6/14/2018

Differentiating the Components of Visual Short-Term Memory with Proactive Interference

An Honors Thesis in Cognitive Science

Robert Walter
University of California, San Diego

Advisors: Dr. Timothy F Brady, Dr. Andrea Chiba and Eric Leonardis

Differentiating the Components of Visual Short-Term Memory with Proactive Interference

Robert Walter and Timothy F. Brady

Proactive interference (PI) is when an item previously held in memory interferes with a new memory item. Previous studies investigating PI in WM with real-world objects have shown it to be critically dependent on the method of presentation. When items are presented sequentially, there is a large PI effect (Endress & Potter, 2014). When objects are presented simultaneously at different spatial locations, there is little effect of repeating items trial to trial (Makovski, 2016), suggesting little role for LTM in the standard change detection task, even with real-world objects.

We asked if PI can be used to distinguish the components of visual short-term memory tasks, such as LTM, WM, and feature binding. While many of our experiments weren't definitive, we do demonstrate that the PI in real-world objects is not dependent on spatial location: participants were as likely to select a foil from the previous trial that did not match the current target location as one that did. We also see that some visual features are more prone to PI than others.

The lack of location specificity may indicate that the main distinction between simultaneous and sequential presentation is not location per se. Instead, participants may rely on different strategies when encoding simultaneous displays and sequential displays, relying more on LTM representations in sequential presentation. This is further backed by our second experiment that shows the importance of location information within WM, and LTM may be recruited when location is uninformative.

Visual working memory (VWM) is operationally defined as a system to hold and manipulate visual information. This manipulation aspect distinguishes it from short-term memory, which is usually held as simply a temporary store; however, the two terms are sometimes used interchangeably (Baddeley, 2012). It is regarded as a core cognitive process, utilized for a wide range of behaviors (Baddeley, 2003; Ma et al., 2014), and is highly correlated with academic performance and fluid intelligence (Alloway & Alloway, 2010; Fukuda et al, 2010). VWM has a limited capacity, usually around 4 objects (Cohen, 2001) and requires active maintenance, leading to a limited duration, usually on the scale of seconds (Schurgin, 2018). There is much debate on the sources of these limits, as the components of VSM are not well understood (Baddeley, 2012).

VWM is often contrasted with visual long-term memory (VLTM), which passively stores visual information over a large amount of time (Brady, Konkle, & Alvarez, 2011), at a near limitless capacity (Standing, 1973) in an extremely detailed manner (Brady et al 2008). VLTM may possibly be neurologically distinct, as it is associated with the medial temporal lobe and hippocampus, whereas VWM is associated with the occipital and parietal cortex (Schurgin, 2018). Another distinction is the degree of interference in VLTM vs VWM, which is when one memory item interferes with the encoding, storage, or retrieval of another memory item (Underwood, 1957).

Because VWM and VLTM almost always function simultaneously and dynamically, they can be very difficult to pull apart (Schurgin, 2018; Baddeley, 2012), and some aspects of VLTM have possibly been misattributed to VWM (Lin & Luck, 2012). Many critical aspects of VWM are still under investigation, such as how it performs object recognition, how individual visual features are bound into whole objects, and many more (Schurgin 2018).

Our study aims to differentiate the components of short-term memory tasks to examine the influencing features of VWM, VLTM, and object binding. We attempt to use proactive interference to help pull apart these differences. Interference can be held in two forms: proactive and retroactive. Retroactive is when new memories interfere with previous memories; and proactive is when old memories interfere with newer memories. Interference theory holds that retrieval competition is the prime limitation on long-term memory and that working memory's primary function is to limit this interference (Engel, 2002). Behavioral evidence suggests that moderate proactive interference (PI) can occur in short-term memory tasks (Makovski, 2008 & 2016); however, some evidence has pointed to a much more pronounced effect of PI (Endress and Potter, 2014). While this PI is objectively evident in these behavioral tasks, there have been suggestions that its origins may stem from the use of VLTM (Makovski 2016; Lin & Luck 2012). Our first series of experiments aims to see how VWM and VLTM contribute to behavioral PI and examine how changes in the stimuli may influence the PI susceptibility or memory strategy. The second series of experiments aims to look at the characteristics of object binding, also using PI. Both series of experiments use an AFC task where some answer choices originate from the current trial, while other originate from the previous trial and a novel choice not seen in either is held as a baseline.

Experiment 1

Differentiating VWM and VLTM Proactive Interference

Does PI interference in short-term visual memory tasks stem from VWM or LTM? This is difficult to separate because when we are presented a visual stimulus, it is likely that this information is simultaneously being encoded in both memory systems (Engel, 2002). However, while both can be utilized for short-term information storage, only VWM is associated with active manipulation, and this can be used to create conditions indicative of only VWM. In this series of experiments, objects will be memorized primarily by spatial location; however, subjects will mentally rotate the objects after visual presentation so that the updated visual-spatial information is only held in a VWM representation. We can then see if this updated location is prone to PI, or if PI is confined only to the unmanipulated visual information.

Experiment 1a

PI in simple stimuli after mental location updates

This experiment was largely to investigate the discrepancy between claims of PI found with simple stimuli (Makovski, 2008) and it not being found (Lin & Luck, 2012) and to identify the source of this PI. Our original hypothesis was that any PI found would be a result of LTM.

Methods

Participants. 42 subjects between 18 and 25 were recruited from UC San Diego's online recruitment system, SONA, who received class credit for their research participation. The only restrictions included corrected or uncorrected normal color vision. Every participant was successful in reciting the numbers from Ishihara's plates 2-8 in order to confirm normal color vision. All 42 subjects completed the task and were included in the study.

Materials All subjects were tested individually, seated in an approximately 1.8x1.5-meter sound-attenuated room with stimuli presented on an iMac pro computer with a 60 Hz, 55 cm diagonal viewable screen. There was no instruction on how far to sit from the monitor and participants sat as they felt comfortable. MATLAB's Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) was used to run the experiment.

Stimuli and Task Figure 1 illustrates stimuli and AFC task. Every trial displayed an array of 4 color disks (2.3 cm in diameter) in a box configuration, approximately 6.2 cm between centers of each disk, presented for 1000 ms. The colors of the disks would then disappear, leaving the outlines of each disk, and the array of 4 disks would rotate clockwise 90 degrees over a period of 1500 ms. Subjects were instructed to imagine the colors of each disk moving with the outlines to the new location. After rotation, the perimeter of one of the disk outlines would distinctively thicken to signify it as the target cue. Seven color disks would appear in a line above the array to signify a 7 AFC task. The participant was instructed to select the disk that had the same color that should be in the cued location after the rotation. Among the 7 AFC task was the following options: the correct color after the rotation; the color at the location before rotation; one of the two remaining colors from the current trial, randomly selected; the correct color after the rotation from the previous trial (N-1); the color at the location before rotation of the previous trial (N-1); one of the two remaining colors from the previous trial(N-1); and a novel color not seen in the current(N) or previous trial(N-1). These disks were presented in a randomly shuffled order. After the participant made a selection, they were asked to press a key to continue to the next trial. There was a total of 10 possible colors used, and no color repeated

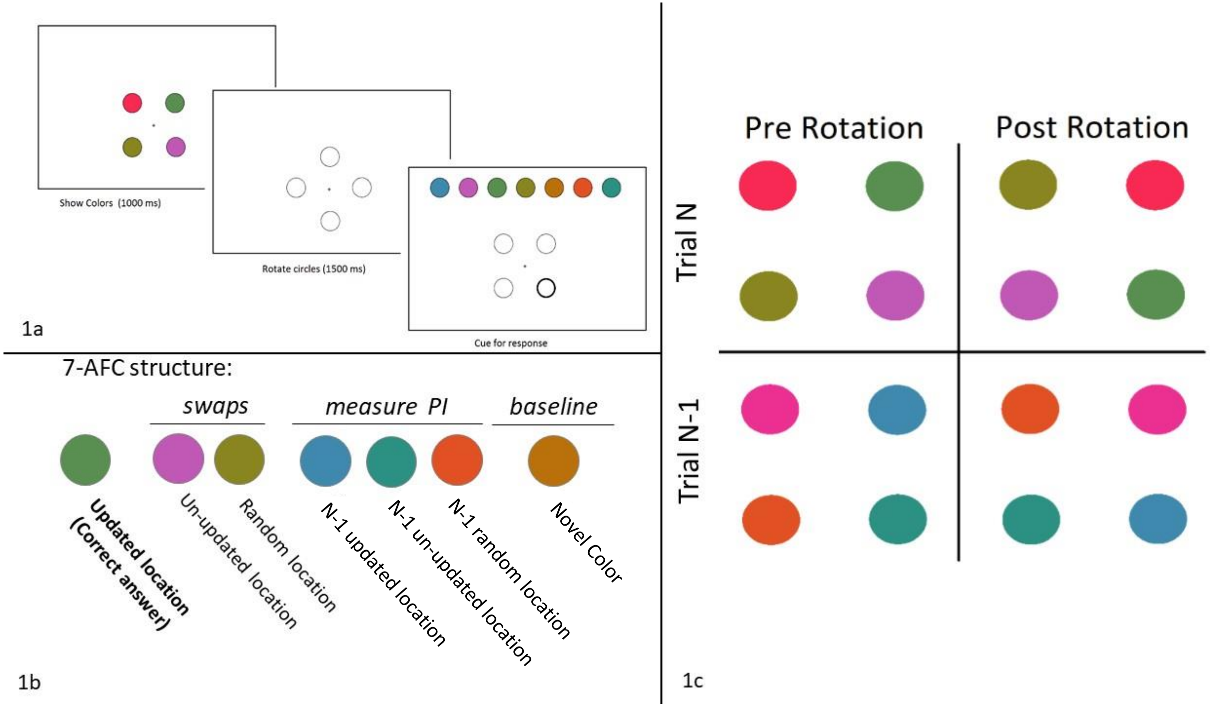

Figure 1: The top left panel shows a schematic illustration of the trial sequences used in experiment 1a. The right panel shows a sample stimulus from a current (N) and previous trial (N-1), before and after the 90-degree clockwise rotation. The bottom left panel shows descriptions of the 7 AFC task options. Analysis is done by comparing PI choices to the novel baseline

from the previous trial. The cued location never repeated but was randomly selected from the .three other possible locations.

Procedure A practice of 5 examples was used to familiarize subjects with the task. The practice stimuli and procedure were the same as the main experiment except the encoding time was increased for the first 3 trials to 2000 ms, and in all trials, the post-rotation colored disks would appear after the participants made their selections so they could see what they got correct. During the practice, a research assistant was present to answer any questions about the task. Immediately after the practice, participants started the main experiment, where they completed 275 trials, divided into 11 blocks of 25. The pace was self-directed; however, all subjects completed the task within an hour, including the practice. Participants were instructed to “not use words” when encoding or making answers, although no formalized verbal interference task was used.

Results

Figure 2 shows a summary of our results, including the averaged accuracy rates of each selection and the error distributions of the novel and PI characteristics. Subjects selected the correct answer 52.8% of the time on average, with the most common error being the current trial unrotated at 12.2% and a current trial random location at 8.2%. The novel and PI selections were all made less than 7% of the time.

Looking at the distribution of errors (meaning correct answers are excluded from analysis), the only PI condition that was significantly higher than the baseline novel foil was the rotated condition, with a P value of .02. While the other PI selections were also higher than the baseline, they were not so at significant levels. These significant PI effects were confined to the first half of the experiment. The T-test for the first half of trials shows a P value of 0.02 when comparing the number of N-1 un-updated

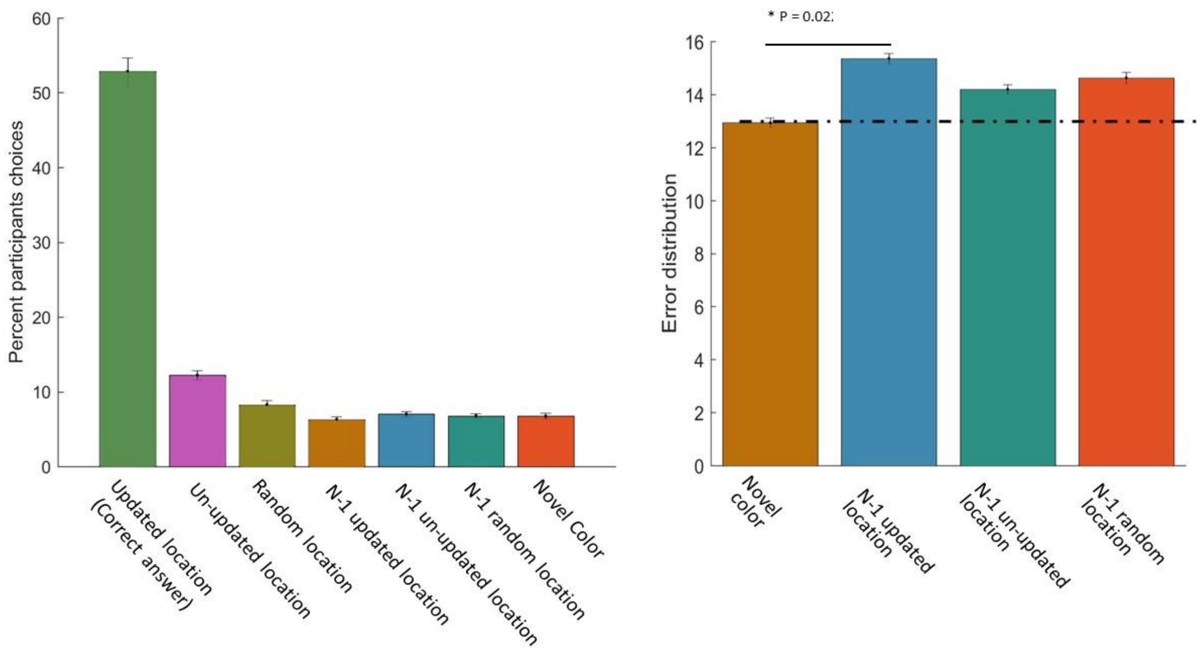

Figure 2: (left) Total distribution of answers from the exp 1a 8afc. (right). Distribution of errors form exp 1a, comparing the novel baseline to the PI conditions. The correct and current trial swap answers are not shown, as they are uninformative to PI rates. Only the N-1 updated location is significantly higher than baseline

selections to baseline, while the same comparison for the second half of trials has a P value of 0.2. However, because the difference between these two values is not significant, the difference could be noise driven.

Discussion

Experiment 1a replicates previous studies in finding that proactive interference for simple stimuli or objects arranged with spatial information is small. However, a surprising result is that this interference is location specific to the mentally updated location. The absence of PI in the N-1 un-updated and N-1 random selections suggests that LTM is minimally utilized in this task; however, the PI only found in the mentally updated location suggests that this PI is within working memory, Combined with there being no significant interference in the other PI conditions, which may be assisted by LTM, this implies that the influence of LTM on proactive interreference is exceedingly low in this testing condition and that this interference can manifest itself in working memory alone. It is possible that there is some proactive interference influencing the choices of the N-1 un-updated and N-1 random conditions, as a complete lack of LTM influence would be surprising; however, this PI is at levels too low to be statistically significant with this experiment. It should be noted that this is not conclusive, as the N-1 updated location isn't significantly different from the other PI conditions, even though it is the only one significantly higher than baseline. The N-1 updated location is significantly different from other N-1 conditions in the first half of the trials, requiring further investigation before making definitive conclusions.

It's unclear why the proactive interference was more pronounced in the first half of the trials; however, it could be due to fatigue not creating strong enough impressions in working memory to cause interference, or due to some kind of adaption to the task to reduce PI. Because performance is almost identical in the first and second halves, fatigue may be unlikely.

Experiment 1b

PI in real-world objects is not location specific

Studies have shown that real-world objects are remembered differently than simple stimuli, either due to differences in lower level processing or by different strategies being used to remember them (Brady, 2008). This has also been seen with real-world objects, where they have much higher rates of PI (Endress & Potter, 2013; Makovski, 2016). However, the stimulus presentation appears to have a very large effect on the degree of PI, as Endress & Potter found very large amounts when objects are shown sequentially at a single location, yet Makovski found much more moderate PI effects for stimuli presented at different spatial locations. We were interested in both studying the higher rates of PI and investigating the possible reasons for the large difference in PI, depending on stimuli presentation.

Method

Participants A new group of 74 subjects between 18 and 35 were selected using the UCDA SONA recruitment tool. However, data from the first 36 of these subjects had to be rejected due to a MATLAB coding error. All subjects were screened for normal color vision, using the Ishihara plates 2-8. All subjects received class credit for their participation.

Stimuli, task, and procedure Experiment 2 was identical to Experiment 1 with the following exceptions: Real-world objects from the data set used in "Visual long-term memory has a massive storage capacity for object details," (Brady et al, 2008) were used instead of color disks. None of these real-world objects was repeated throughout the task. The number of trails was also reduced to 10 blocks of 12 to decrease the length of the experiment to under 30 mins and so objects would not need to be repeated. Figure 3a shows a schematic illustration of the experiment 2 task and procedure.

Results

Participants selected the correct object 63.1% of the time, with current-trial unrotated and random selections at 12.6% and 9.52% respectively. The novel condition was selected 2.59% while PI measures for rotated, unrotated and random were 4.27%, 3.86% and 4.00% respectively. 32.0% of the error was from a form a PI selection, which equates to 11.2% when guesses are factored out. Figure 3 shows the novel and proactive interference selections as distributions of error. All categories of PI were significantly higher than the baseline, but each category was not significantly different from the other.

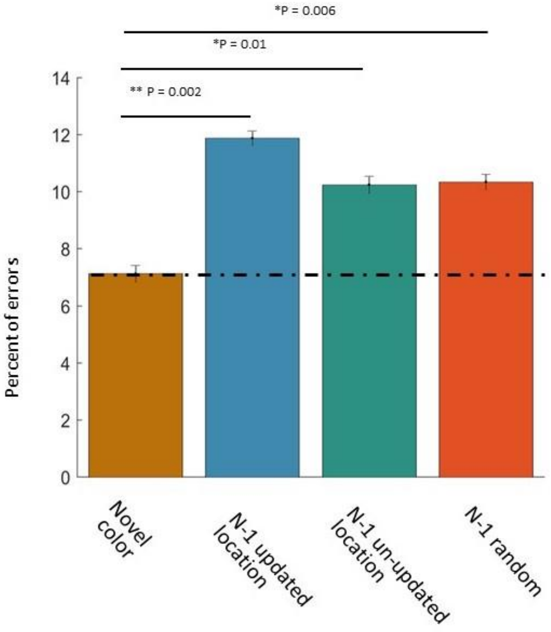

Figure 3. Showing the exp. 1b distribution of errors, comparing the PI conditions with baseline. All three PI conditions are significantly higher than baseline and are not significantly different from each other, suggesting that PI in real-world objects is not location specific

Discussion

While the number of total errors and PI errors is lower than the simple stimuli, the distribution of errors showed much more PI than with simple stimuli. Interestingly, the PI was not location specific, meaning an object seen in the N-1 trial is equally likely to be picked, regardless of it being the correct object or not from that trial, rotated or unrotated. This suggests that subjects incorporate VLTM as a strategy for real-world objects, but this VLTM doesn't carry significant spatial information. However, with simple stimuli, perhaps VLTM is less helpful. This also explains the higher rates of PI in the Endress Porter study, which showed real-world objects serially in the same location, as VLTM was highly utilized when VWM didn't have any spatial information to help differentiate objects or the number of stimuli was beyond capacity.

Experiment 1c

Controlling for the effects of mental rotation

We wanted to confirm that the results of non-location specific PI in real-world objects from experiment 1b was not a result of the mental rotation procedure. It may have been possible that the rotation somehow broke the spatial representation held in VLTM.

Method

Participants – A new group of 15 volunteers were selected from UCSD's SONA, under the same conditions as the previous experiments. One subject was excluded due to their answers being more than 2 standard deviations from the mean.

Stimuli, task, and procedure – The procedure for half of the trials for Experiment 1c were exactly the same as 1b. For the other half of trials, the only difference was that there was no mental rotation phase. In its place, when the objects disappeared, the box outline around the objects remained stationary for the same length of time as the rotationary period, to leave the same inter-stimulus interval between conditions. The rotated and unrotated tasks were given in two separate and sequential testing phases, with half the subjects performing the rotated trials first and the other half the non-rotated trials. A practice run and instructions were given before each testing phase to ensure the participants understood the task.

Results

Figure 4 summarizes the results. The rotation portion of this experiment replicated the results from 1b with almost identical data; however, due a smaller subject sample size, only one of the PI measures, the N-1 updated location, was significantly greater than the Novel color. In the non-rotated experiment, participants did extremely well, with an average of 86.0% correct answers, both current trial random answers averaging to 4.17%, the N-1 correct at 1.64%, both N-1 random answers averaging to 1.50% and the novel at 1.02%. Figure 1b shows the distribution of errors, with the novel color making up 6.64% of the error, the N-1 correct making 11.8% of the error, and the two N-1 random averaging to 9.46% of the error. Only the N-1 correct answer was significantly different from the novel.

Discussion

Although a small sample size limited our statistical power, this experiment effectively replicated our previous findings and showed that the rotation had no visible effect on the PI distribution or the distribution of errors. The largest surprise is how similar and consistent these results are. While only the N-1 correct answers are significantly above baseline, we believe that this is due to the small sample size, as the distributions match so closely to the results of 1b. None of the N-1 conditions are significantly different from each other, suggesting that they are all non-location specific.

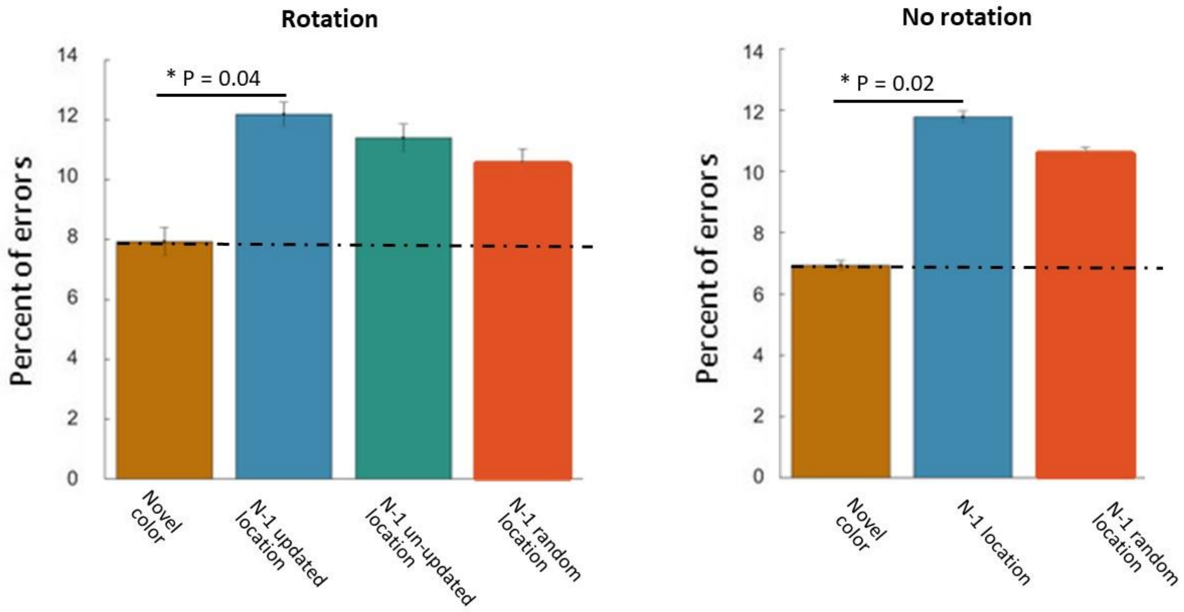
Rotation

No rotation

Figure 4: (left) The distribution of errors for the mental rotation block in exp 1c, focusing on the novel baseline and N-1 trials. All N-1 trials are significantly higher than baseline, replicating 1b. (right) The distribution of errors for the unrotated condition of exp 1c, truncated to only show the novel and N-1 conditions. These results show no effect from mental rotation or no rotation on PI.

Experiment 1d

Simple Stimuli with modified methods

After the previous experiments, we asked if increasing the complexity of the mental manipulation for simple stimuli would change the rate and type of PI. Specifically, if we removed the competition between updated and un-updated location representations, would we still see PI. This experiment has a random clockwise or counterclockwise rotation at 180 degrees. We reduced the presented objects to 3 because the random rotation to such a degree was too difficult at 4 objects, and it removed overlapping positions of the pre and post rotated objects

Method

Participants – 56 students between 18 and 35 were recruited from UCSD's SONA, similar to the previous experiments. No participants were excluded

Stimuli, task, and procedure Experiment 1d followed a similar procedure to Experiment 1a with the following exceptions: Only 3 stimuli were displayed in an equilateral triangular array; the mental rotation randomly occurs clockwise and counterclockwise and at 180 degrees; and the AFC task is reduced to 5 because with no pre/post rotated location overlap, there is no unrotated condition. A schematic is provided in figure 5.

Results

We found the following average selection distribution: correct – 61.98%, N random – 14.97%, N-1 location – 7.87%, N-1 random – 8.18%, and novel – 7.00%. Figure 5 shows this represented as the distribution of errors, comparing the PI conditions with novel condition. Significant PI was found for both PI conditions; however, the PI conditions are not significantly different from each other.

Discussion

This shows non-location specific PI with the simple stimuli, which is different from the location specific PI we found in experiment 1a. One possible explanation is that the more complex rotation reduced the strength of the location representation of the previous trial, causing VLTM traces to influence the guess rate. The increased difficulty of the task may utilize more VWM resources, encouraging a more thorough memory discharge of the previous trial's spatial information.

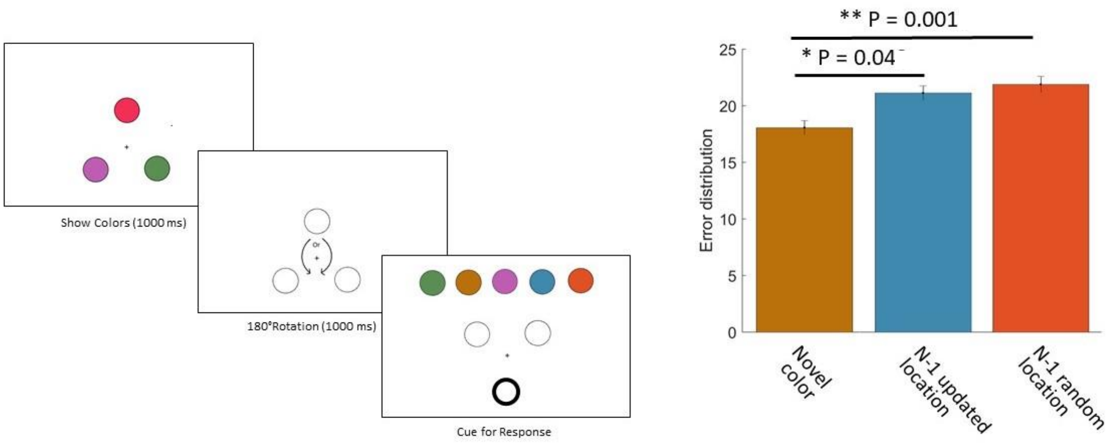

Figure 5: (left) a visual schematic for the stimuli of Exp 1d. (right) The error distribution from exp. 1d, showing only the comparison between the novel and N-1 conditions. Both conditions are higher than the baseline, suggesting non-location specific interference.

Experiment 2

Feature binding and Proactive Interference

While experiment 1 shows that real-world objects show different types of proactive interference from simple stimuli, it doesn't reveal what aspects of the real-world objects cause this difference. Real-world objects have semantic content, differing degrees of salience, and multiple complex features, whereas simple stimuli generally do not, and it is unclear which of these might drive the differences in PI. In experiment 2, we aim to look at how PI interacts with the multiple feature aspect, specifically asking if PI is ascribed to individual features or to an object as a whole.

Experiment 2a

Is proactive interference tied to whole objects or individual features?

The challenge in this experiment is to separate memory of an object from the memories of its component parts. While real-world objects have multiple features which can be separated, such as color, location, size, orientation, etc., there are too many variations in the number and depth of features. Real objects are more or less salient, trigger multiple brain processes, and it is more difficult to separate traces from episodic memory formed from subvocalization. Therefore we will use semi-simple stimuli which are limited to three characteristics: location, shape, and color. We will show a series of objects, give one characteristic as a cue, and ask participants to select the other two. From these answers, we can see if characteristics from objects on the previous trial are selected more than a novel foil. Because this experiment isn't focused on differentiating VLTM and VWM, the mental rotation aspect has been removed.

Methods

Participants A new group of 58 subjects between 18 and 35 was selected using the UCSD SONA recruitment tool, however, the data was lost for one subject due to one subject's data overwriting another's. Subjects were screened for normal color vision, as in previous experiments. Compensation of one class credit was awarded for completion of the study.

Materials The same materials were used in this experiment as with the previous studies in this paper.

Stimuli and Task Three objects were simultaneously displayed to the subjects for 750 ms, upon a neutral grey background. These objects each have three characteristics: shape, color, and location, and each characteristic has one of 8 possible options. Each of the possible locations are drawn with 8 thinly bordered black squares with a 4 cm length arranged in a circular pattern with a radius of 8.8cm, spaced out every Π/4 around the circle, starting at 0. The 8 possible colors were black [0,0,0]; white [255,255,255], red [254,0,0]; green [0,254,0], blue [0,0,254], yellow [255,255,0], turquoise [1, 255, 255]; and purple [201,0,200]. The possible objects included a circle, arrow, exclamation point, 6-pointed star, hourglass, triangle, diamond, and equilateral cross.

After a 500ms neutral gray masking screen, the cue screen would display all eight color options on the right-hand side in a vertical 2x4 array, all the shapes would be on the left in a 2x4 array, and the locations would be marked with the same thin black squares used in the stimulus screen. Both the shapes and colors were displayed in 4 cm sections, with 1.25 cm spacing. The cued characteristic, either a color, shape, or location from the current trial's objects, would be highlighted with a thick box, and the subject would use the mouse to click on the other two characteristics that match the object whose given characteristic is shown. This created two simultaneous 8 AFC tasks. Subjects could only select one of each characteristic per trial and could not change their original answer. After the second characteristic was chosen, a prompt would ask the subject to press “any key” to move on to the next trial.

None of the characteristics on a single trial were repeated (there were no two objects that were red, or circles) and only one characteristic was repeated from the current trial and the trial before. This repeated characteristic was used as the cue for each trial, except for the first trial of each block, which had a random cue. The cue was randomly chosen to be either a color, shape, or location cue, allowing for repeated cues of the same type.

Because all 8 options of each characteristic were available, one was the correct answer for the current trial, two were from one of the objects seen on the current trial, one matched the current trial’s cued characteristic with the object from the previous trial, another two were from the previous trial, and two were novel options not seen in the current or immediately preceding trial. The order of the color and shape characteristics within their arrays was randomized for each trial.

Procedure Subjects first performed a practice to familiarize themselves with the task, under the guidance of a researcher. The practice included 5 trials of randomized stimuli with no characteristics overlapping from the first to second trial. The cues included at least one of each characteristic. After making selections for each practice trial, the original stimuli were superimposed on the screen so subjects could compare it to their answers. Once the practice was completed, they completed 350 trials, divided into 10 blocks of 35. The pace was self-directed, with all subjects completing the task in under 30 mins. Participants were instructed not to use words but were given no verifiable verbal interference tasks. See figure 7 (left).

Results

Preliminary results are shown in figure 6, although there are many more data points and analysis that are not yet completed. Figure 6 shows the distribution of choices, separating bound selections from unbound. A selection is considered bound if each selected choice was from the same object, regardless of its correspondence to the cue. Because there are only single N and N-1 correct choices, but pairs of Novel and N and N-1 random, these pairs were averaged together for analysis. In all cases, when selection characteristics were bound to

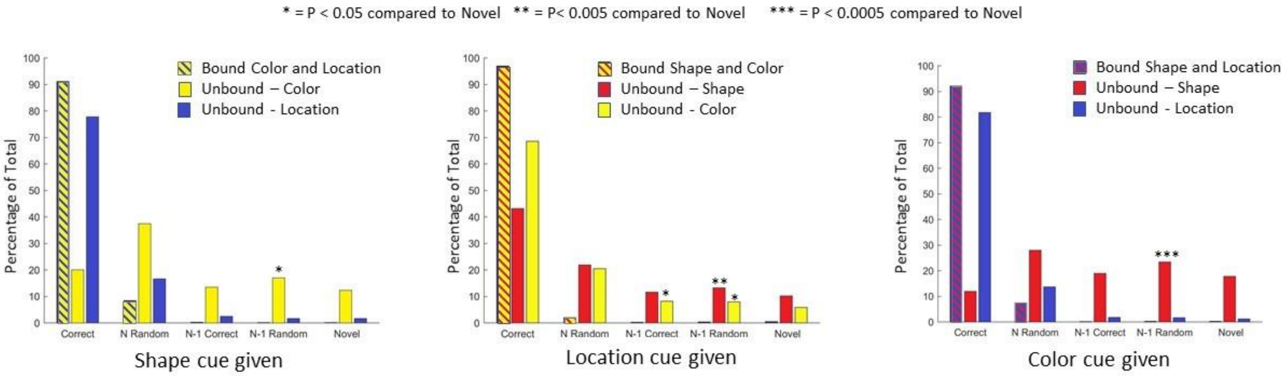

Figure 6: All from experiment 2a, the graph shows the distribution of choices for a shape cue when the selections from the two 8 AFC tasks are from the same object (bound), or the answers for each of the unbound characteristics. The left graph is when Shape is a cue, the middle for a location cue and the right for a color cue. Bound characteristics are very accurate, as is unbound location, while unbound shape or color accuracy is either right above chance, at chance or even bellow chance

one presented object, the accuracy was extremely high, producing a ceiling effect.

Looking at unbound selections, location was still very accurate, regardless of its binding, either to shape or color. However, unbound location selections were much more accurate for a shape cue than a color cue. When a color cue was given, if the shape wasn't bound to location, it's choices were well distributed. Answer selections from this condition were 66% less likely to be correct than the baseline (p = 0.0003). Both unbound color and shape correct responses were significantly higher than other responses when location was a cue.

Ceiling effects prevented the analysis of PI for bound and location characteristics, since there were so few errors, PI influenced or otherwise. This also reduced the statistical power for unbound color and shape analysis. Some PI over baseline was found in unbound color when location was a cue, for both the N-1 correct (P = 0.05) and N-1 random (P = 0.02); however, these weren't statistically different from each other. Unbound color also had PI which shape as a cue for N-1 Random (P=0.005), and this was significantly different from the N-1 Correct as well (P = .05). Unbound shape with a location cue had a modest significant increase over baseline for only the N-1 Random (P= 0.05). The largest PI effect over the baseline was found for the N-1 random unbound shape with a color cue (P=0.0002) and this was also higher than the non-significantly different N-1 Correct (P= 0.02).

Discussion

While this experiment didn't show strong PI effects, most likely due to the ceiling effect, there is lots of useful object binding data that will be further analyzed. Memory of spatial characteristics proved to be the highest indicator of a correct answer, as correct color and shape responses were heavily dependent on them being bound to the location characteristic. In other words, the location was known most of the time, so if the color or shape was bound to the location, it got a "free" ride to the correct answer. However, if color or shape was not bound to the correct location, the answer selection was almost as good as chance, or was worse if the color or location information was mistakenly bounded to an incorrect location. It

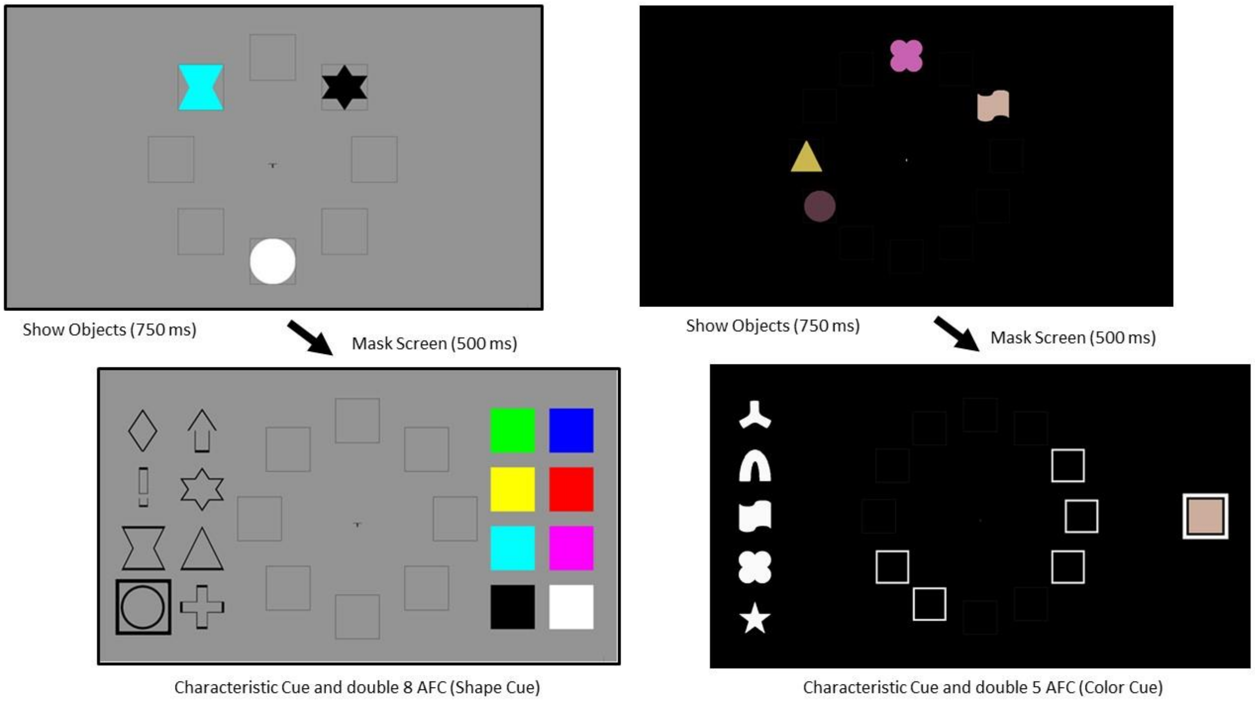

Figure 7: (left) – a schematic for the experiment in 2a, done in 10 blocks of 35 trials each. (right) – a schematic for the experiment in 2b, done in 10 blocks of 25 each.

is unusual that when the shape characteristic is not bound with location, it has a correct answer rate below chance for both the N and N-1 trials. Some of this may be explained by it being bound to the wrong location, thus pulling it away from the correct answer, or there may be a mental strategy in play that works to minimize PI errors. Binding characteristics to location but not to shape or color seems to be the primary strategy used.

Experiment 2b

Increasing the difficulty with more objects

Methods

Participants A new group of 62 UCSD undergraduates, aged between 18 and 35, were recruited from SONA and were screened for normal colored vision. They all received class credit for their participation.

Materials, Stimuli, Task and Procedure This experiment is a modification of experiment 2a, in which four objects were used as stimuli, with each object having one of twelve distinct color, shape, and location characteristics, all over a black background. Each of the possible locations was outlined in a thin white lined box with a length and width of 3 cm and was positioned in a circular pattern with a radius of 9 cm and locations at every Π/6 starting at 0. The colors were generated from "I Want Hue" online color generator (Jacomy, 2016) and eight of the shapes were taken from another object binding experiment from "Feature binding and attention in working memory: A resolution of previous contradictory findings" (Allen et al., 2012), as well as 4 other shapes. A schematic of the procedure is in figure 7 and all the stimuli shapes and colors are presented in figure 8.

Results

The main results are shown in chart 1, with breakdown of the distribution by feature binding in figures 9 and 10. Color errors when provided a location cue and Shape errors when provided a color cue produced the most PI, none of which was binding specific. This means that if someone is pulled by a foil from the previous trial, it doesn't matter what that feature was bound to; only that it was seen on the previous trial. Overall, PI was very slight for most conditions but always trended higher than the novel baseline, even if not at statistically significant levels. Separating answers by their feature binding shows that unbound color with a shape

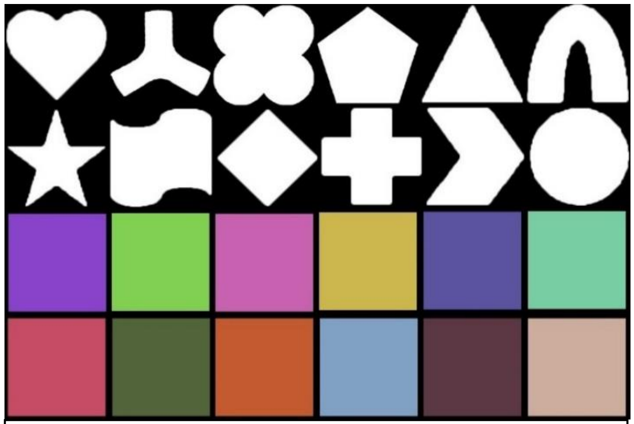

Figure 8. The stimuli shapes and colors used for experiment 2b. The RGB values are as follows (from left to right, top down): [137,67,203]; [130,208,81];[201,97,177]; [204,183, 79]; [91, 82, 160]; [121,205,163]; [198, 76, 102]; [82, 100, 57]; [196, 91, 48]; [128, 161, 195]; [91, 56, 68]; and [205, 173, 157].

	Correct	N Rand	N-1 Correct	N-1 Rand	Novel
Color Cue: Location	62	20	6	6	6
Color Cue: Shape	48	20	11	11	10
Shape Cue: Location	66	17	6	6	5
Shape Cue: Color	44	21	12	13	10
Location Cue: Shape:	59	14	10	9	8
Location Cue: Color:	51	19	10	11	9

Chart 1 This gives the percentages of responses, separated by the cue characteristic and characteristic selection and organized by answer type. A cue was either a location, shape or color and subjects had to pick the corresponding other two characteristics that matched the object with the cued characteristic. The correct answer matched the cue, while the N rand matched a different object from the current trial. N-1 correct matched the cued characteristic with an object from the previous trial, while N-1 was for a characteristic from the previous trial not associated with the cue. The novel characteristic is one not seen on this or the previous trials.

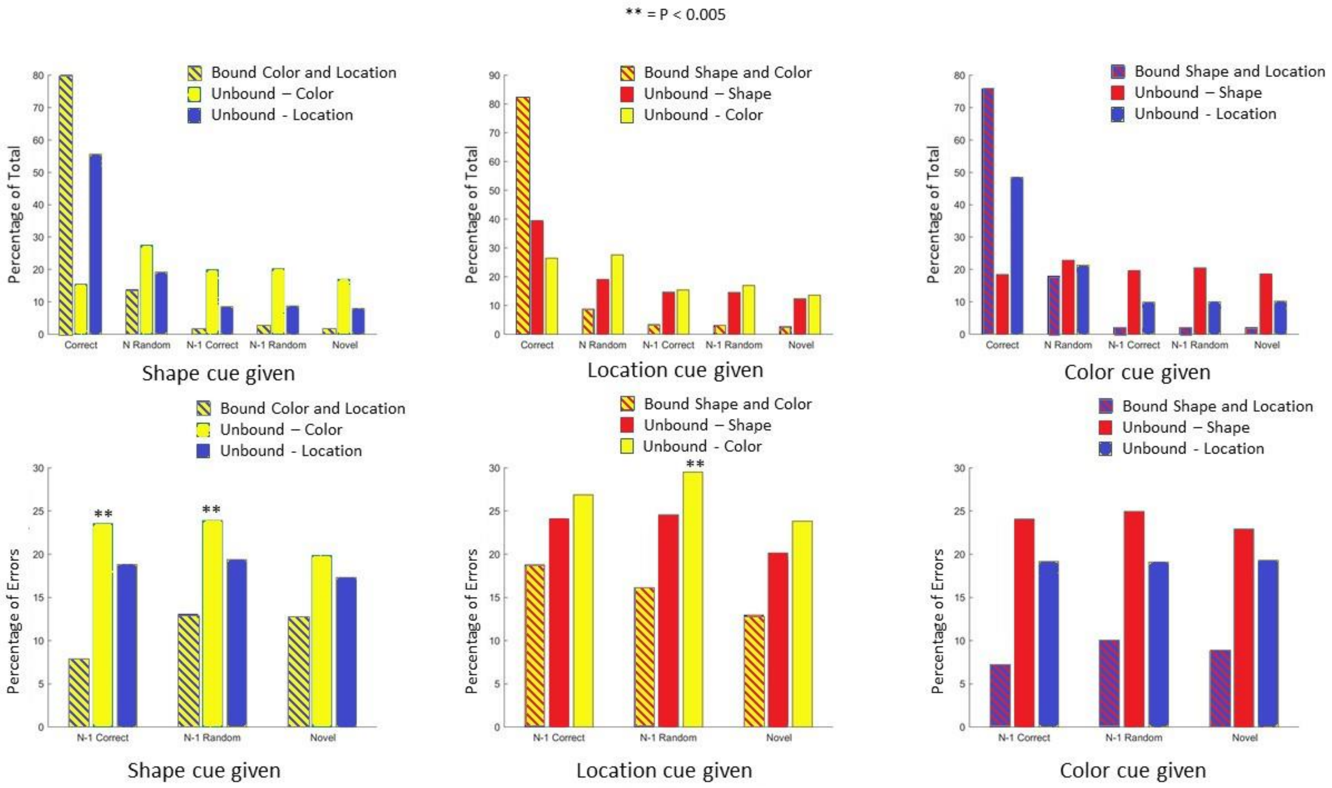

Figure 9: The top half shows the distribution of answer selections based on if the cue was given as a shape, location or color. One data set is of answers that are bound, meaning both selected characteristics are associated with the same object, and the two unbound data sets, where the answers referred to other answer sets. The bound novel condition is when both characteristic selections were not present in the current or previous trial. The bottom half shows a reformulation of the same data, focusing on the distribution of errors (disregarding the correct answers), focusing only on the N-1 and novel conditions. PI is marked with **, showing it is significantly higher than the novel condition

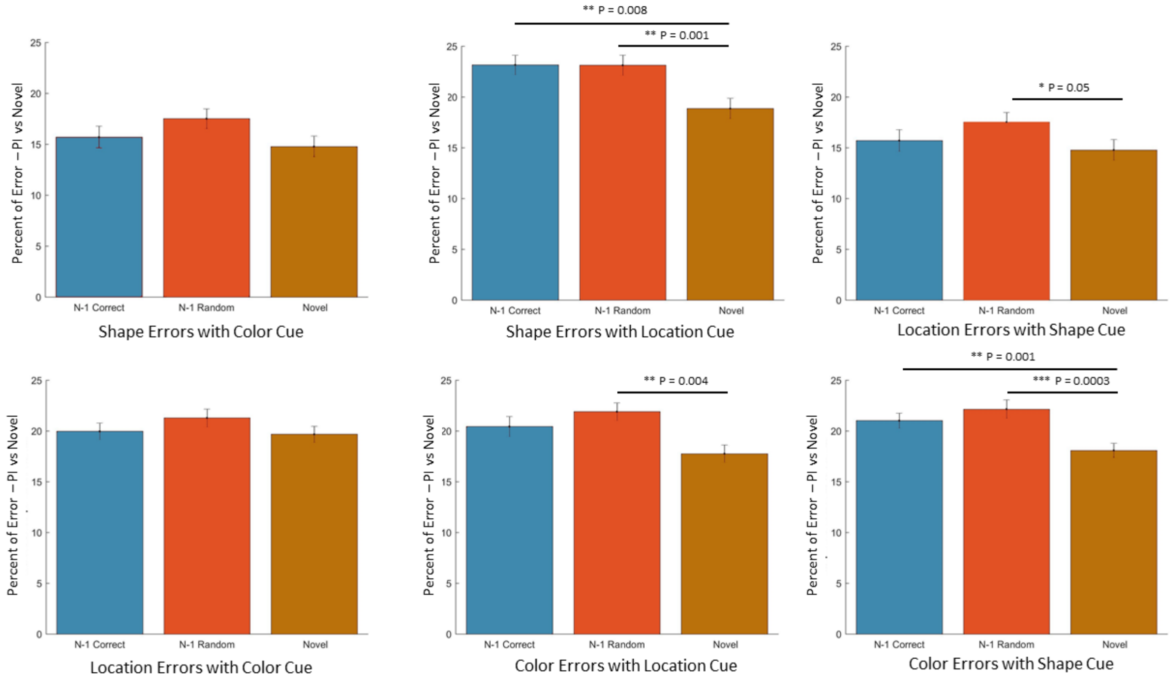

Figure 10. This figure shows the distribution of errors, comparing the N-1 and novel conditions. They are organized by the answer and cue types (i.e. for a location cue, there are shape and color responses) and are agnostic to feature binding. Small amounts of PI was found, however PI was not universally found for any particular characteristic. Further analysis is needed to determine if there are any clear patterns in the distributions of significant PI

characteristic had the highest PI, equally for the N-1 correct and N-1 random as well as unbound color for a location cue, but only with the N-1 random (although this is not significantly different from the N-1 correct).

The feature binding results were similar to the previous study, in that color or shape were equally distributed when not bound to location information. Bound characteristics were always very accurately to the cued characteristic (correct answers) and unbound location was still accurate as well.

Discussion

Much further analysis of this data is required; however, minimal PI results make this difficult.

Like experiment 2a, subjects were much more accurate than expected, giving less chances for proactive interference. A follow-up experiment using real-world objects, looking at binding between location, color, size, and orientation, all using a continuous report instead of AFC task might show stronger signs of PI. However, the feature binding aspects are interesting, further showing the importance of location information for successful use of VWM.

General Discussion & Conclusion

While difficult to detect and not present in all conditions, PI shows promise in helping differentiate some aspects of memory in short-term memory tasks that are otherwise very difficult to separate. Our results on object binding are inconclusive and we are not yet confident enough to fully state that PI is found in VWM and isn't only derived from influences of VLTM, although the data suggests so. The biggest takeaway from this study is that the PI that is found in real-world objects is non-location specific when stimuli are presented in a location-specific manner. This is puzzling that the PI is not location specific when the amount of location information is critical to the amount of PI. Our conjecture is that when location is not a good source of information to differentiate objects, VLTM is employed to help out, thus increasing the PI. The importance of location information for VWM is demonstrated in our feature binding experiments 2a and 2b. Other studies in our lab, particularly Dr. Schurgin's forthcoming paper "Episodic Memory Replaces Active Maintenance in Working Memory When Available" gives strong electrophysiological evidence that VLTM can be and is employed to help VWM. Research by Brady Lab undergraduate researcher Zeljana Babic further advances the concept the strategy in how VWM is used plays an important role, more so than capacity.

It should be noted that our results may not be generalizable to the average population, due to the selection bias from using UCSD students. These subjects are W.E.I.R.D. and tend to have above average IQ's, which is highly correlated to working memory.

Future Research

There are many directions for future study in PI in VWM. It is unclear if the PI difference between real-world objects and simple stimuli is due to multiple complex features, increased salience, semantic content, episodic memory triggers, or other characteristics. We can use different stimuli to try to differentiate these factors by using fractals or greebles that have complex features but no semantic content, or we can use the same real-world object as stimuli multiple times to see if the PI changes over time as the utility of VLTM of that object diminishes. Eye tracking can be used to see if PI is a lure for attention, even if it isn't selected, and we can combine the feature binding experiment with mental rotation to see if object characteristics are easily updated when the location information is updated. Comparing PI from different stimuli presentations, such as location specific sequential, simultaneous, and non-location specific sequential may show interesting results, as would allowing for a continuous report to see if PI has "gist" influences that are not seen in an AFC or change detection task.

Acknowledgements

I would like to sincerely thank everyone in the UCSD Vision and Memory Lab, particularly Dr. Tim Brady and his Post Doc, Mark Schurgin, who both dedicate an extraordinary amount of time and effort into the education of their undergraduate researchers. I was also immensely helped by my department advisors Dr. Andrea Chiba and her graduate student Eric Leonardis. UCSD's cognitive science dept. chair, Dr. Marta Kutas, was instrumental in helping my research presentation skills and offering guidance for academic writing. I give a very deserved thank you to our undergraduate advisor Katy Schoeneweis, whose positivity, encouragement, and helpfulness in all academic matters goes beyond expectations.

Most important of all, I'm thankful for the love and encouragement from my fiancée, Adrienne, without whom I wouldn't have gone back to school or pursued research.

Reverences

Allen, R. J., Hitch, G. J., Mate, J., & Baddeley, A. D. (2012). Feature binding and attention in working memory: A resolution of previous contradictory findings. Quarterly Journal of Experimental Psychology, 65(12), 2369-2383. doi:10.1080/17470218.2012.687384

Alloway TP, Alloway RG (2010) Investigating the predictive roles of working memory and IQ in academic attainment. J Exp Child Psychol 106(1):20–29. 4.

Baddeley A (2012) Working memory: Theories, models, and controversies. Annu Rev Psychol 63:1–29.

Brady, T. F., Konkle, T., Alvarez, G. A. and Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. Proceedings of the National Academy of Sciences, USA, 105 (38), 14325-14329.

Brady, T. F., Konkle, T, and Alvarez, G.A. (2011). A review of visual memory capacity: Beyond individual items and towards structured representations. Journal of Vision, 11(5):4, 1-34

Brainard, D. H. (1997). The Psychophysics Toolbox. Spatial Vision, 10, 433– 436. http://dx.doi.org/10.1163/156856897X00357

Cowan N (2001) The magical number 4 in short-term memory: A reconsideration of mental storage capacity. Behav Brain Sci 24(1):87–114, discussion 114–185.

Endress, A. D., & Potter, M. C. (2013, August 12). Large Capacity Temporary Visual Memory. Journal of Experimental Psychology: General. Advance online publication. doi: 10.1037/a0033934

Engle, R. W. (2002). Working Memory Capacity as Executive Attention. Current Directions in Psychological Science, 11(1), 19-23. doi:10.1111/1467-8721.00160

Fukuda K, Vogel E, Mayr U, Awh E (2010) Quantity, not quality: The relationship between fluid intelligence and working memory capacity. Psychon Bull Rev 17(5):673–679

Jacomy, M. (2016). I want hue: Colors for data scientists. Sciences-Po Medialab. Retrieved June 12, 2018, from http://tools.medialab.sciences-po.fr/iwanthue/

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. Nature Neuroscience, 17(3), 347–356.

Makovski, T., & Jiang, Y. V. (2008). Proactive interference from items previously stored in visual working memory. Memory & Cognition, 36(1), 43-52. doi:10.3758/mc.36.1.43

Makovski, T. (2016). Does proactive interference play a significant role in visual working memory tasks? Journal of Experimental Psychology: Learning, Memory, and Cognition, 42(10), 1664-1672. doi:10.1037/xlm0000262

McDowd, J. M., OSEAS-KREGER, D. M., & FILION, D. L. (1995). Inhibitory processes in cognition and aging. In F. N. Dempsteree C. 1. Brainerd (Eds.), Interference and inhibition in cognition (pp. 363- 401). San Diego: Academic Press

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. Spatial Vision, 10, 437– 442. http:// dx.doi.org/10.1163/156856897X00366

Standing, L. (1973). Learning 10000 pictures. Quarterly Journal of Experimental Psychology, 25(2), 207–222.

Underwood, B. J. (1957). Interference and forgetting. Psychological review, 64(1), 49.