the nondominant quadrants (62%), $t(23)=4.82$, $p<0.01$ (see Figure 2), with an average bias effect of 11%. A bias effect started to appear in the 2nd block (8%), and ended with 13% in the last block. A 2 (probability manipulation) by 8 (block) analysis of variance showed a significant main effect of both the probability manipulation, $F(1,23)=23.68$, $p<.01$, Partial $\eta^2=.51$, and block, $F(7,161)=3.90$, $p<.01$, Partial $\eta^2=.15$. The effect of block appears to reflect a rise in accuracy after the first block. Again, there was no significant interaction between probability manipulation and block, suggesting that the bias effect was relatively stable across blocks. Only 1 out of the 24 subjects correctly identified the dominant quadrant in the post-experiment questionnaire, with the other subjects either reporting not to have noticed any probability manipulation ($n=9$), selecting incorrect quadrants, or a broad region not centered on the dominant quadrant.

Experiment 1c

The average capacity estimate was 3.5 objects. Similar to Experiments 1a and 1b, accuracy in the dominant quadrant (65%) was significantly higher than in the nondominant quadrants (53%), $t(11)=2.39$, $p=0.04$ (see Figure 2). The average bias effect was 12%. A 2 (probability manipulation) by 10 (block) analysis of variance revealed a significant main effect of probability manipulation, $F(1,11)=5.72$, $p=.04$, Partial $\eta^2=.34$, no effect of block and no interaction between probability manipulation and block. Two out of 12 subjects correctly identified the dominant quadrant (numerically below chance levels), with the rest reporting not to have noticed any probability manipulation ($n=4$), selecting incorrect quadrants, or a broad region that was not centered on the dominant quadrant.

Finally, we analyzed the combined data from Experiments 1a, 1b and 1c. This two by two ANOVA with experiment (1a, 1b, or 1c), and probability manipulation (dominant or non-dominant) as between- and within-factors, respectively, revealed a main effect of probability manipulation, $F(1,45) = 29.09$, $p<.01$, Partial $\eta^2=.39$, but not of experiment, $F(2,45) = 2.19$, $p=.12$, Partial $\eta^2=.09$. Unexpectedly, we did not observe an interaction between condition and experiment ($F[2,45] = .41$, $p = .67$, Partial $\eta^2=.02$), suggesting that the effect was not sensitive to the parametric variation in the strength of the probability manipulation across these procedures. Admittedly, we were surprised at the failure to observe reliable differences

in the size of the encoding bias across relatively large changes in the strength of the probability manipulation. Although there was a moderate numerical difference between the bias effect in Experiment 1a and Experiments 1b and 1c, this was not statistically reliable. We cannot offer any conclusive evidence for why this null result was obtained, but our working hypothesis is that measurement noise prevented a sensitive assessment of the predicted relationship between the strength of the bias and the strength of the probability manipulation. The factors that may have increased measurement noise included relatively low numbers of trials in which the critical stimulus appeared in the non-dominant quadrants in Experiment 1a, and relatively low numbers of subjects in Experiment 1a and 1c. Nevertheless, the three experimental groups did serve of the purpose showing that the core effect was robust across three independent groups of observers. Finally, the analysis of reaction time from all three experiments ruled out a speed-accuracy tradeoff explanation of encoding bias; reaction times were significantly faster in the dominant quadrant ($M=1065$, $SD=233$) than in the nondominant quadrants ($M=1125$, $SD=270$), $F(1,45)=8.84$, $p=.01$, Partial $\eta^2=.16$. This main effect of condition did not interact with experiment. Additionally, there was a main effect of the experiment factor, such that reaction time decreased as the strength of the probability manipulation increased, $F(2,45)=6.17$, $p<.01$.

To summarize the results of Experiments 1a, b, and c, three separate groups of subjects replicated the finding that change detection performance was enhanced in regions that were more likely to contain the changed target. The post-experiment questionnaire suggested that subjects were not aware of the probability manipulation, in that only a very small proportion of subjects (5 of 48) identified the dominant quadrant when asked whether any specific regions of the screen were more likely to contain targets. Moreover, subjects who correctly identified the dominant quadrant did not show a larger bias effect (6%) than those who did not ($n=43$; 14%). Nevertheless, we cannot conclude that subjects had no explicit knowledge of the dominant quadrant, because a more sensitive forced choice procedure in Experiments 2 and 3 revealed above-chance identification of the dominant quadrant. That fact notwithstanding, the failure of observers to indicate the dominant quadrant when initially queried, combined with the disconnect between successful identification and the size of the bias effect argues that the observed encoding bias was not a product of voluntary encoding strategies.

Experiment 2

The results so far clearly show higher accuracy for change detection in the dominant quadrant. This finding suggests that knowledge of the likely target position biased memory encoding towards the items in the dominant quadrant. However, because the probability manipulation could only be examined during “change” trials, we were not able to directly compare the rate of false alarms (i.e., trials in which observers incorrectly indicated a change) in the dominant and nondominant conditions of Experiments 1a, 1b and 1c. Thus, the previous experiments cannot rule out the possibility of a shift in decision criterion for indicating changes in the dominant and nondominant quadrants. For example, if observers reduced their response threshold for indicating a change when the suspected change was in the dominant quadrant, then change detection rates could improve in the dominant quadrant even if the amount of information retained from that quadrant was not affected. Such a reduction in response thresholds for indicating a change in the dominant quadrant would also increase false alarms in the dominant quadrant, but this could not be tested in Experiments 1, because the dominant and nondominant quadrants were not distinguished during the “no change” trials when false alarms were possible. In Experiment 2, we addressed this limitation by cueing a specific item in the test display and asking observers to indicate whether it was the same color as the item that had occupied that position in the sample array. This provided separate measurements of hit and false alarms rates in the

dominant and nondominant quadrants, enabling an unambiguous test of whether sensitivity to changes was higher in the dominant quadrant.

Methods

One methodological concern was that the use of a single probe would make the probability manipulation much easier to notice. To avoid this problem, observers were first run through the same procedure used in Experiment 1b, followed by a neutral procedure in which observers were cued to indicate the status (changed or not) of a single item in the test array. Because previous pilot data had established that these bias effects linger for a substantial period even after the bias in target frequencies is eliminated, we expected the probability effect from the first part of the procedure to linger throughout the second neutral phase of the study. As the data will show, the predicted encoding biases were robust throughout the neutral phase of the experiment.

Subjects

Twenty students from the University of Oregon participated in a one hour experiment for course credit ($n=10$) or monetary compensation ($n=10$). All subjects reported normal or corrected-to-normal visual acuity.

Stimulus

All aspects of the stimuli were equivalent to those in Experiment 1b with the exception of the post-cue box used during the neutral condition. In the neutral condition, the test array included a brown box cue ($2.1^\circ \times 2.5^\circ$) that encircled the critical item in the test display, and this cue appeared 750ms after the onset of each test array. The cued item was the only item that could have changed, and it appeared equally often in each of the four quadrants.

Procedures

Four blocks (60 trials each) of the biased condition preceded four blocks (64 trials each) of the neutral condition. In the neutral condition, subjects were instructed to indicate whether the cued item had changed or not. The cued item was different from the sample array with probability 0.5. A post-experiment questionnaire was administered after completing the entire session to assess subjects' awareness of the probability manipulation. To provide a more sensitive test of whether subjects had any explicit knowledge of the dominant quadrant, the post-experiment questionnaire was modified so that those who reported not to have noticed the probability manipulation or did not specify a quadrant when first queried (e.g., chose a broad region), were also required to choose one of the four quadrants where they thought changes were most likely to occur by marking a quadrant on a square diagram.

Results and Discussion

Accuracy in the 4-item trials was 90% on average, so that ceiling effects precluded the emergence of the predicted encoding bias. Thus, the analysis reported here focuses on the 8-item trials. Replicating the previous experiments, the biased condition from the initial phase of the experiment revealed better change detection accuracy for the dominant quadrant (65%) than for the nondominant quadrants (52%), resulting in a bias effect of 13%, $t(19)=3.85$, $p<.01$. A bias effect was already present during the 1st block (11%). A 2 (probability manipulation) by 4 (block) analysis of variance to examine the timecourse of the bias revealed a significant main effect of the probability manipulation, $F(1,19)=15.3$, $p<.01$, Partial $\eta^2=.45$, no main effect of block and no interaction between the probability manipulation and block, suggesting that the bias was similar in size throughout the four blocks of the biased condition. The size of the encoding bias was the same whether

participants received payment or course-credit (both 13%). Therefore, the analysis below did not distinguish these two groups. Next, we analyzed the data from the subsequent neutral condition, with each quadrant classified as “dominant” or “non-dominant” based on the assignments in the initial biased procedure. The bias effect in the neutral phase of the study (11%) was sustained at the same level as during the initial biased procedure. Accuracy in the change trials was 56% in the dominant quadrant and 45% in the nondominant quadrant, $t(19)=2.64$, $p=.02$. Accuracy in no-change trials was equivalent in the dominant (92%) and non-dominant (94%) trials, indicating that false alarms were no more likely in the dominant quadrant than in the nondominant quadrant. Because separate hit and false alarm (HR and FA, respectively) rates were available for the dominant and non-dominant quadrants, we examined hit rate (HR) minus false alarm rate (FA) to obtain a bias-free measure of performance. This estimate revealed a significantly higher performance for the dominant quadrant ($M=.48$, $SD=.22$) than for the nondominant ($M=.39$, $SD=.14$) quadrants, $t(19)=2.23$, $p=.04$, indicating greater sensitivity to changes in the dominant quadrant (see Figure 3). A 2 (probability manipulation) by 4 (block) analysis of variance using this dependent measure revealed a significant main effect of probability manipulation, $F(1,19)=4.87$, $p=.04$, Partial $\eta^2=.20$. There was no significant main effect of block ($F(3,57)=1.06$, $p=.37$), or interaction between probability manipulation and block ($F(3,57)=.27$, $p=.85$).

Finally, we examined the possibility of differential response bias in the two conditions by calculating B'' (Grier, 1971). This analysis showed that response bias was indistinguishable between the dominant ($M=.64$, $SD=.37$) and nondominant quadrants ($M=.65$, $SD=.28$), $t(19)=-.07$, $p=.94$. This suggests that, although subjects seemed more inclined to indicate “no change” than “change”, this bias was equally likely for both conditions. The same result was obtained using two alternate measures of bias (i.e., B''_D ; Donaldson, 1992; or B''_H ; Hodos, 1970). Thus, the results suggest that improved change detection in the dominant quadrant was not caused by a selective reduction in the response threshold for indicating a change in that quadrant. Instead, we conclude that the increased probability of targets in the dominant quadrant elicited a bias to encode the items from that location, thereby enhancing later sensitivity to changes.

Two subjects out of 20 were able to correctly identify the dominant quadrant when first queried about whether they had noticed any location that was more likely to contain changes. 15 subjects reported having no knowledge about the probability manipulation. However, when these subjects were forced to select a quadrant, 12 of 20 observers correctly guessed the dominant quadrant (exceeding chance performance: $\chi^2(1)=9.8$, $p<.01$). Thus, observers were able to access explicit knowledge about the probability manipulation. Yet, those subjects who guessed the dominant quadrant correctly did not show a larger bias ($n=12$, $M=13\%$, $SD=.18$) than those who did not ($n=8$, $M=7\%$, $SD=.19$), $t(18)=.75$, $p=.47$; this conclusion is reinforced by a combined analysis of Experiment 2 and 3 below. Moreover, no observer reported using a biased encoding strategy. Thus, although subjects clearly had some explicit knowledge of the dominant quadrant, it remains unlikely that this knowledge motivated a voluntary encoding bias towards the dominant quadrant.

Experiment 3

Experiments 1 and 2 suggest that the probability manipulation influenced the allocation of limited resources for storage in working memory, so that more information was retained from the dominant quadrant. The goal of Experiment 3 was to further characterize how these statistical regularities influenced the allocation of resources in visual working memory. According to “flexible resource” models, capacity in working memory is determined by a central pool of resources that can be flexibly divided between the items to be stored, such

that items of greater importance or complexity can be granted a larger proportion of this shared resource pool (e.g., Alvarez & Cavanagh, 2004; Bays and Hussein, 2008; Eng, Chen & Jiang, 2005, Wilken & Ma, 2004). A key assumption of flexible resource models is that the precision or resolution of each representation in working memory is determined by the relative proportion of resources devoted to that item. From this perspective, the optimal encoding strategy given a spatial bias in target position would be to allocate a larger share of resources to items in the dominant quadrant, leading to higher resolution representations of those items. By contrast, discrete “slot-based” models suggest that capacity in working memory is determined by a limited number of discrete “slots”, each of which can be assigned to a single item (Barton, Ester, & Awh, 2009; Rouder, Morey, Cowan, Zwillling, Morey & Pratte, 2008; Zhang & Luck, 2008). In this case more slots could be devoted to the dominant quadrant, leading to an increased probability that those items would be stored. Biases in the assignment of discrete storage slots, however, would lead to equivalent mnemonic resolution for all items that were assigned to a slot, regardless of whether they appeared in the dominant or the non-dominant quadrant.

To examine which of these models provides the best description of how statistical learning biases encoding into working memory, Experiment 3 employed a procedure that enabled separate estimates of the number and the resolution of the memory representations stored from the dominant and non-dominant quadrants. This was achieved by manipulating the similarity between the critical item in the sample array and the new item that replaced it in the test array. The rationale for this design follows from recent evidence (Awh et al., 2007; Barton, Ester and Awh, 2009; Jiang, Shim and Makovski, 2008; Scolari, Vogel and Awh, 2008) that the primary limiting factor for change detection performance depends critically on the size of the changes between the sample and test arrays. When the changes are large, accurate change detection depends primarily on whether or not the critical item has been stored in working memory (consistent with the underlying assumptions of the formula used to calculate capacity estimates (Pashler, 1988; Cowan, 2000)). Thus, the rate at which large changes are detected provides an estimate of the number of items encoded into working memory. By contrast, when changes are small (i.e., when the psychological similarity of the sample and test is high), then even when the critical item is stored there is a higher probability of comparison errors in which observers fail to perceive the difference between the sample and test items. Given that increasingly precise representations of the sample are required to detect such small changes, the incidence of such comparison errors may provide a useful operational definition of mnemonic resolution (Awh et al, 2007; Barton, Ester and Awh, 2009).

Multiple studies of delay-specific neural activity bolster our claim that change detection with low sample-test similarity is limited by the number of items that can be stored in working memory. These studies have shown that neural activity in the parietal cortex reaches asymptote at the same set size where behavioral estimates of the number of stored items reach asymptote (Vogel et al, 2004; Vogel et al, 2005; Todd and Marois, 2004; Todd and Marois, 2005) establishing a critical link between parietal activity and behavioral estimates of capacity in working memory. Importantly, Xu and Chun (2006) used both simple and complex items in a similar task, and found once again that the same parietal region was sensitive to the number rather than the complexity of the stored items. Thus, parietal activity reflects the number rather than the complexity of the stored items, and parietal activity is predicted by success in the detection of big changes. By contrast, activity in distinct cortical regions maps onto performance when changes are smaller and mnemonic resolution becomes the limiting factor. Thus, neural data converge with behavioral data to suggest that it is productive to distinguish between the number and the resolution of the representations in working memory, and that the detection of big and small changes can

provide useful estimates of these two aspects of memory ability (Awh et al, 2007; Xu and Chun, 2006).

To distinguish between flexible resource and slot-based models of how encoding into working memory was biased by the probability manipulation, we estimated both the number and the resolution of the stored representations in the dominant and non-dominant quadrants. To reiterate, flexible resource models predict that the optimal strategy would be to allocate a larger proportion of resources to items in the dominant quadrant, thereby boosting mnemonic resolution for those items. By contrast, discrete slot-based models predict that items in the dominant quadrant should be more likely to be stored, but that resolution should be equivalent for each stored item regardless of which quadrant it occupies.

Methods

Subjects

Sixteen students from the University of Oregon participated in a 90 min. experimental session for psychology course credits. All subjects reported normal or corrected-to-normal visual acuity.

Stimuli

All aspects of stimuli were equivalent to those in Experiment 1b with the following exceptions. In addition to colored squares, a new set of stimuli were included (in separate trials) that allowed a straightforward manipulation of sample-test similarity. The new stimuli (illustrated in Figure 4) included two ovals ($1.1^\circ \times 2.9^\circ$) and two rectangles ($2.4^\circ \times 1.7^\circ$) with patterned surfaces. These stimuli were randomly sampled to create 8-item sample arrays with the constraint that no specific object appeared more than 3 times within a single display. During change trials, the critical item was replaced by selecting at random one of the remaining three items. We reasoned that when an oval changed into a rectangle (or vice versa) that the large size of the change would minimize the incidence of comparison errors; thus performance in this condition was used to estimate the number of items that could be maintained in that condition. By contrast, when an oval changed into an oval (or a rectangle changed into a rectangle) the changes were relatively small, such that even when the critical item was stored we expected an increased probability of comparison errors. Thus, based on the assumption that clearer representations are needed to detect smaller changes, performance during these small change trials was used to derive an estimate of mnemonic resolution (details of this derivation described in the Results section below). Just as with the color stimuli, changes occurred with probability .5. During change trials with the ovals and rectangles, big and small changes occurred equally often. Finally, only 8-item displays were presented in Experiment 3.

Procedures

The general procedure was similar to that of Exp.1b. Six blocks (90 trials each) were administered to the participants. For 30 of those trials, we presented color stimuli identical to those in Experiment 1b (15 change and 15 no-change trials). For the remaining 60 trials, the ovals and rectangles were presented (30 no-change trials, 15 big-change trials, and 15 small-change trials). The duration of the sample display was increased to 500 ms to ensure adequate time to encode the more complex rectangle and oval stimuli. When the test display was presented, subjects indicated whether one of the items had changed or not. The same post-experiment questionnaire as in Experiment 2 was administered after the session to assess subjects' awareness of the probability manipulation.

Results and Discussion

Accuracy during change trials in Experiment 3 is illustrated in Figure 5. Replicating the previous observations, a reliable bias effect was obtained for the color trials, such that accuracy during change trials was significantly higher for the dominant quadrant (79%) than for the nondominant quadrant (69%), $t(15)=2.99$, $p=.01$. A bias effect reached 11% by the 3rd block, and a 2 (probability manipulation) by 6 (block) analysis of variance showed a significant main effect of probability manipulation, $F(1,15)=9.20$, $p=.01$, $\text{Partial}\eta^2=.38$, no main effect of block ($F(5,75)=.68$, $p=.64$), and no interaction between probability manipulation and block ($F(5,75)=.87$, $p=.51$).

Because previous evidence suggests that distinct limiting factors determine performance during trials with big (i.e., oval-to-rectangle or rectangle-to-oval changes) and small (i.e., oval-to-oval or rectangle-to-rectangle) changes (Awh et al., 2007; Barton et al., 2009) separate analyses were carried out for big and small changes in the oval and rectangle displays. First, we note that the pattern of individual differences in the big and small change trials supported previous finding (Awh et al., 2007) and suggestions that these conditions – despite employing the same sample stimuli within the same blocks of trials – measure distinct aspects of working memory ability. Specifically, capacity estimates from the big change condition were reliably correlated with capacity estimates obtained with the color stimuli ($n=16$, $r = .65$, $p = .01$), consistent with our hypothesis that performance in both conditions was determined by a common limiting factor, the number of items that could be stored in working memory. By contrast, although a split-half correlation on the small change condition during odd and even trials revealed that reliability was good for this measure ($n=16$, $r = .77$, $p < .01$), capacity estimates from the small change condition were not correlated with those from the color stimuli ($n=16$, $r = .05$, $p = .87$), consistent with the idea that performance in the small change condition is limited by a qualitatively different aspect of memory ability (i.e., mnemonic resolution) from the other conditions.

The key results, however, concerned how the probability manipulation influenced the number and the resolution of the representations encoded from the dominant quadrant. Here, the big change trials showed the same pattern of bias that was observed in the previous experiments, with 70% accuracy for changes in the dominant quadrant and 60% accuracy for changes in the nondominant quadrant, $t(15)=2.74$, $p=.02$. A 14% bias effect emerged during the 1st block, and a 2 (probability manipulation) by 6 (block) analysis of variance revealed a significant main effect of the probability manipulation, $F(1,15)=7.52$, $p=.02$, $\text{Partial}\eta^2=.33$, no main effect of block ($F(5,75)=1.41$, $p=.23$) and no interaction between probability manipulation and block ($F(5,75)=.26$, $p=.93$). Thus, when the limiting factor for change detection was the number of items that could be held in working memory (i.e., in both the color trials, and the big change trials with ovals and rectangles), a clear advantage was observed when the critical item appeared in the dominant quadrant. This mirrors the findings in the earlier experiments to suggest that a larger number of items were retained from the dominant quadrant. A different result emerged when we examined performance in the small-change trials where performance was limited by mnemonic resolution. No reliable difference was found between resolution-limited change detection in the dominant (41%) and non-dominant (37%) trials, $t(15)= 1.8$, $p = .09$. Taken at face value, this result suggests that the implicit encoding bias did not lead to higher resolution representations in the dominant quadrant, a result that falls in line with the predictions of discrete slot-based models of resource allocation. Nevertheless, there was a numerical advantage for change detection accuracy in the dominant quadrant, and a trend towards a reliable difference between accuracy in the dominant and non-dominant quadrants. Recall, however, that a precise analysis of the small-change trials should acknowledge that errors could have occurred either due to a failure to store the critical item, or due to errors in comparing the

sample and test items (because of limited mnemonic resolution); that is, accuracy in the small-change trials was really a composite measure that was influenced by both the number and the resolution of the stored representations. Thus, to obtain a clearer estimate of whether mnemonic resolution was influenced by the implicit bias, it was necessary to correct for errors in the small change trials that were due to storage failures. We accomplished this by using performance in the large change trials to estimate (for each observer) the number of ovals and rectangles that were stored in working memory. This estimate allowed us to calculate the probability of storage failures in the small change condition, and then to derive the probability of correct comparison between the sample and test *given* that the critical item was stored (Barton, Ester & Awh, 2009).

Here we describe the logic of this analytic approach. Accuracy (Acc) in the small-change condition was assumed to be equal to the probability that the object actually was encoded into working memory (P_{mem} , where $P_{\text{mem}} = k/\text{set size}$, and k is estimated using the big change and no-change trials (76%) with the oval/rectangle stimuli) multiplied by the probability that the sample and test were compared correctly (C), plus a correction for guessing based on the assumption that subjects would guess correctly half of the time when the object was not stored.

$$\text{Acc} = (P_{\text{mem}} * C) + (1 - P_{\text{mem}}) / 2$$

Solving for C :

$$C = \{\text{Acc} - (1 - P_{\text{mem}}) / 2\} / (P_{\text{mem}})$$

Using the probability of correct comparison (C) as an operational measure of mnemonic resolution (see Figure 6), we found that resolution was no higher for items represented in the dominant quadrant ($C = 69\%$) than for those in the non-dominant quadrant ($C = 67\%$), $t(15) = .39$, $p = .70$. Thus, although the results from the big change and color trials suggested that more items were stored from the dominant quadrant, we found no evidence that resolution of those representations was influenced in a similar way. These findings are inconsistent with a flexible resource model of the encoding bias, because those models suggest that a shared resource pool can be allocated in a continuous fashion, with a larger share of resources allocated to higher priority items. In this case, a bias towards the dominant quadrant should have caused a disproportionate allocation of resources to items in the dominant quadrant, leading to higher mnemonic resolution for those items. Instead, our results suggest that statistical regularities in target position influenced encoding into visual working memory in a discrete slot-based fashion. It is worth noting that the relatively low performance for the small change condition was unlikely to be due to a strategic difference in which interleaving both the big and small changes within a single block might have encouraged participants to adopt a “low-resolution” strategy for encoding sample arrays. That is, subjects might have chosen to remember as many items as they could (i.e., outline shapes) at the expense of storing detailed information about each object (i.e., surface patterns). If subjects had chosen not to store the details about surface patterns, then a failure to find changes in their detection of changes in those patterns would be difficult to interpret; in that case, a null result might simply reflect performance at floor. The observers in this experiment did store enough detail to detect these small changes, however, as shown by above-chance performance in the small change conditions ($k=1.2$), and a mean probability of correct comparison well above 50% (see Figure 6). Awh et al. (2007) also provided some evidence against a shift to a “low-resolution” mode in mixed blocks by showing no significant differences in the detection of small changes between pure blocks of small

change trials (that would not elicit this putative strategy) and blocks in which approximately equal numbers of changes were large and small.

Finally, only 2 of 16 subjects successfully identified the dominant quadrant in response to a query about whether any location had been more likely to contain a change item, and half the subjects reported not to have noticed any probability manipulation. However, when we administered a more sensitive forced choice test, subjects were able to guess the correct quadrant at above chance levels (11 of 16 subjects), $\chi^2(1)=12.25$, $p<.01$. Yet once again, those who correctly identified the dominant quadrant in the forced choice test did not show a larger bias effect ($n=11$, $M=9\%$, $SD=.15$) than those who did not ($n=5$, $M=13\%$, $SD=.08$), $t(14)=-.55$, $p=.59$. Thus, although subjects could access some explicit knowledge about the probability manipulation, that knowledge was unlikely to have motivated a voluntary strategy to orient towards the dominant quadrant. Crucially, no subject reported any overt strategy to bias encoding towards the dominant quadrant, just as in all of the previous studies. Additionally, the combined data from the neutral condition of Experiment 2 (when a reliable bias was already established) and Experiment 3 also showed that subjects who could correctly guess the dominant quadrant ($n=23$, $M=.11$, $SD=.16$) did not show a larger bias effect than those who could not ($n=13$, $M=.09$, $SD=.16$), $t(34)=.35$, $p=.73$. Thus, consistent with the previous studies, the overall results argue against the hypothesis that the observers' explicit knowledge of the probability manipulation motivated a voluntary encoding bias towards the dominant quadrant.

General discussion

A growing body of evidence suggests that the ability to exert voluntary control over what is encoded into working memory is strongly intertwined with the effective capacity of that system (Awh, Vogel and Oh, 2006; Engle, 2002; Kane, Bleckley, Conway and Engle, 2001; Vogel, McCollough and Machizawa, 2005). In the present work, we demonstrate and further characterize how statistical regularities in target position can also exert an influence on which items will gain entry into this highly capacity-limited system. When targets were more likely to appear in a specific quadrant of the display, observers were more likely to encode items from that region into working memory. This is in accord with previous research showing a facilitated deployment of attention towards locations containing higher probability of target occurrence, enhancing speed or quality with which single targets were processed, and biasing which items are encoded into memory (Olson et al., 2005; Beck et al., 2008).

The encoding bias in the current study was observed in five independent samples of subjects, even though subjects were not explicitly informed about the probability manipulation. A surprisingly small number of participants were able to articulate the nature of the probability manipulation (9 out of 84), and none reported any kind of explicit strategy to preferentially encode items from the dominant quadrant. Nevertheless, forced choice tests in Experiments 2 and 3 provided clear evidence of above-chance discrimination of the dominant quadrant; thus, observers were able to access some explicit knowledge about the bias in target position. This fact notwithstanding, multiple pieces of evidence argue against the hypothesis that the observed encoding bias was a product of a voluntary bias in observers' encoding strategies. Perhaps the most central piece of evidence against a voluntary encoding bias is that none of the 84 subjects reported a voluntary encoding bias towards the dominant quadrant. In addition, when observers were split by virtue of whether they succeeded or failed in identifying the dominant quadrant during the forced choice test, the bias effect was no larger for subjects who had succeeded than for those that had not. Similarly, when the data were sorted by whether observers had shown a bias towards the dominant quadrant (as operationalized by better change detection in the dominant quadrant),

those who showed the bias did not perform better on the forced choice test. Finally, it should be emphasized that the mere presence of some explicit knowledge does not provide strong evidence that that knowledge elicited a voluntary strategy on the part of observers. That is, although it is prudent to keep in mind that explicit knowledge may be present in various statistical learning procedures, there is also strong evidence that statistical learning may influence cognitive processing in a non-strategic fashion. These considerations motivate our hypothesis that the encoding bias was a product of the same kind of statistical learning operations that previous research has shown to influence the deployment of visual attention and mnemonic resources.

This kind of dissociation between encoding bias and the explicit goals of the observer is a striking possibility in the context of visual working memory, to the extent that this “online” memory system is often seen as synonymous with the current contents of awareness (e.g., Cowan, 1988; Jonides, Lewis, Nee, Lustig, Berman and Moore, 2008; Oberauer, 2002). Nevertheless, the plausibility of this hypothesis is highlighted by a substantial body of previous research showing that statistical regularities can be apprehended and used to guide behavior – whether intentionally or not -- during sequence learning (Lewicki, Czyzewska, and Hoffman, 1987; Nissen & Bullmer, 1987; Lewicki, Hill, and Bizot, 1988; Frensch, 1998), visual search (Chun and Jiang, 1998; Chun and Jiang, 1999; Chun, 2000), temporal order judgments (Turk-Browne, Junge, and Scholl, 2005), linguistic processing (Gebhart, Aslin, and Newport, 2009) and task switching (Mayr and Bryck, 2005). For instance, Lewicki et al. (1987) employed a paradigm in which specific sequences of stimulus presentations predicted the upcoming location of a target. Faster reaction times at the predicted locations suggested that participants could benefit from regularities in the order with which targets were presented, although participants were unable to articulate such statistical regularities. The present work suggests that a similar apprehension of statistical regularities in target position may influence encoding into visual working memory, and that this knowledge may be applied without the explicit intent of the observer.

Finally, Experiment 3 revealed that while items in the dominant quadrant were more likely to be stored in working memory, the resolution of those representations was no better than for items encoded from the other quadrants. These findings suggest that statistical learning biased encoding into working memory in a discrete slot-based fashion, rather than causing an asymmetric division of resources between items stored from the dominant and non-dominant quadrants. This account of how statistical learning biases encoding into working memory makes a clear prediction. Specifically, this biasing effect should be evident only when the total number of items to be stored exceeds the working memory capacity of the observer. This is because biases in the selection of items to be stored – as opposed to biases in the relative quality of a given representation – should not influence performance when there is sufficient capacity to store all the relevant items.

Although it would be premature to assume that statistical learning and voluntary selection are equivalent with respect to flexible versus discrete modes of resource allocation, we note that previous studies that have examined the flexibility of voluntary resource allocation in visual working memory have yielded similar conclusions. Zhang and Luck (2008) used a cueing procedure in which observers were informed that a specific item in a three item display was most likely to be probed at the end of the trial. Their results showed that the cued items were more likely to be stored in working memory, but they found only modest variations in resolution that could not disconfirm a slot-based account. Likewise, Barton, Ester and Awh (2009) examined whether more complex items in a display command a disproportionate share of resources as predicted by flexible resource models. This was accomplished by measuring whether resolution for a given item was affected by large variations in the complexity of the other to-be-stored items, as predicted by the hypothesis

that more complex items demand a larger share of a central resource pool. Inconsistent with a flexible resource account, mnemonic resolution for a given item was completely insensitive to the complexity of the other items to be stored. Thus, neither valid precues that clearly specify object priority, nor large disparities in information load across items were able to elicit evidence of asymmetric resource allocation across stored items. Instead, these strong manipulations had a selective influence on the probability of storage, consistent with a discrete slot-based model of resource allocation.

Our results suggest that statistical learning may help to optimize the allocation of limited resources in working memory by biasing encoding towards behaviorally relevant items. In the current experiments, this bias enhanced performance by raising the probability that the changed item was encoded into memory. Thus, while the selection of which items to store from supra-span displays may typically be relatively haphazard³, statistical regularities in target position may shape this process by highlighting a subset of the to-be-stored items. Given the evidence we've presented that the bias may not have been voluntary in nature, it is possible that statistical learning provides one mechanism for easing the burden on a highly capacity limited system for voluntary selection and storage in working memory. As Lewicki, Hill, and Czyzewska (1992) pointed out, statistical learning processes are sensitive to structurally complex patterns that may be difficult if not impossible to apprehend solely through consciously controlled learning processes. For example, Lewicki, Czyzewska and Hoffman (1987) found evidence that observers nonconsciously acquired knowledge about sequences of locations that were described in terms of a four-way interaction (Lewicki, et al., 1992). The fact that statistical learning is sensitive to such complex patterns in the observers' experience suggests that this process may provide a powerful complement to voluntary control processes that are not only highly limited in capacity, but which may also be less sensitive to informative regularities in visual experience.

Finally, we highlight a number of remaining questions regarding the boundary conditions of the statistical learning effects documented in the present work. For instance, further work is needed to determine the specificity of the interactions between statistical learning and mnemonic encoding. Is encoding into working memory biased only by patterns that are behaviorally relevant to the observer (e.g., Turk-Browne, Junge, and Scholl, 2005), or would encoding be biased towards any region that consistently contained salient or attention-capturing events, regardless of whether those events were relevant to the ongoing goals of the observer? Would biases obtained with one set of stimuli also influence encoding patterns with novel stimuli or a new task context (Jiang and Song, 2005 JEP:HPP; Jiang and Song, 2005)? In the present studies, we manipulated the spatial position of the change target; would similar effects be observed if similar manipulations were applied to non-spatial dimensions such as color and shape (e.g., Chun and Jiang, 1999; Fiser and Aslin, 2001; but see also Beck, Angelone, Levin, Peterson and Varakin, 2008). Such explorations of the boundary conditions of this bias effect may provide a richer understanding of how statistical learning influences the deployment of limited resources for storage in working memory.

Acknowledgments

Portions of this work were supported by NIH-R01MH087214-01 to E.A. and E.K.V.

³Unless bottom-up stimulus characteristics such as perceptual grouping (Woodman, Vecera and Luck, 2003) or attentional capture (Schmidt, Vogel, Woodman and Luck, 2002) are at work.

References

- Awh E, Barton B, Vogel E. Visual Working Memory Represents a Fixed Number of Items Regardless of Complexity. *Psychological Science* 2007;18:622–628. [PubMed: 17614871]
- Awh E, Vogel EK, Oh S-H. Interactions between attention and working memory. *Neuroscience* 2006;139:201–208. [PubMed: 16324792]
- Alvarez GA, Cavanagh P. The Capacity of Visual Short-Term Memory Is Set Both by Visual Information Load and by Number of Objects. *Psychological Science* 2004;15:106–111. [PubMed: 14738517]
- Barton B, Ester E, Awh E. Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance* 2009;35:1359–1367. [PubMed: 19803642]
- Bays PM, Hussain M. Dynamic shifts of limited working memory resources in human vision. *Science* 2008;321:851–854. [PubMed: 18687968]
- Beck MR, Angelone BL, Levin DT, Peterson MS, Varakin DA. Implicit learning for probable changes in a visual change detection task. *Consciousness and Cognition* 2008;17:1192–1208. [PubMed: 18684642]
- Chun MM. Contextual cueing of visual attention. *Trends in Cognitive Sciences* 2000;4:170–178. [PubMed: 10782102]
- Chun MM, Jiang Y. Contextual Cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology* 1998;36:28–71. [PubMed: 9679076]
- Chun MM, Jiang Y 7. Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science* 1999;10(4):360–65.
- Cowan N. Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information processing system. *Psychological Bulletin* 1988;104:163–191. [PubMed: 3054993]
- Cowan N. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 2000;24:87–185. [PubMed: 11515286]
- Donaldson W. Measuring recognition memory. *Journal of Experimental Psychology: General* 1992;121:275–277. [PubMed: 1402701]
- Eng HY, Chen D, Jiang Y. Visual working memory for simple and complex visual stimuli. *Psychonomic Bulletin and Review* 2005;12:1127–1133. [PubMed: 16615339]
- Engle R. Working memory capacity as executive attention. *Current Directions in Psychological Science* 2002;11(1):19–23.
- Fiser J, Aslin RN. Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science* 2001;12:499–504. [PubMed: 11760138]
- Fiser J, Aslin RN. Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2002;28:458–467.
- Frensch, PA. One concept, multiple meanings. In: Standler, M.; Frensch, P., editors. *Handbook of implicit learning*. Sage; Thousand Oak: 1998.
- Gebhart AL, Aslin RN, Newport EL. Changing structures in midstream: Learning along the statistical garden path. *Cognitive Science* 2009:1–30.
- Geng JJ, Behrmann M. Spatial probability as an attentional cue in visual search. *Perception and Psychophysics* 2005;67:1252–1268. [PubMed: 16502846]
- Grier JB. Nonparametric indexes for sensitivity and bias: Computing formulas. *Psychological Bulletin* 1971;75:424–429. [PubMed: 5580548]
- Hodos W. Nonparametric index of response bias for use in detection and recognition experiments. *Psychological Bulletin* 1970;74:351–354.
- Hoffmann J, Kunde W. Location-specific target expectancies in visual search. *Journal of Experimental Psychology: Human Perception and Performance* 1999;25:1127–1141.
- Jiang Y, Shim WM, Makovski T. Visual working memory for line orientations and face identities. *Perception and Psychophysics* 2008;70(8):1581–1591. [PubMed: 19064500]

- Jiang Y, Song JH. Hyperspecificity in visual implicit learning: Learning of spatial layout is contingent on item identity. *Journal of Experimental Psychology: Human Perception and Performance* 2005;31:1439–1448. [PubMed: 16366800]
- Jiang Y, Song JH. Spatial context learning in visual search and change detection. *Perception & Psychophysics* 2005;67:1128–1139. [PubMed: 16502835]
- Jonides J, Lewis RL, Nee DE, Lustig CA, Berman MG, Moore KS. The mind and brain of short-term memory. *Annual Review of Psychology* 2008;59:193–224.
- Kane MJ, Bleckley MK, Conway ARA, Engle R. A controlled-attention view of WM capacity. *Journal of Experimental Psychology: General* 2001;130:169–83. [PubMed: 11409097]
- Lewicki P, Czyzewska M, Hoffman H. Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1987;13:523–530.
- Lewicki P, Hill T, Bizot E. Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology* 1988;20:24–37. [PubMed: 3338266]
- Lewicki P, Hill T, Czyzewska M. Nonconscious acquisition of information. *American Psychologist* 1992;27:796–801. [PubMed: 1616179]
- Luck SJ, Vogel EK. The capacity of visual working memory for features and conjunctions. *Nature* November 20;1997 390:279–281. [PubMed: 9384378]
- Mayr U, Bryck R. Sticky rules: Integration between abstract rules and specific actions. *Journal of Experimental Psychology: Learning, Memory and Cognition* 2005;31(2):337–350.
- McNab F, Klingberg T. Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience* 2008;11:103–107.
- Miller J. Components of the location probability effect in visual search tasks. *Journal of Experimental Psychology: Human Perception and Performance* 1988;14:453–471. [PubMed: 2971773]
- Miller J. Components of the location probability effect in visual search tasks. *Journal of Experimental Psychology: Human Perception and Performance* 1988;14:453–471. [PubMed: 2971773]
- Nissen MJ, Bullemer P. Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology* 1987;19:1–32.
- Oberauer K. Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory and Cognition* 2002;28:411–421.
- Olson IR, Jiang Y, Moore KS. Associative learning improves visual working memory performance. *Journal of Experimental Psychology: Human Perception and Performance* 2005;31:889–900. [PubMed: 16262486]
- Pashler H. Familiarity and visual change detection. *Perception and Psychophysics* 1988;44:369–378. [PubMed: 3226885]
- Reeves A, Sperling G. Attention gating in short term visual memory. *Psychological Review* 1986;93:180–206. [PubMed: 3714927]
- Rouder JN, Morey RD, Cowan N, Zwilling CE, Morey CC, Pratte MS. An assessment of fixed-capacity models of visual working memory. *Proceedings of the National Academy of Sciences* 2008;105:5975–5979.
- Schmidt BK, Vogel EK, Woodman GF, Luck SJ. Voluntary and automatic control of visual working memory. *Perception and Psychophysics* 2002;64:754–763. [PubMed: 12201334]
- Scolari M, Vogel E, Awh E. Perceptual Expertise Enhances the Resolution But Not the Number of Representations in Working Memory. *Psychological Bulletin and Review* 2008;15:215–222.
- Shaw ML. A capacity allocation model for reaction time. *Journal of Experimental Psychology: Human Perception and Performance* 1978;4:586–598.
- Shaw ML, Shaw P. Optimal allocation of cognitive resources to spatial locations. *Journal of Experimental Psychology: Human Perception and Performance* 1977;3:201–211. [PubMed: 864393]
- Sperling G. The information available in brief visual presentations. *Psychological Monographs: General and Applied* 1960;74(11) Whole No. 498.
- Todd JJ, Marois R. Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature* 2004;428:751–754. [PubMed: 15085133]

- Todd JJ, Marois R. Posterior parietal cortex activity predicts individual differences in visual short-term memory capacity. *Cognitive, Affective, & Behavioral Neuroscience* 2005;6:144–155.
- Turk-Browne NB, Junge JA, Scholl BL. The automaticity of visual statistical learning. *Journal of Experimental Psychology: General* 2005;134:552–564. [PubMed: 16316291]
- Vogel EK, Machizawa MG. Neural activity predicts individual differences in visual working memory capacity. *Nature* 2004;428:749–751.
- Vogel EK, McCollough AW, Machizawa MG. Neural measures reveal individual differences in controlling access to visual working memory. *Nature* 2005;438:500–503. [PubMed: 16306992]
- Vogel EK, Woodman GF, Luck SJ. The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance* 2006;32:1436–1451. [PubMed: 17154783]
- Wilken P, Ma WJ. A detection theory account of change detection. *Journal of Vision* 2004;4:1120–1135. [PubMed: 15669916]
- Xu Y, Chun MM. Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature* 2006;440:91–95. [PubMed: 16382240]
- Woodman GF, Vecera SP, Luck SJ. Perceptual organization influences visual working memory. *Psychonomic Bulletin and Review* 2003;10:80–87. [PubMed: 12747493]
- Zhang W, Luck SJ. Discrete fixed-resolution representations in visual working memory. *Nature* 2008;453:233–235. [PubMed: 18385672]

Figure 1.
The sequence of events in a single trial of the change detection task.

Figure 2.
Change detection accuracy as a function of the quadrant that held the changed item in Experiments 1a, 1b, and 1c.

Figure 3. Performance (HR-FA) during the neutral phase in Experiment 2. Accuracy calculated by the hit-rates minus false-alarms during the neutral phase in Experiment 2.

Figure 4.
The geometric stimuli used in Experiment 3.

Figure 5.
Accuracy during change trials as a function of the type of change in Experiment 3.

Figure 6.
The probability of correct comparison (C) given that the item is stored in the dominant and non-dominant quadrants of Experiment 3.