Spring 2018

The Guidance of Visual Attention Through Learned Feature Probabilities

Eli Schmidt
University of Iowa

THE GUIDANCE OF VISUAL ATTENTION THROUGH LEARNED FEATURE PROBABILITIES

by

Eli Schmidt

A thesis submitted in partial fulfillment of the requirements
for graduation with Honors in the Psychology

Shaun Vecera
Thesis Mentor

Spring 2018

All requirements for graduation with Honors in the
Psychology have been completed.

Toby Mordkoff
Psychology Honors Advisor

THE GUIDANCE OF VISUAL ATTENTION THROUGH LEARNED FEATURE PROBABILITIES

by

Eli Schmidt

This thesis has been reviewed and approved.

____

Shaun P. Vecera
Thesis Advisor

____

Brad T. Stilwell
Second Reader

Spring, 2018

Abstract

Visual attention can be influenced through statistical learning of information in the environment, and over time, extracted visual patterns relevant to the current task can be used to guide attention. Specifically, within a visual search paradigm, statistical learning of feature information (e.g. color) can be implicitly extracted to make attentional guidance more efficient. In addition, we know that the system can use explicitly provided information, such as the contents of visual working memory (VWM) to bias attention. We hypothesized that feature-based statistical learning might interact with the contents of VWM. In the present study, participants searched through displays containing target features (i.e. shape) that were more likely to contain the target of the search than another feature. In order to examine the interaction between VWM and previously learned attentional biases, in Experiment 2, we trained participants on the same implicit learning task as Experiment 1. We introduced a color to be remembered that either matched the color of the target, of the distractor, or was a color that was not present in the search display. We found that when a color was stored in VWM and present in the search display, all attentional biases based on the previously learned statistics disappeared. Therefore, we hypothesize that VWM dominates attentional guidance. In other words, feature-based statistical learning disappears when there is a concurrent strong VWM bias.

Humans live and operate in a complex visual environment that contains too much information at any given moment for one to perceive and respond to. Therefore, in order to successfully interact with the world around us, we need to prioritize the parts of the external environment that are most important. One mechanism the visual system uses to prioritize the incoming information is to focus attention towards inputs from the environment extracted through statistical learning.

Implicit statistical learning refers to an unconscious process that extracts repeated patterns or regularities in the environment in order to increase the efficiency of the allocation of attentional resources (Goujon, Didierjean, & Thorpe, 2015). Previous research on statistical learning through visual attention has focused on location-based cueing. That is, the visual system can implicitly learn associations between target and distractor locations in space implicitly. More specifically, Chun and Jiang (1998) had participants search for a target among distractor letters. Some displays contained repeated configurations of distractor locations that were exactly predictive of that target location. This association between repeated distractor arrays and target locations led participants to become more efficient in their search (Chun & Jiang, 1998). This bias to preferentially attend to a relevant location, in order to more efficiently locate the target, is quickly learned and persists for at least a week (Jiang, Swallow, Rosenbaum, & Herzig, 2012).

Location is not the only visual quality that can be implicitly learned to aid in target selection. The visual system can also prioritize visual features such as color through statistical learning. When the information is provided explicitly, the visual system has repeatedly been shown to prioritize task-relevant colors (Wolfe, 1994). However, recently Sha, Remington, &

Jiang (2017) demonstrated feature (i.e. color) prioritization through implicit statistical regularities. Participants were tasked with searching through a display where the target was a color 75% of the time and another color 25% of the time. The researchers found that participants were faster at locating the target when the target was the more frequently presented color than when it was the less frequent color. This suggests that you can implicitly learn to preferentially attend to not just a location, but also to a feature (e.g. color) in order to increase target selectivity.

In addition to prioritizing features for target selection, the visual attentional system can increase search efficiency by learning to ignore irrelevant distractors (Stilwell, Bahle, Vecera, in prep; Vatterott & Vecera, 2012). When the visual system is presented with reliable distractor regularities, the visual attentional system can extract these feature regularities and use the information to ignore distractors. Taken together with facilitation of target selection, implicitly learned distractor rejection suggests that the visual system extracts statistics from the environment and implicitly uses those statistics to guide attention more efficiently. We know the visual attentional system is capable of extracting statistical information to guide attention and these effects persist, but we do not know how these effects interact with explicit sources of information or items held in visual working memory.

Visual working memory (VWM) is a limited-capacity system that hold visual representations of the visual form of relevant objects over time (Stigchel & Hollingworth, 2018). In other words, to guide attention, the visual system must maintain some sort of active representation of the relevant visual features. For example, if you are looking for your phone, you may recall an image of your phone, hold that image in VWM and use it as a template to guide your attention to different items within your environment. However, any items that match that template would be flagged as targets during search and you would inspect them more carefully (e.g. perhaps a similar looking phone that belongs to your colleague). Thus, the visual attentional system is strongly guided by the contents of VWM but, as discussed previously, the visual system is also capable of guiding attention through learned statistical regularities. A question therefore remains: how do these sources of information interact? More specifically, does established statistical learning persist even with strong VWM guidance? Or, does VWM completely drive attentional guidance irrespective of previous learning?

To test these competing hypotheses, we first needed to replicate previous findings that a relevant feature dimension can be implicitly learned in order to guide attention. Rather than using color as the relevant dimension, however, we instead used shape. This was necessary in order to test the interaction of long-term learning with the contents of VWM (discussed in more detail below, Experiment, 2). Thus, in Experiment 1, participants searched through displays containing target features (i.e. shape) that were more likely to contain the target of the search than another feature. In order to examine the interaction between VWM and previously learned attentional biases, in Experiment 2, we trained participants on the same implicit learning task as Experiment 1. We introduced a color to be remembered that either matched the color of the target, of the distractor, or was a color that was not present in the search display.

To preview the results, in Experiment 1, participants learned to preferentially attend to the shape that was more likely to contain the target but that learning was quickly extinguished.

In Experiment 2, the VWM load extinguished the bias learned during the training session even when the probabilities remained.

Experiment 1: Learning a High Probability Shape

Method

The experiment consisted of a short practice, a training phase and a testing phase. The practice phase was to allow participants to become accustomed to the task. The training phase attempted to implicitly teach the participants that a certain shaped stimulus was more likely to contain the target. Once the association was learned, the testing phase examined whether the learned bias to attend to the relevant shape remained when preferentially attending to that shape was no longer beneficial to the task.

Participants

Participants in both experiments were students from the University of Iowa. They all had normal or correct-to-normal vision. Participants completed the experiments for course credit.

There were 25 participants in Experiment 1, 15 females and 10 males. Their mean age was 20 years old.

Apparatus

Participants completed the experiment individually in a dimly lit room. Stimuli were displayed on a 21-inch computer monitor placed 1.5 feet away from the seated participant with a normal keyboard placed comfortable in front of them for collecting responses. All participants were instructed to place their head again a forehead rest for the duration of the experiment. Each array was presented on a black background and had 2 stimuli displayed randomly around a cross at the center.

Procedure and Design

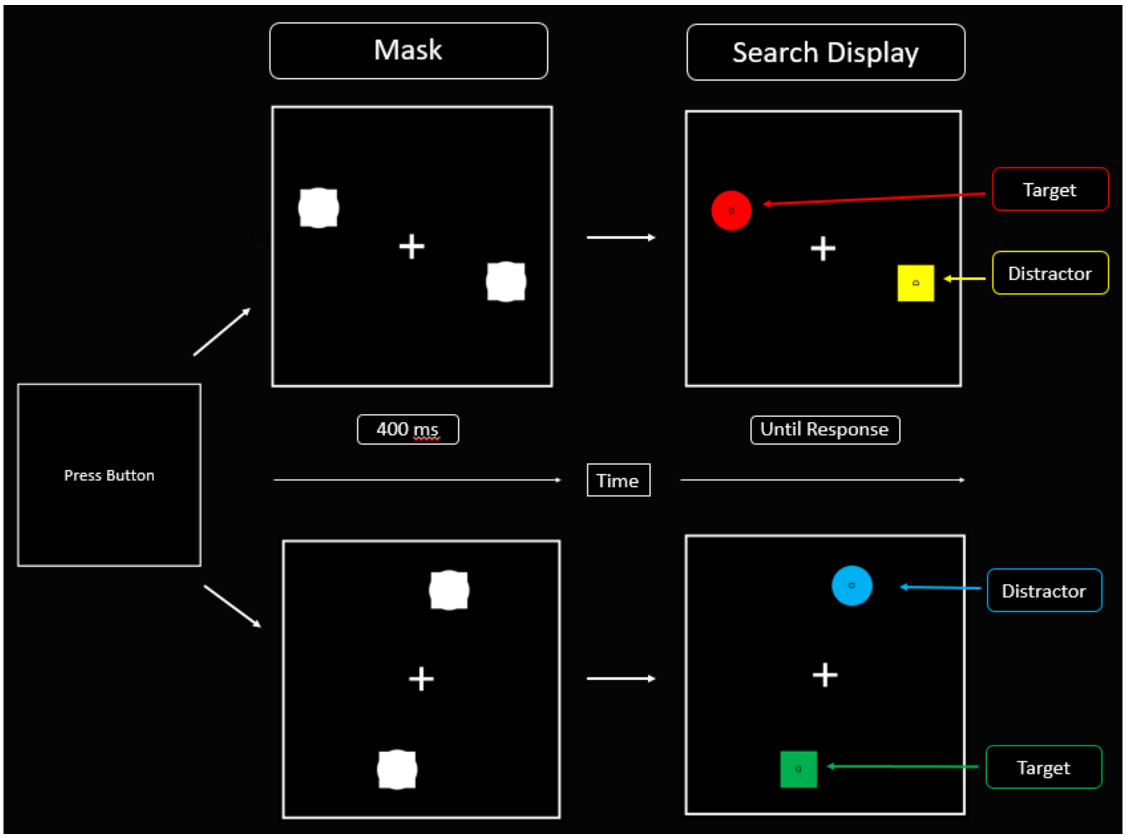

Figure 1. Sequence of events in a trial for Experiment 1. The task for Experiment 2 was identical, except for an added VWM task

Participants completed 10 trials of practice, 192 trials of training and 192 trials of testing. After the practice phase, participants continued through the training and testing phases without further instructions. The beginning of the testing phase was continuous with the end of the training phase and participants were not aware of the adjustment.

The task was to find a normal or a backwards "D" that was embedded in either a circle or a square (Figure 1). The shape that did not contain the target 'D' contained a distractor 'D' that was facing either facing upwards or downwards. The target was equally likely to appear in any colored shape making color an irrelevant feature for the entirety of Experiment 1. Before each trial, a mask appeared for 400 ms in the locations that the stimuli would appear. This mask was white in color and made of the circle and square overlaid on each other. The mask was necessary in order to minimize the visual system preferring one shape over another due to low-level saliency. The stimuli were displayed until response with 'Incorrect' showing for 1500 ms for an incorrect response. No feedback was given for correct responses.

Practice. The 10 trials of practice were used to familiarize participants with the task. During this phase, the target was equally likely to appear in either shape.

Training. In the 192 trials of the training phase, the target appeared in a certain shape on 75% of trials (High Probability). The target appeared in the other shape on 25% of trials (Low Probability). In this way, shape became a relevant feature that the visual system can use in order to complete the task more efficiently.

Testing. In the 192 trials of the testing phase, the target appeared in each shape during an equal amount of trials (50% circle, 50% square). The Training Phase progressed to the Testing Phase without participants' knowledge and without a break in the trials.

Results

Search Accuracy

Overall search accuracy was 97.5%. There was no effect of probability condition on search accuracy, $t(23) = 0.62$, $p = .541$, with equally high rates of accuracy for both the High Probability condition (97.4%) as the Low Probability condition (97.6%).

Manual RT: Training Phase

The critical measure was mean response time (RT) as a function of shape probability in the training and testing phase. The analysis was limited to correct search trials only, and trials with RTs more than 2.5 SD from a participant's mean in each condition were removed (5.3% of trials). The pattern of results was not influenced by trimming in either experiment in this study. In the training phase, participants were faster at finding the target letter when it was located in the High Probability (825 ms) shape than in the Low Probability shape (861 ms), $t(23) = 3.81$, $p < .001$.

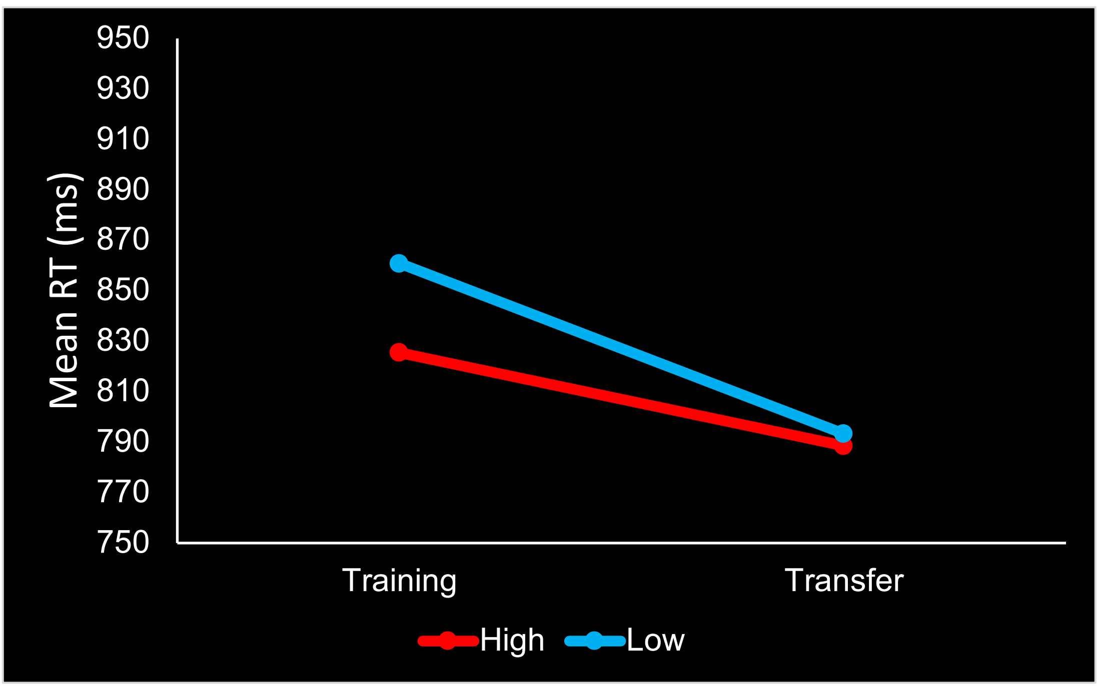

Figure 2. Results from Experiment 1. Mean RTs as a function of whether the target was more likely to be in a certain shape (Training) or equally likely to be in either shape (Transfer)

Manual RT: Testing Phase

However, learned bias from training did not persist into the testing phase, $t(23) = 0.47$, $p = .645$, as participants were equally as fast in finding the target letter when it appeared in the

High Probability (789 ms) shape than in the Low Probability shape (793 ms). Moreover, even in the first block of testing, the difference in RTs between probability conditions was not significantly different, $t(23) = 1.16$, $p = .259$, although there was a numerical trend toward a difference (High: 763 ms; Low: 792 ms). In sum, participants learned to preferentially attend to the High Probability shape in the training phase, but this difference was rapidly extinguished in the testing phase.

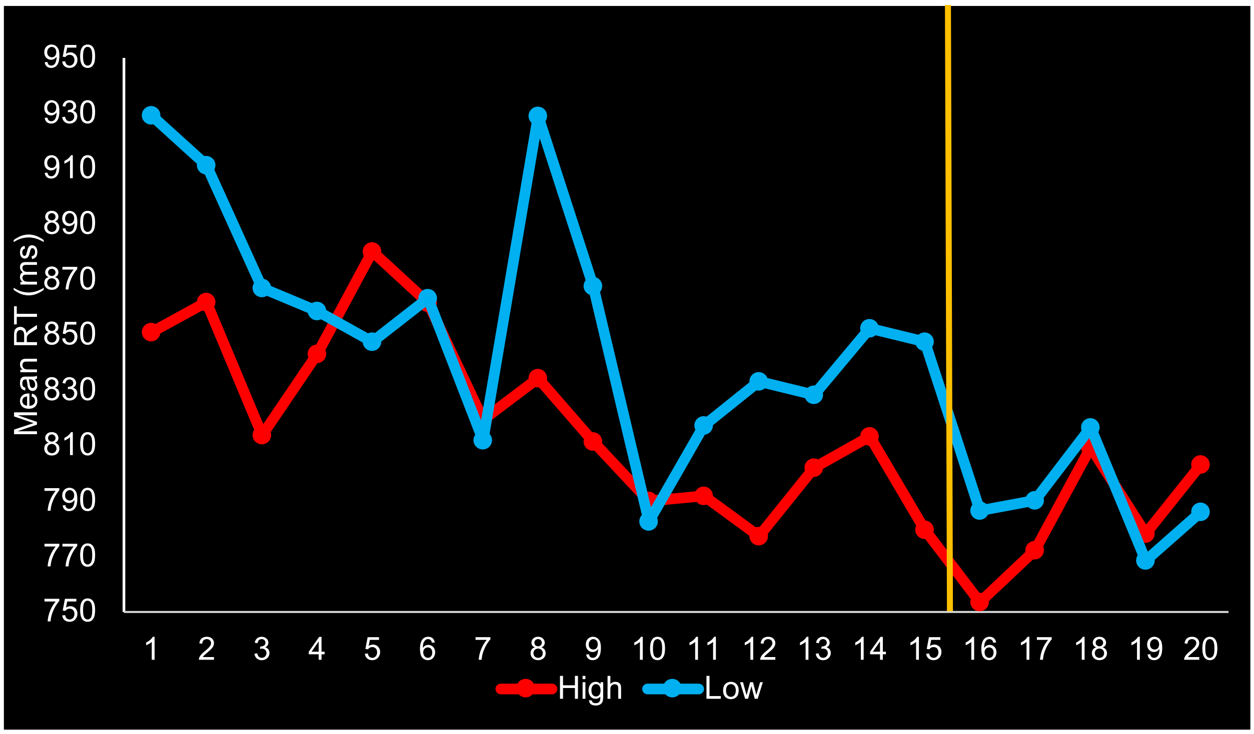

Figure 3. Results from Experiment 1. In the Training Phase (Blocks 1-15), the target appeared in a certain shape 75% of the time. In the Testing Phase (Blocks 16-20), the target was equally likely to appear in either shape. Each block contained 24 trials.

Discussion

Experiment 1 both replicated and extended previous research on visual statistical learning of a feature dimension. First, the results indicated that the difference in the distribution of target shape probability across trials led to faster RT in the High Probability condition than in the Low Probability condition. Therefore, the participants were able to learn to preferentially attend to the shape that most likely contained the target. This extends the findings of Sha et al., (2017) that feature dimensions can be implicitly, statistically learned in that both color and shape probabilities can be extracted. Second, again in support of Sha et al., (2017), the bias to attend to the previously rich probability shape during the testing phase was quickly extinguished. In other words, when the target was equally likely to appear in either color, the RT advantage disappeared. Therefore, we conclude that the visual system can extract feature-based information and use it to guide attention.

Experiment 2 (Pilot Study): Previously Learned Biases vs. VWM

The results of Experiment 1 suggest that it is possible to learn to preferentially attend to a certain feature (i.e. shape) that is relevant to the current task. This preference is implicitly learned quickly and is extinguished in a relatively short amount of trials when compared to previous research on location-based statistical learning (Jiang et al., 2012; Jiang, Swallow, Won, Cistera, & Rosenbaum, 2014). While both sources of information (feature and location) can be learned, attentional biases based on location persist for longer than it does for attentional biases based on features. One possible explanation for this difference is that, for features, the visual system uses the contents of VWM to guide attention while relying on more long-term learned biases for locations. On the other hand, it is possible the system learns to extract the feature-based information so quickly and efficiently that when the probabilities of the task change (from 75/25 to 50/50), the system quickly adjusts and guides attention accordingly.

Attention is guided by the contents of VWM and learned regularities of the task environment. A question therefore remains: what is the nature of, if one exists, the interaction between the contents of VWM and statistical learning on the guidance of attention? To the investigate this question, we trained participants on the same implicit learning task as Experiment 1, but introduced a color to be held in VWM. This color either matched the target, the distractor, or was a color that was not present in the display.

If the visual system only uses the contents of VWM, then we would expect the most distraction (slowest mean RTS) when the distractor matches the contents of VWM, the most facilitation (fastest mean RTs) when the target matches the contents of VWM. Both compared to when no VWM color is present in the display. But, the magnitude of this match effect (VWM match – VWM mismatch) should not be affected by the shape probabilities (75% versus 25%). This pattern of results would provide evidence for VWM driving attentional guidance, not statistical learning.

Alternatively, if the system does not rely on the contents of VWM, and attentional guidance is driven primarily by learned biases, we would expect faster mean RTs on the trails where the target was in the high probability shape than when the target was in the low probability shape (a replication of the training phase in Experiment 1). Critically, this learned bias (high probability – low probability) would not change as a function of the contents of VWM: the learned bias would be equal across VWM match and mismatch conditions.

If there is an interaction between the contents of VWM and learned biases on the guidance of attention, we would expect the fastest mean RTs when both the target matched the color held in VWM and the target was in the more probable shape. We would expect the slowest mean RTs when both the distractor was the color held in VWM and the target was in the low probability shape. All other conditions would fall in between those two extremes.

Method

Participants

There were 11 participants in Experiment 1, 10 females and 1 male. Their mean age was 19 years old.

Procedure and Design

Participants completed 10 trials of practice and 192 trials of training for Part 1. Then they completed an additional 10 trials of practice and another 192 trials of the VWM task in Part 2. After the first practice phase, participants continued through the training phase without further instruction. Once the participants reached Part 2, more instructions were given with another practice phase before moving onto the VWM phase.

Part 1: Practice, Training Phase

The task for Part 1 was identical to the training phase for Experiment 1. The target 'D' appeared in a certain shape on 75% of trials (High Probability) and in the other shape on the remaining 25% of the trials (Low Probability). As in Experiment 1, shape is a relevant feature that the visual system can use in order to complete the task more efficiently.

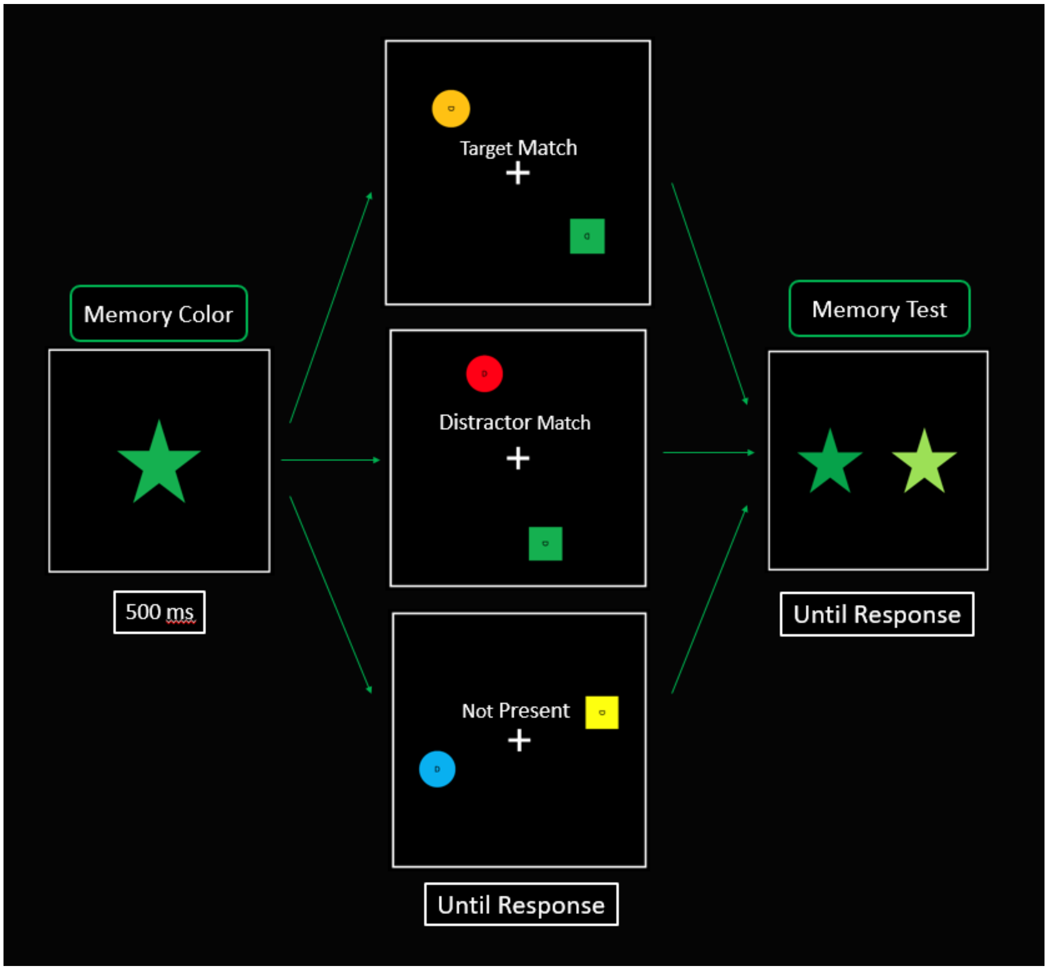

Figure 3. Sequence of events in a trial for Experiment 2.

Part 2: Practice, VWM Phase

The search task is the same as the one used in Experiment 1 and Part 1 of Experiment 2 with an added VWM component. Before each search display, a round-tipped star appeared that was a certain color (Fig. 2). The participant was asked to remember that color and be able to recall that color. Then the visual search task was completed as it was in Experiment 1. The contingency also remained from Experiment 1 throughout the entirety of Experiment 2. That is, including during the VWM component, the target appeared in the same previously learned shape on 75% of the trials and in the other shape on 25% of the trials. Therefore, shape remained a reliable source of information that can be used to guide attention to the target.

Overall, in 50% of the total trials the color held in VWM was in the display (match) and in the other 50% of trials, the color was not present in the display (mismatch). Within the match trials, 50% of those trials had the color match the target (target-match) and 50% had the color match the distractor (distractor-match). After the response had been made, two round-tipped stars appeared that were two different colors. One star was the color that the participant was asked to remember and the other was a different hue from the same color category as the remembered color. The participant was instructed to choose the star that matched the one they saw at the beginning of the trial. In this manner, the color was held in VWM during the search task.

Results

Training Phase

Search Accuracy

Overall search accuracy was 97.6%. There was no effect of probability condition on search accuracy, $t(23) = 0.09$, $p = .926$, with equally high rates of accuracy for both the High Probability condition (97.6%) as the Low Probability condition (97.7%).

Manual RT: Training Phase

To assess whether participants learned which shape contained the target more frequently, we first examined mean RT as a function of probability condition in the training phase. The analysis was limited to correct search trials only, and trials with RTs more than 2.5 SD from a participant's mean in each condition were removed (2.4% of trials). As in Experiment 1, participants were faster at finding the target letter when it was located in the High Probability (801 ms) shape than in the Low Probability shape (834 ms), $t(7) = 2.50$, $p = .041$.

Thus, participants learned which shape was more likely to contain the target letter and used this information to preferentially attend to that shape. However, the critical question was what would happen to this learning when a VWM load was introduced.

Testing Phase (added VWM load)

Search Accuracy

Overall search accuracy was 98.8%. Data for both search accuracy and RT were entered into a 3 (match condition: target-match, distractor-match, mismatch) X 2 (probability condition: high, low) ANOVA. There was no main effect of match condition, $F(2,14) = 2.094$, $p = .160$, nor probability condition, $F(1,7) = 0.482$, $p = .510$, on search accuracy, $t(23) = 0.62$, $p = .541$. The interaction also did not reach significance, $F(2,14) = 0.594$, $p = .565$.

Memory Accuracy

Overall memory test accuracy was 85.1%. There was no main effect of match condition, $F(2,14) = 0.378$, $p = .692$, nor probability condition, $F(1,7) = 3.118$, $p = .121$, on memory accuracy, $t(23) = 0.62$, $p = .541$. The interaction also did not reach significance, $F(2,14) = 1.808$, $p = .200$.

Manual RT

The critical measure was RT as a function of color match and probability condition. Trials were again removed according to the same criteria as previous experiments (1.4% of trials). In contrast to the training phase, there was no longer an effect of probability condition, with participants equally as fast at finding the target letter when it appeared in either the High Probability (766 ms) or Low Probability (767 ms), $F(1,7) = 0.001$, $p = .985$. Surprisingly, there was no reliable main effect of color match, $F(2,14) = 1.388$, $p = .282$.[1] However, there was a numerical trend to find the target faster when the color of the target-possessing shape matched the contents of VWM compared with when it did not: 751 ms for target-match, 776 ms for distractor-match, 771 for mismatch. Finally, the interaction between probability condition and color match did not reach significance, $F(2,14) = 0.625$, $p = .550$.

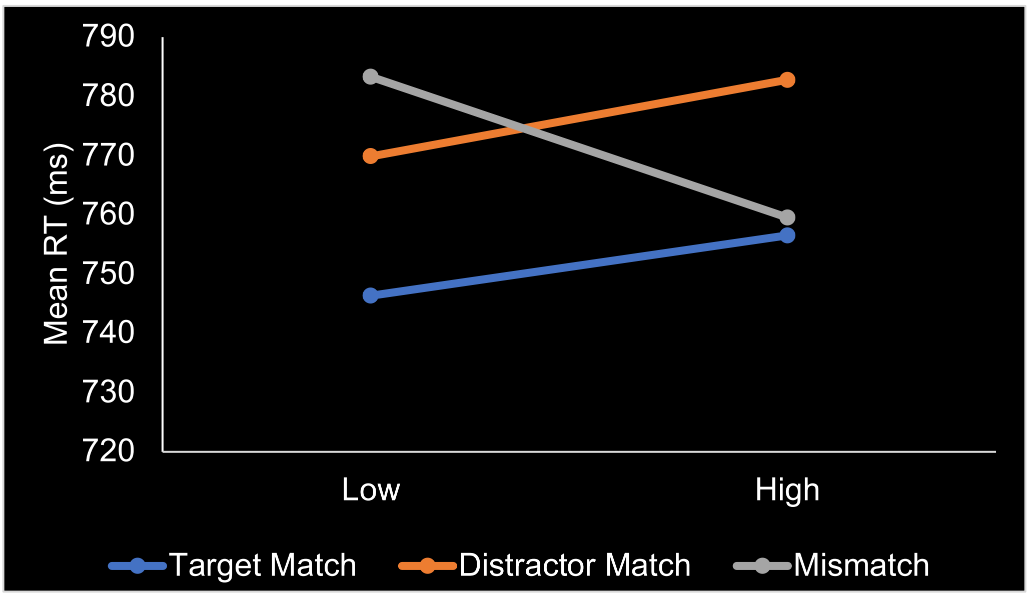

Figure 4. Results from Testing Phase of Experiment 2. VWM condition as a function of whether target was in the less probable shape (Low) or the more probable shape (High) measured as RT.

[1] However, previous research has demonstrated that color match effects on RT require 60+ participants to achieve 80% power (Bahle, Beck, & Hollingworth, 2018). Thus, this pilot study was substantially underpowered.

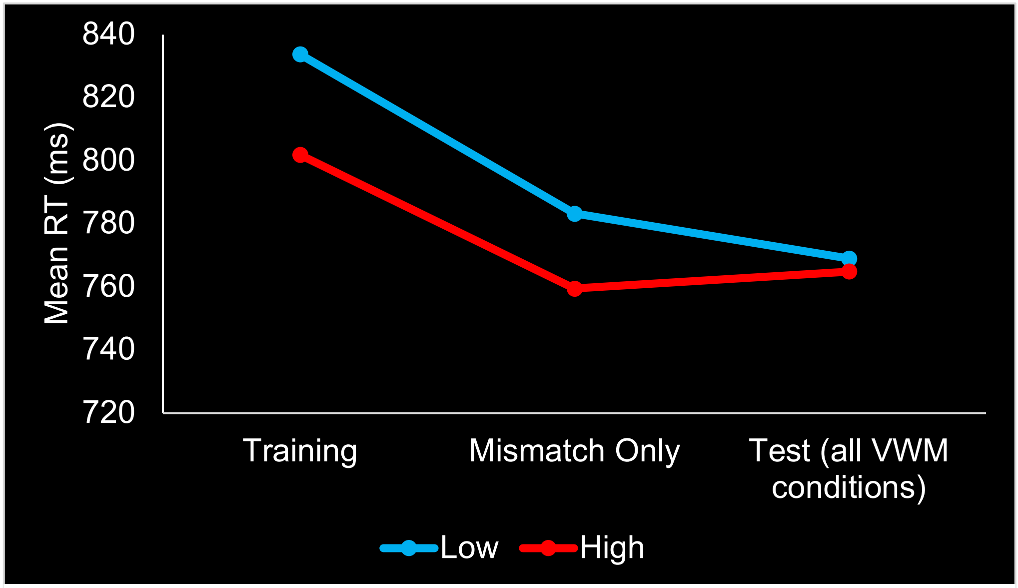

Figure 5. Results from Testing Phase of Experiment 2. Probability condition (Low vs. High) as a function of VWM measured in RT. Training contained no VWM component. Mismatch Only are the trials where the display did not contain the color held in VWM. The Test trials are all trials that contained the color held in VWM within the display.

Discussion

Experiment 2 examined the interaction between the bias to attend to the previously learned High Probability shape, and the color held in VWM. Even though the probabilities remained from the training phase, participants were as fast at finding the High Probability target as the Low Probability target. Therefore, including a simultaneous VWM load extinguished the learning from the training phase, even though the probabilities remained the same as in the training phase. This suggests that the statistical learning of a relevant feature was trumped by the contents of VWM because when another, more salient item was held in VWM, attentional guidance through statistical learning did not occur.

Although not significant, participants were slightly faster when the color held in VWM matched the color of the target-possessing shape than when it did not. This suggests that the color held in VWM guided attention while the previous attention biases based on statistical learning did not. Additionally, the slowest responses came when the color held in VWM matched the distractor. This further supports the idea that the contents of VWM overpowered any previously learned attentional bias.

General Discussion

Experiment 1 both replicated and extended previous research on visual statistical learning of a feature dimension by suggesting that participants were able to learn to preferentially attend to the shape that most likely contained the target. In contrast with previous research on location-based statistical learning, the attentional bias towards the high probability shape was quickly extinguished. Therefore, the question emerges: Why does location-based learning have long lasted attentional biases while the attentional biases of feature-based learning are quickly diminished?

One explanation is that learned feature-based biases occur on a short time scale: the system operates in such a way that instead of forming longer-term biases for features (biases likely residing in long-term memory), the system keeps track of what has been occurring frequently and recently (the few last dozen trials). The system likely quickly adapts to changes in environmental statistics, so much so that when the probabilities of the feature-based information change from 75% and 25% to 50/50 (Experiment 1), the system might quickly detect this change and modify attentional guidance accordingly. Due to the shorter-term nature of these biases, any available contents of VWM likely outweighs attentional guidance based on learned feature-based statistical learning.

When the system has access to more active representations (i.e. templates) in VWM, the system may rely less on longer term biases and emphasize the current trial's objective.

Even when the visual system is given feature-based statistical learning (the persistence of the probability manipulation in Experiment 2's testing phase), the contents of VWM seem to be fully in control of attentional guidance. One major limitation of this conclusion is the nature of Experiment 2. We were severely underpowered, and thus we cannot make strong conclusions based on the data presented. However, despite this limitation, the data were trending in the direction of strong VWM guidance outweighing statistical learning.

Further research is needed to further examine the interaction between previously learned feature-based attentional biases and the content held in VWM. One possible direction is to examine the effect of statistically learned distractor suppression concurrent with a VWM load. Will the visual system fail to successfully suppress the distractor once a competing feature is held in VWM? Further research should also address the aforementioned limitation and conduct a sufficiently powered investigation.

To conclude, the results of Experiment 1 and Experiment 2 suggest an explanation for the quickly terminating attentional biases of feature-based statistical learning. Features held in VWM outcompetes any previously learned attentional biases that rely on sources that can be maintained within the contents of VWM, such as statistical learning of features.

Reference

Beck, V. M., Hollingworth, A., & Luck, S. J. (2012). Simultaneous control of attention by multiple working memory representations. Psychological science, 23(8), 887-898.

Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. Cognitive psychology, 36(1), 28-71.

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. Journal of Experimental Psychology: Human perception and performance, 18(4), 1030.

Goujon, A., Didierjean, A., & Thorpe, S. (2015). Investigating implicit statistical learning mechanisms through contextual cueing. Trends in Cognitive Sciences, 19(9), 524-533.

Hollingworth, A., & Luck, S. J. (2009). The role of visual working memory (VWM) in the control of gaze during visual search. Attention, Perception, & Psychophysics, 71(4), 936-949.

Jiang, Y. V., Swallow, K. M., Rosenbaum, G. M., & Herzig, C. (2013). Rapid acquisition but slow extinction of an attentional bias in space. Journal of Experimental Psychology: Human Perception and Performance, 39(1), 87.

Jiang, Y. V., Swallow, K. M., Won, B. Y., Cistera, J. D., & Rosenbaum, G. M. (2015). Task specificity of attention training: the case of probability cuing. Attention, Perception, & Psychophysics, 77(1), 50-66.

Sha, Z. L., Remington, R. W., & Jiang, Y. V. (2017). Short-term and long-term attentional biases to frequently encountered target features. Attention, Perception, & Psychophysics, 79(5), 1311-1322.

Stilwell, Bahle, Vecera, Manuscript in preparation

Van der Stigchel, S., & Hollingworth, A. (2018). Visuospatial Working Memory as a Fundamental Component of the Eye Movement System. Current Directions in Psychological Science, 0963721417741710.

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. Psychonomic bulletin & review, 19(5), 871-878.

Wolfe, Jeremy M. "Guided search 2.0 a revised model of visual search." Psychonomic bulletin & review 1.2 (1994): 202-238.