Published in final edited form as:

J Exp Psychol Hum Percept Perform. 2014 August ; 40(4): 1654–1664. doi:10.1037/a0037267.

The Role of Reward Prediction in the Control of Attention

Anthony W. Sali, Brian A. Anderson, and Steven Yantis

Department of Psychological and Brain Sciences, Johns Hopkins University

Abstract

Previously rewarded stimuli involuntarily capture attention. The learning mechanisms underlying this value-driven attentional capture remain less understood. We tested whether theories of prediction-based associative reward learning explain the conditions under which reward feedback leads to value-based modulations of attentional priority. Across four experiments, we manipulated whether stimulus features served as unique predictors of reward outcomes. Participants received monetary rewards for correctly identifying a color-defined target in an initial search task (training phase) and then immediately completed a second, unrewarded visual search task in which color was irrelevant (test phase). In Experiments 1–3, monetary reward followed correct target selection during training, but critically, no target-defining features carried uniquely predictive information about reward outcomes. Under these conditions, we found no evidence of attentional capture by the previous target colors in the subsequent test phase. Conversely, when target colors in the training phase of Experiment 4 carried uniquely predictive information about reward magnitude, we observed significant attentional capture by the previously rewarded color. Our findings show that value-based attentional priority only develops for stimulus features that carry uniquely predictive information about reward, ruling out a purely motivational account and suggesting that mechanisms of reward prediction play an important role in shaping attentional priorities.

Keywords

attentional capture; associative learning; perceptual learning; reward and attention

Attentional selection determines what information receives cognitive processing, gating access to awareness, decision making, and storage in memory (e.g., Desimone & Duncan, 1995; Mack & Rock, 1998; Rensink, O'Regan, & Clark, 1997; Yantis & Johnston, 1990). The physical salience of a stimulus as well as an individual's current goals together determine the likelihood of attentional selection, known as attentional priority (Serences & Yantis, 2006; Yantis & Egeth, 1999). An increasing body of work has suggested that reward history also modulates the computation of attentional priority (e.g., Della Libera & Chelazzi, 2009; Hickey, Chelazzi, & Theeuwes, 2010a, 2010b; Raymond & O'Brien, 2009) such that stimuli previously associated with reward involuntarily capture attention (e.g., Anderson, Laurent, & Yantis, 2011a, 2011b). Attentional capture resulting from reward history has been referred to as *value-driven attentional capture* (Anderson et al., 2011b; for a review, see Anderson, 2013). The learning mechanisms responsible for linking reward history to the

development of value-driven attentional priority remain largely unexplored. In the current series of experiments, we tested whether prediction-based associative learning between stimulus features and reward outcomes plays an important role in value-based modulations of attentional priority.

Value-driven attentional capture is demonstrated when task-irrelevant but previously rewarded stimuli persistently capture attention in the absence of further reward feedback (e.g., Anderson et al., 2011a, 2011b). In a typical experiment, a two-part design is employed in which participants first receive monetary rewards for correctly selecting color-defined targets in a *training phase* and then subsequently complete a *test phase* involving a second search task in which color is task-irrelevant (Anderson et al., 2011b). Critically, on half of the trials during the test phase, one of the non-targets is rendered in the color of a formerly rewarded target from the training phase, which serves as a task-irrelevant distractor. Using this design, response times are slowed by the presence of a formerly rewarded distractor (e.g., Anderson et al., 2011b), therefore supporting the conclusion that reward history modulates attentional priority. Subsequent studies have shown that these modulations of attentional priority transfer across different tasks of attentional control (Anderson et al., 2012), can result in oculomotor capture (Anderson & Yantis, 2012), and persist for long periods of time without intervening reward learning (Anderson & Yantis, 2013).

Animal models have demonstrated that a learned association between a particular stimulus and reward develops only when that stimulus offers uniquely predictive information about the receipt of reward. Rodents fail to learn an association between reward and stimuli that provide reward-predictive information that is redundant with other stimuli in the environment (e.g., Kamin, 1969, Pearce & Hall, 1980). Associative reward learning is mediated by dopaminergic neurons originating within the ventral tegmental area and substantia nigra, which project to the ventral striatum and encode both reward predictions and subsequent errors in the prediction of reward (e.g., Schultz, Dayan, & Montague, 1997; Waelti, Dickinson, & Schultz, 2001). One possibility is that when these midbrain reward-prediction signals are bound to a specific stimulus feature, the use of that feature as a unique predictor of reward biases attention to select that feature in the future.

Previous accounts of value-driven attentional capture have all used stimulus sets in which the training phase target colors provided unique information about the magnitude of upcoming reward that was otherwise unpredictable. One explanation for the development of value-based attentional priority under these conditions is that over the course of training, participants used the target color to generate a reward-prediction signal, and this prediction signal served an important role in modulating attentional priority. Prediction-based associative learning consequently provides one mechanistic account of how reward history modulates subsequent attention.

Mechanisms of learning that do not involve reward prediction could also account for previous demonstrations of value-driven attentional capture. In particular, one alternative possibility is that reward, which serves as a motivator of accurate performance, may reinforce the development of automatic selection by speeding the rate at which perceptual learning unfolds. In the absence of reward, the accumulation of experience over the course

of thousands of trials influences the tuning of the visual system, such that the ability to discriminate a particular orientation or direction of motion becomes sharper (Roelfsema, Ooyen, & Watanabe, 2010; Watanabe et al., 2001). In the context of visual search, repeated selection of a particular target stimulus gives rise to later impairment in visual search if the previous target serves as a distractor, suggesting that selection history can give rise to automatic selection (Shiffrin & Schneider, 1977; see also Kyllingsbaek, Schneider, & Bundesen, 2001). To the extent that the motivation provided by reward enhances the development of automatic selection, such perceptual learning could account for the influence of reward history on subsequent attention without implicating mechanisms of reward prediction.

Under unrewarded conditions, attentional capture resulting from perceptual learning occurs on a much longer timescale than previous demonstrations of value-driven attentional capture, often unfolding over the course of thousands of trials of training (e.g. Shiffrin & Schneider, 1977). Attentional priority for previously rewarded targets develops much more quickly and cannot be fully explained by reward-independent perceptual learning. An attentional bias for former targets is either weaker or does not occur under equally brief training conditions that do not involve reward (Anderson et al., 2011a, 2011b, 2012; Qi, Zeng, Ding, & Li, 2013; Wang, Yu, & Zhou, 2013) and value-based attentional priority is modulated by the magnitude of reward (Anderson et al., 2011a, 2011b, 2012; Anderson & Yantis, 2013; Theeuwes & Belopolsky, 2012). However, it is known that reward incentives can modulate the voluntary deployment of attention to task-relevant stimuli, improving the efficiency of target localization processes (e.g., Kiss, Driver, & Eimer, 2009; Kristjansson, Sigurjonsdottir, & Driver, 2010; Pessoa & Engelmann, 2010). One possibility, then, is that the introduction of monetary rewards for accurate target detection in a visual search task facilitates stronger or more rapid perceptual learning for target items by motivating the enhancement of goal-directed processing.

Existing demonstrations of value-driven attentional capture are consistent with both associative reward-learning and motivated perceptual learning accounts. For example, targets associated with a high magnitude reward during training show behavioral response times in a subsequent test phase to a greater degree than targets that were associated with a low magnitude reward (e.g. Anderson et al., 2011b). According to an associative learning account, the high reward magnitude target acquires greater value over the course of training than does the low reward magnitude target as the result of its ability to predict when larger rewards occur. However, it is also possible that the asymmetric reward schedule leads to greater motivation to find and report the high-reward target and consequently stronger perceptual learning for that stimulus. In the present study, we tested whether stimulus features must carry uniquely predictive information about reward magnitude for the rapid development of value-based attentional priority.

To rule out a non-associative learning mechanism of value-driven attentional capture, we conducted a series of experiments in which participants were rewarded for accurate visual search, but target stimuli did not differentially predict reward outcomes. If incentivizing target selection with the prospect of receiving a reward is sufficient to modulate attentional priority regardless of the predictive relationship between targets and reward, previously

rewarded target stimuli should involuntarily capture attention even under conditions in which associative learning for target features does not occur. Conversely, if associative reward learning is necessary for value-driven attentional capture, previous target items should only involuntarily capture attention when a feature such as color uniquely predicts reward outcomes. Using this logic, we tested non-associative and associative reward learning accounts of value-driven attentional capture across four experiments. The design of each experiment was similar to that used by Anderson and colleagues (2011b). Participants completed an initial training phase in which they received monetary rewards for accurate performance in a visual search task for color-defined targets. All participants then immediately completed an unrewarded test phase involving visual search for a shape singleton target in which color was irrelevant to the task and stimuli possessing the color of a former target served as physically non-salient distractors.

In Experiment 1, participants searched for a consistent target color throughout the training phase and received unpredictable high and low magnitude rewards, while in Experiment 2 participants received a fixed amount of reward for accurately identifying a target from a set of two possible colors. In a third experiment, rewards for identifying a target of a particular color were based on the speed of participants' behavioral response times (RTs), providing strong incentive for fast and accurate selection. Critically, although training phase targets were defined by color as in previous tests of value-driven attentional capture (e.g., Anderson et al., 2011b), target color did not serve as a unique predictor of reward in Experiments 1–3. Instead, target color was redundant with a variety of other factors such as target shape and button press accuracy in predicting subsequent reward. Under these conditions in which the target features would not be expected to generate reward-prediction signals, we found no evidence of value-driven attentional capture. Finally, in Experiment 4, we replicated value-driven attentional capture under otherwise equivalent conditions in which target color uniquely predicted the magnitude of reward. We conclude that through reward prediction mechanisms, associative learning between stimuli and reward gives rise to value-based attentional priority.

Experiment 1

If reward prediction mechanisms are not necessary for the rapid development of value-based attentional priority, merely providing a reward incentive for target selection during training should bias attention toward stimuli possessing target-defining features in a subsequent test phase. We therefore modified the paradigm used by Anderson and colleagues (2011b) such that there was only a single target color during the initial training phase. Participants received monetary rewards for correctly reporting the orientation of a line segment within a color-defined target (either red or green, counterbalanced). All participants then completed a shape singleton search task in which color was irrelevant and no further monetary rewards were issued. If the mere presence of reward incentives during the initial color search task is sufficient to modulate the attentional priority assigned to objects, RTs should be slowed by the presence of former-target-color distractors. Conversely, if associative learning between stimulus features and reward outcomes is necessary for value-based modulations of attentional priority, RT should not differ between the distractor present and distractor absent conditions.

Method

Participants—Eighteen participants (16 female) ranging in age from 19–23 years ($M = 20.06$, $SD = 1.06$) completed the study in exchange for monetary compensation. All participants were right handed and completed a consent form approved by the Johns Hopkins University Institutional Review Board prior to participating in the study. One participant was removed from the analysis for having a mean behavioral accuracy below 75 percent during the test phase and another was excluded for previously participating in a different study of reward and attention (see Anderson & Yantis, 2013).

Apparatus—Participants viewed the stimuli on an Asus VE247 monitor connected to a Mac Mini computer. The monitor was positioned approximately 76 cm from the participant. Stimulus presentation and participant response recording was controlled with the Psychophysics Toolbox for Matlab (Brainard, 1997). Participants made all responses using a standard computer keyboard positioned in front of the monitor.

Stimuli

Training phase: Each search array consisted of 6 colored circles ($3.01^\circ \times 3.01^\circ$) arranged in an imaginary circle (radius = 6.01°) around a central fixation cross ($0.75^\circ \times 0.75^\circ$) (see Figure 1). On each trial, participants searched for a target circle that was defined by color. For half of the participants, the target circle was always the color red; for the remaining participants, the target was always green. Across all participants, the colors of the non-target circles were randomly selected from blue, yellow, white, orange, purple, and red/green on each trial without replacement. The line segment within the target was either vertical or horizontal (equally-often), while the line segment within each non-target circle was randomly tilted 45 degrees to either side of the vertical meridian. The target appeared in each of the array locations an equal number of times throughout the training phase.

Test phase: Participants searched for the unique shape in an array of 6 objects (either a single circle among 5 diamonds or a single diamond among 5 circles). As in the training phase, participants reported the orientation of the line segment displayed within the target shape. The size and position of the stimuli were identical to the training phase. Critically, on half of the trials, one of the distractor shapes was the previous target color from the training phase (distractor present trials). For the remaining trials, none of the shapes were the previous target color (distractor absent trials). The ordering of distractor present and distractor absent trials was randomized over the course of the test phase. The target and non-targets other than the previously rewarded distractor were selected from the same set of colors as in the training phase; the target in the test phase was never the previous target-defining color from the training phase.

Procedure

Training phase—Participants received instructions to indicate the orientation of the line segment within the target circle on each trial as rapidly as possible while still maintaining a high level of accuracy, using the "z" key for vertical and the "m" key for horizontal. Each of 144 total trials began with the presentation of a central fixation cross for 400, 500, or 600 ms (randomly determined). The search array then appeared and was displayed until the

participant made a response or 1000 ms had elapsed, after which the trial timed out. Reward feedback, which consisted of the earnings on the current trial in addition to a running total of accumulated earnings, was then presented for 1500 ms, following a blank 1000 ms interval. Lastly, there was a blank 1000 ms inter-trial-interval (ITI) (see Figure 1A). Participants received a monetary reward for each correct response they made during the training phase task. Participants could receive a high magnitude reward of 6 cents on 75 percent of the trials and a low magnitude reward of 2 cents on the remaining 25 percent of trials. The reward magnitude was randomized across the training phase trials using these parameters. If participants responded incorrectly or too slowly, the reward feedback indicated that 0 cents had been earned. Responses that were too slow (i.e., the trial timed out) were additionally followed by a 1000 Hz 250 ms tone. Since the target color remained the same across all trials, no aspect of the search array predicted whether the participant would receive a high or low magnitude reward for a correct response on a trial-by-trial basis.

Test phase—Immediately after completing the training phase, all participants completed a second visual search task in which they identified the orientation of a line segment within the shape singleton. Participants were informed that they would no longer receive monetary rewards for correct responses. They were instructed that color was irrelevant in this task and that they should search for the unique shape only. Each of 120 total trials began with the presentation of a central fixation cross for 400, 500, or 600 ms (randomly determined). The search array was then presented until the participant made a response or 1500 ms had elapsed, after which the trial timed out. A white "X" ($1.28^\circ \times 1.28^\circ$) was presented centrally for 1000 ms following incorrect responses (this display was omitted following correct responses) and there was a 500 ms blank ITI. As in the training phase, a 1000 Hz 250 ms tone followed responses that were too slow.

Results

We tested whether previously target-defining features from the training phase involuntarily captured attention when they served as distractors in a subsequent visual search task. Specifically, we contrasted RT for trials in which the previous target color served as a distractor in the test phase (distractor present) against those test phase trials in which the previous target color was absent from the array (distractor absent). We computed a mean RT for distractor present and distractor absent trials in the first and second half of the test phase for all participants. Furthermore, RTs more than 2.5 SD above or below the subject's mean RT for each condition were considered outliers and trimmed from the analysis, resulting in the removal of less than 2 percent of trials. We conducted a split-half analysis to examine the possibility of rapid extinction of learning accounting for a null effect when collapsing across the entire test phase. Consistent with an associative learning explanation of value-driven attentional capture, a repeated measures analysis of variance (ANOVA) with 2 levels of distractor presence (present vs. absent) and 2 levels of experimental half (first vs. second) revealed that there was no significant main effect of distractor presence, $F(1,15) = 0.22$, $p = .648$, 95% CI: -29.28 – 18.78 (see Figure 2A). We computed the JZS Bayes factor for the main effect of distractor presence (see Rouder, Speckman, Sun, Morey, & Iverson, 2009) to quantify the likelihood that the null hypothesis was true. The JZS Bayes factor was 4.77, suggesting that the null hypothesis (that the distractors had no influence on performance)

was over 4 times as likely to be true as was the alternative hypothesis (that there was a difference in RT based on distractor presence). There was a significant main effect of experimental half, $F(1,15) = 6.61$, $p = .021$, $\eta_p^2 = .306$, such that participants were faster overall in the second half of the test phase than in the first half. Lastly, the interaction between distractor presence and experimental half failed to reach statistical significance, $F(1,15) = 0.23$, $p = .642$.

We also examined whether participants differed in accuracies for distractor absent and distractor present trials across both halves of the test phase. The main effect of distractor presence failed to reach statistical significance, $F(1,15) = 0.38$, $p = .546$ (see Figure 2B). Participants were more accurate in the second half of the test phase than in the first half, $F(1,15) = 12.94$, $p = .003$, $\eta_p^2 = .463$. As with the RT analysis, the distractor presence by experimental half interaction failed to reach statistical significance, $F(1,15) = 0.09$, $p = .766$.

Discussion

Both the RT and accuracy results of Experiment 1 showed no evidence of value-driven attentional capture; rewarding target selection during the training phase did not provide a lasting modulation of attentional priority for the previous target color. In this experiment, target color did not predict reward magnitude on a trial-by-trial basis (as in, e.g., Anderson et al., 2011b). Rather, participants could predict the upcoming reward equally well without considering the target color, and under such conditions prediction-based associative learning between a stimulus and reward does not occur (e.g., Kamin, 1969).

A perceptual learning effect on automatic selection such as that demonstrated by Shiffrin and Schneider (1977) would likely have emerged with sufficient training exposure in our paradigm. In the current experiment we used only 144 training trials, whereas Shiffrin and Schneider as well as Kyllingsbaek and colleagues (2001) both found evidence of perceptual learning after over 2,000 instances of selecting a target. Nevertheless, previous work on reward and the control of attention has suggested that modulations of attentional priority persist into extinction after as few as 240 trials of training (Anderson et al., 2011b). The results of the current study provide evidence that reward-induced motivation is alone not sufficient to increase the robustness of perceptual learning such that attentional priority is modulated following such brief training.

In addition to the lack of uniquely predictive information carried by training phase targets, a difference between the current paradigm and that used by Anderson and colleagues (2011b) is that participants searched for a single target-defining color during the training phase in the current study. It is possible that motivation-based non-associative learning may influence attentional priority in the case in which participants must search for more than a single target-defining feature on each trial, which requires more effortful and deliberate goal-directed attentional control. To rule out this possibility, we conducted Experiment 2 in which all participants searched for either a red or green circle on each trial of the training phase as in earlier studies, but target stimuli did not carry predictive information regarding the magnitude of reward (Anderson et al., 2011b).

Experiment 2

Given that we did not find evidence of value-driven attentional capture following search for a single target color, in Experiment 2 we modified our training task such that participants searched for two potential target colors. As in previous studies of value-driven attentional capture, participants reported the orientation of a line segment within either a red or green target circle during the training phase (Anderson et al., 2011b). However, participants always received a reward of 6 cents following each correct training phase trial. Participants therefore searched for a target set of more than a single item, but color did not differentially predict either the likelihood of receiving reward or the magnitude of the reward on each trial. As in Experiment 1, a previous target color served as a distractor on half of the test phase trials.

Method

Participants—Eighteen healthy adults (11 female) ranging in age from 18–33 years ($M = 22.72$, $SD = 5.13$) participated in the study in exchange for monetary compensation. All participants but one were right handed. None of the participants had completed Experiment 1. All participants signed a consent form approved by the Johns Hopkins University Institutional Review Board. One participant was removed from the analysis for having a mean test phase behavioral accuracy below 75 percent, and another was excluded because they had previously completed a similar study on reward and attention with differing stimulus-color reward contingencies.

Visual Search Task

Training phase: The stimuli and procedure used in the training phase were identical to those used in Experiment 1 with the exception that participants searched for a red or green target circle on each trial and cyan was included in the set of non-target colors. Each search array contained only a single target, which was rendered in red or green equally-often. Participants received a consistent reward of 6 cents for each correct response made during the training phase.

Test phase: As in Experiment 1, participants searched for the unique shape in an array of circles and diamonds. All aspects of the task were the same as in Experiment 1, except as noted below. On half of the trials, none of the non-target shapes were the previous target colors of red or green. On one quarter of the trials, a non-target shape was red, and on the remaining quarter of trials, a non-target shape was green. The target was never red or green, and cyan was now included in the set of colors used.

Results

We again compared RT for trials in which a previous target color (either red or green) was present in the test phase search array versus those trials in which the previous target colors were absent. RTs were trimmed as in Experiment 1 to remove outliers, resulting in a reduction of less than 2 percent of the data. We conducted a repeated measures ANOVA to determine whether RT for distractor present and absent trials differed across both halves of the test phase. As in Experiment 1, there was no significant main effect of distractor

presence, $F(1,15) = 1.78$, $p = .202$, 95% CI: -6.73 – 29.31 , therefore providing further evidence against a non-associative learning account of value-driven attentional capture (see Figure 3A). The JZS Bayes factor for the main effect of distractor presence was 2.36, suggesting that the null hypothesis of no value-driven attentional capture was over 2 times as likely to be true as was the alternative hypothesis. Unlike Experiment 1, the main effect of experimental half failed to reach statistical significance, $F(1,15) = 0.69$, $p = .420$, in addition to the interaction between distractor presence and experimental half, $F(1,15) = 0.38$, $p = .547$.

We also examined whether accuracies differed based on the presence of the previous target color in the search array. The main effect of distractor presence failed to reach statistical significance, $F(1,15) = 0.41$, $p = .531$ (see Figure 3B). The main effect of experimental half approached significance, $F(1,15) = 3.89$, $p = .067$, $\eta_p^2 = .206$, such that participants tended to be more accurate in the second half of the test phase. The interaction of experimental half and distractor presence also failed to reach statistical significance, $F(1,15) = 1.39$, $p = .257$.

Discussion

Experiment 2 again failed to find evidence of value-driven attentional capture in the absence of associative learning between stimulus color and reward. As in Experiment 1, distractor features that had not carried uniquely predictive reward information during training did not involuntarily capture attention during the test phase as indicated by RT or accuracy. The results of Experiment 2, in conjunction with those from Experiment 1, suggest that motivating visual search with monetary reward incentives is alone not sufficient to modulate the attentional priority of targets after a brief training exposure.

Experiment 3

In Experiments 1 and 2 we found no evidence of value-driven attentional capture following training in which visual search targets did not uniquely predict trial-by-trial reward magnitude. In both of these experiments, the reward participants received on a given trial was unrelated to their behavior (beyond the need to respond correctly within the time limit), which may have limited motivation. In Experiment 3, we provide a strong test of whether the motivation provided by reward incentive can create a persistent attentional bias. Participants searched for a target defined by a consistent color, and the monetary reward received for correct target identification was a direct function of how quickly the target was reported. Therefore, participants had strong incentive to select the target as rapidly and efficiently as possible. If incentivizing the selection of a stimulus with reward persistently biases attention to select that stimulus in the future, we should see evidence of robust value-driven attentional capture in the test phase of this experiment. If, however, value-driven attentional priority is mediated by mechanisms of reward prediction, the results should mirror those of Experiments 1 and 2 and show no evidence of capture.

Method

Participants—Twenty-one healthy adults (17 women) ranging in age from 18–28 years ($M = 21.4$, $SD = 2.75$) participated in the experiment in exchange for monetary compensation.

Twenty participants were right handed. All signed a consent form that was approved by the Johns Hopkins Institutional Review Board prior to participating and none had participated in either Experiments 1 or 2. One participant was excluded from all analyses because they had previously completed a different study that used a similar training paradigm.

Visual Search Task

Training Phase: The stimuli and procedure were identical to those used in Experiment 1, except where noted below. Participants received a monetary reward that was based solely on RT for each correct button press. Reward, measured in cents, was based on the following algorithm: $\text{Reward} = \text{round}[(1000 - \text{RT}) * .01]$ (see Miranda & Palmer, 2013 for a similar calculation). For example, an RT of 600 ms would yield a reward of 4 cents, whereas an RT of 400 ms would yield a reward of 6 cents. Participants were explicitly informed that rewards were proportional to the speed of correct responses. Overall, RTs within the timeout limit ranged from 232 to 989 ms ($M = 533$, $SD = 101$) making trial-by-trial feedback magnitudes similar to those used in the first two experiments. Participants saw written feedback of “too slow,” or “incorrect” following trials in which they failed to respond within the timeout limit or made the wrong button press, respectively.

Test Phase: All aspects of the test phase were identical to Experiment 1.

Results

As in Experiment 1, we compared RT for trials in which the previous target color was present as a distractor to those in which the previous target color was absent from the search array with a 2×2 repeated measures ANOVA with factors of distractor presence (absent vs. present) and experimental half (first vs. second). RTs were trimmed as in the first two experiments, resulting in the removal of less than 2 percent of trials. Consistent with the associative reward learning account, there was no significant main effect of distractor presence, $F(1,19) = 0.58$, $p = .454$, 95% CI: -28.94 – 13.46 (see Figure 4A). The JZS Bayes factor was 4.44, indicating that the null hypothesis was more than 4 times as likely to be true than was the alternative hypothesis. The main effect of experimental half reached statistical significance, $F(1,19) = 13.73$, $p = .001$, $\eta_p^2 = .420$, indicating that participants responded with faster RTs in the second half, but there was no significant interaction of distractor presence and experimental half, $F(1,19) = 0.32$, $p = .581$.

Next we tested behavioral accuracies with a 2X2 repeated measures ANOVA with factors of distractor presence and experimental half. There was a significant main effect of experimental half, $F(1,19) = 17.21$, $p = .001$, $\eta_p^2 = .475$, such that participants were more accurate in the second half of the test phase (see Figure 4B). The main effect of distractor presence, $F(1,19) = 3.05$, $p = .097$, as well as the interaction of distractor presence by half, $F(1,19) < .01$, $p = .965$, failed to reach statistical significance.

Discussion

Experiment 3 provided converging evidence that the motivation provided by reward incentive cannot alone account for the rapid instantiation of value-based attentional priority. Unlike the first two experiments, trial-by-trial reward magnitudes were based on behavioral

performance, therefore providing a stronger test of the motivated perceptual learning hypothesis. Even under these conditions, we found no evidence of value-based modulations of attentional priority.

One limitation of the first three experiments that could potentially account for the lack of significant attentional capture by previous target stimulus features was that there were fewer training phase trials than in previous studies (Anderson et al., 2011a, 2011b, 2012; Anderson & Yantis, 2012). Previous accounts of value-driven attentional capture have used as many as 1,008 and as few as 240 training trials. In the current experiments, we reduced training to only 144 trials. We therefore ran Experiment 4 to replicate previous demonstrations of value-driven attentional capture using only 144 trials of training, and to provide a basis for comparing the results of the first three experiments to those from a condition in which target color predicted reward magnitude.

Experiment 4

We conducted Experiment 4 to determine whether associative reward learning modulates attentional priority with only 144 trials of training. Evidence of value-driven attentional capture in Experiment 4 would therefore suggest that the null effects of Experiments 1–3 were not merely the result of inadequate training exposure. Rather, the involuntary capture of attention by previously rewarded stimuli in Experiment 4 would provide evidence that associative learning and not motivation per se is necessary for persistent modulations of attentional priority.

In Experiment 4 we directly tested the importance of reward prediction (e.g., Schultz et al., 1997) for value-based modulations of attentional priority. Participants searched for the red or green circle on each trial as in Experiment 2. However, unlike Experiment 2, one target color (red for half of the participants) was associated with a high magnitude reward for the majority of trials and a low magnitude reward for the remaining trials. Conversely, the other target color was associated with a low magnitude reward for the majority of trials (as in, e.g., Anderson et al., 2011b). The differing reward contingencies for red and green targets therefore allowed for the prediction of reward magnitude based on the target color in the search array on a trial-by-trial basis. Such a predictive relationship provides the foundation for the use of stimulus feature information to reduce reward prediction errors, which may ultimately lead to reward-learning-based modulations of attentional priority. As in the first three experiments, participants completed a test phase in which they searched for the unique shape in each array and were instructed to ignore color. We predicted that in support of associative reward learning accounts of value-driven attentional capture, participants would be slower to detect singleton targets in the test phase when the previously high-value distractor was present in the search array than when it was absent.

Method

Participants—Seventeen participants (12 female) ranging in age from 18–28 years ($M = 21.47$, $SD = 2.81$) completed the experiment in exchange for monetary compensation. All participants but one were right handed and all signed a consent form approved by the Johns Hopkins University Institutional Review Board prior to beginning the study. One participant

was excluded from the analyses because of mean behavioral accuracy below 75 percent during the test phase.

Visual Search Task

Training phase: The stimuli and procedure were identical to those used in Experiment 2, except as noted below. For each participant, one target color (red or green, counterbalanced) yielded a high magnitude reward of 10 cents on 83.33 percent of correct trials and a low magnitude reward of 2 cents on 16.67 percent of correct trials, while for the other color these contingencies were reversed. The average trial-by-trial reward was the same as in Experiment 2.

Test phase: The test phase used the same design as Experiment 1. On half of the trials, a non-target shape appeared in the previously high-value target color (either red or green, counterbalanced across participants). For the remainder of the trials, the previously highly-rewarded target color was absent from the search array. Non-targets and targets were selected from the same array of colors from Experiment 1, such that the test phase target was sometimes rendered in the low-value target color from the training phase; the previously high-value target color was never used for a target in the test phase, and was unrelated to which color served as the target color on a given trial.

Results

We tested whether the previous high-value distractor captured attention in the test phase following training in which the two target colors differentially predicted reward magnitude on a trial-by-trial basis. RTs were trimmed as in the previous experiments, resulting in a removal of less than 3 percent of all trials. We compared trials in which the previous high-value target color served as a distractor against those in which the high-value color was absent from the array. We conducted a repeated measures ANOVA with 2 levels of distractor presence and 2 levels of experimental half to examine whether the high-value distractor involuntarily captured attention in either half of trials. Unlike Experiments 1–3 in which target color did not differentially predict the magnitude of reward during the training phase, we observed a significant main effect of distractor presence in Experiment 4, $F(1,15) = 6.77$, $p = .020$, $\eta_p^2 = .311$, 95% CI: 3.99–40.23 (see Figure 5A). The JZS Bayes factor was 0.37. There was a significant main effect of experimental half, $F(1,15) = 6.86$, $p = .019$, $\eta_p^2 = .314$, while the distractor by experimental half interaction, $F(1,15) = 2.00$, $p = .178$, failed to reach statistical significance.

We also analyzed accuracy for distractor present and absent trials across both halves of the test phase with a repeated measures ANOVA. As in Experiments 1 and 3, there was a significant main effect of experimental half, $F(1,15) = 5.52$, $p = .033$, $\eta_p^2 = .269$. The main effect of distractor presence, $F(1,15) = 2.38$, $p = .144$, and the distractor by half interaction, $F(1,15) = 0.14$, $p = .710$, both failed to reach statistical significance (see Figure 5B).

Discussion

Experiment 4 replicated previous demonstrations of value-driven attentional capture (e.g., Anderson et al., 2011b). Our results indicate that previously reward-predictive target

features involuntarily capture attention after only 144 trials of stimulus-reward pairings. Unlike Experiments 1–3 in which we found no evidence of value-driven attentional capture, target stimuli during the training phase of Experiment 4 carried uniquely predictive information regarding the magnitude of upcoming reward on a trial-by-trial basis. Taken together, Experiments 1–4 suggest that associative reward learning, rather than reward-motivated perceptual learning, produces lasting modulations of attentional priority following brief exposure to stimulus-reward pairings.

Between-Experiment Comparison

We first compared the slowing of RT on distractor-present relative to distractor-absent trials in the test phases of Experiments 1–3. In each of these experiments, reward incentive was present during training but there was no predictive relationship between target color and reward magnitude. Since there were no significant interactions between distractor presence and experimental half in any of the experiments, we collapsed across experimental half to compute an attentional capture score (distractor present minus distractor absent RT) for each participant. A univariate ANOVA confirmed that there were no significant differences in attentional capture among Experiments 1–3, $F(2,49) = 1.02$, $p = .367$; therefore, we collapsed the capture scores from these experiments and compared them against capture scores of Experiment 4 where target color did predict reward magnitude during training. Participants in Experiment 4 demonstrated significantly greater attentional capture than participants in Experiments 1–3, $t(66) = 2.01$, $p = .048$. The JZS Bayes factor for the between-experiments comparison was 0.80. Although the homogeneity of variance assumption was not violated in this comparison (Levene's Test: $F = 1.33$, $p = .252$), given the unequal sample sizes, we ran an additional t -test in which homogeneity of variance was not assumed, which produced quantitatively similar results, $t(30.47) = 2.25$, $p = .032$. The magnitude of attentional capture was significantly greater following training in which stimulus color uniquely predicted reward magnitude ($M = 22.11$ ms, $SD = 34.00$) than following training in which stimulus features did not predict trial-by-trial reward magnitude ($M = -1.12$ ms, $SD = 42.06$).

General Discussion

In the current series of experiments, we tested whether motivating visual search with monetary reward incentives was sufficient to modulate the attentional priority of target stimulus features in the absence of prediction-based associative learning. Previous research has demonstrated that previously rewarded target stimuli involuntarily capture attention when presented as distractors in a second experimental task (Anderson et al., 2011a, 2011b; Anderson & Yantis, 2013, 2012). The results of the current experiments suggest that such modulations of attentional priority are not the result of motivationally magnified non-associative learning mechanisms such as perceptual learning. Providing participants with a variable magnitude reward for identifying a predictable feature-defined target or providing rewards of a consistent magnitude for identifying targets defined by a varying feature both left no lasting modulation of attentional priority in a second visual search task. Furthermore, we found no evidence of attentional capture following training in which reward magnitude varied based on the speed of behavioral responses to identify a color-defined target.

Conversely, associative reward learning for stimulus features that predicted reward magnitude was sufficient to modulate the attentional priority of previous target stimuli even after only 144 trials of training exposure. Target colors only contained predictive information about trial-by-trial reward outcome in Experiment 4. Consequently, attentional biases of target colors only developed when reward prediction errors occurred and could be resolved specifically and uniquely by learning the color–reward contingencies.

Our findings suggest that value-based modulations of attentional priority develop rapidly in comparison to automatic attentional selection that arises from perceptual learning (e.g., Kyllingsbaek, Schneider, & Bundesen, 2001; Shiffrin & Schneider, 1977). In the absence of reward, attentional priority for target items increases over the course of thousands of trials of exposure. In the current study, we found that even when individuals were motivated by reward, perceptual learning did not occur under brief training conditions. Conversely, the results of Experiment 4 suggest that the computation of reward prediction error offers one potential mechanism through which the brain rapidly updates attentional priority in response to the changing environment.

In addition to stimuli associated with secondary rewards such as the monetary rewards used in the current study, stimuli associated with primary rewards such as food also receive enhanced attentional priority (Pool, Brosch, Delplanque, & Sander, 2014). Physiological changes in the need for a particular reward influence both behavioral preferences for that reward as well as the attentional priority for stimuli previously associated with the reward. For example, both stimulus-evoked reward signals in the brain and corresponding attentional biases for food-related stimuli are robust when hungry but decrease when satiated (Pool et al., 2014; Siep, Roefs, Roebroek, Havermans, Bonte, & Jansen, 2009). In the present study, we show that associative learning, rather than motivation per se, is responsible for the rapid modulation of attentional priority, which is consistent with these prior reports. However, these studies also suggest a potential role for motivation in gating the influence of associative learning on attentional selection, such that learned value will have a reduced influence on information processing when the associated rewards are devalued. In this way, associative reward learning and motivation may play complementary roles in the control of attention.

Future research is needed to further understand the neural circuitry implicated in value-based modulations of attentional priority. The dopaminergic reward circuitry of the basal ganglia has been a prominent topic of research on associative reward learning (e.g., Cools, 2008; Schultz et al., 1997; Waelti et al., 2001). The interaction of this circuitry with attentional control networks in the brain, however, is more poorly understood. The lateral intraparietal area is one frequently studied map of attentional priority within the primate brain that is sensitive to reward history (Balan & Gottlieb, 2006; Bisley, & Goldberg, 2003; Glimcher, 2003; Gottlieb & Balan, 2010; Loui, Gratton, & Glimcher, 2011; Peck, Jangraw, Suzuki, Efem, & Gottlieb, 2009). Furthermore, transient signals within the dorsal frontoparietal network, including the intraparietal sulcus and the frontal eye fields, as well as a ventral network of areas including the temporoparietal junction and the ventral prefrontal cortex, have been implicated in the involuntary orienting of attention to stimuli that are behaviorally relevant (Corbetta, Patel, & Shulman, 2008; Corbetta & Shulman, 2002;

Kincade, Abrams, Astafiev, Shulman, & Corbetta, 2005; Serences, Shomstein, Leber, Golay, Egeth, & Yantis, 2005; Serences & Yantis, 2006; Serences & Yantis, 2007; Shulman et al., 2009). Future research is necessary to understand how signaling of reward prediction error is implicated in modulations of activity within such frontal and parietal regions involved in the control of attention.

The current study holds important implications for understanding the control of attention in both healthy individuals as well as in patient populations. Disorders of cognitive control such as attention deficit hyperactivity disorder as well as drug addiction have been linked to sensitivity to rewards and corresponding deficiencies within the reward circuitry of the midbrain and prefrontal cortex (Anderson, Faulkner, Rilee, Yantis, & Marvel, 2013; Bédard et al., 2010; Cools, 2008). A greater understanding of the modulatory role of reward learning on attentional control mechanisms may therefore improve knowledge of the etiology of these disorders as well as provide a foundation for improvements in treatment.

In the current study we tested whether motivating visual search with monetary rewards was sufficient to persistently modulate the attentional priority assigned to previous target features. Across four experiments, we found that associative reward learning but not reward-motivated perceptual learning led to an increase in attentional priority for previous target features. Associative reward learning only occurred in the case in which target features carried uniquely predictive information about the magnitude of reward on a trial-by-trial basis. Learned associations between stimulus features and reward history therefore play an important role in the rapid and persistent instantiation of value-based attentional priority through mechanisms of reward prediction.

Acknowledgments

We thank M. Hall and A. Blank for assistance with data collection. This research was supported by U.S. National Institutes of Health grant R01-DA013165 to S.Y., National Science Foundation GRFP DGE-0707427 to A.W.S., and U.S. National Institutes of Health NRSA Fellowship F31-DA033754 to B.A.A.

References

- Anderson BA. A value-driven mechanism of attentional selection. *Journal of Vision*. 2013; 13(3):1–16. 7.
- Anderson BA, Faulkner ML, Rilee JJ, Yantis S, Marvel CL. Attentional bias for non-drug reward is magnified in addiction. *Experimental and Clinical Psychopharmacology*. 2013; 21:499–506. [PubMed: 24128148]
- Anderson BA, Laurent PA, Yantis S. Learned value magnifies salience-based attentional capture. *PLoS ONE*. 2011a; 6:e27926. [PubMed: 22132170]
- Anderson BA, Laurent PA, Yantis S. Value driven attentional capture. *Proceedings of the National Academy of Sciences, USA*. 2011b; 108:10367–10371.
- Anderson BA, Laurent PA, Yantis S. Generalization of value-based attentional priority. *Visual Cognition*. 2012; 20:647–658.
- Anderson BA, Yantis S. Value-driven attentional and oculomotor capture during goal-directed, unconstrained viewing. *Attention, Perception, and Psychophysics*. 2012; 74:1644–1653.
- Anderson BA, Yantis S. Persistence of value-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*. 2013; 39:6–9. [PubMed: 23181684]
- Balan PF, Gottlieb J. Integration of exogenous input into a dynamic salience map revealed by perturbed attention. *Journal of Neuroscience*. 2006; 26:9239–9249. [PubMed: 16957080]

- Bédard A-C, S KP, Cook EH, Fan J, Clerkin SM, Ivanov I, Halperin JM, Newcorn JH. Dopamine transporter gene variation modulates activation of striatum in youth with ADHD. *NeuroImage*. 2010; 53:935–942. [PubMed: 20026227]
- Bisley JW, Goldberg ME. Neuronal activity in the lateral intraparietal area and spatial attention. *Science*. 2003; 299:81–86. [PubMed: 12511644]
- Brainard DH. The Psychophysics Toolbox. *Spatial Vision*. 1997; 10:433–436. [PubMed: 9176952]
- Cools R. Role of dopamine in the motivational and cognitive control of behavior. *The Neuroscientist*. 2008; 14:381–395. [PubMed: 18660464]
- Corbetta M, Patel G, Shulman GL. The reorienting system of the human brain: From environment to theory of mind. *Neuron*. 2008; 58:306–324. [PubMed: 18466742]
- Corbetta M, Shulman GL. Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*. 2002; 3:201–215.
- Della Libera C, Chelazzi L. Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*. 2009; 20:778–784. [PubMed: 19422618]
- Desimone R, Duncan J. Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*. 1995; 18:193–222.
- Glimcher PW. The neurobiology of visual-saccadic decision making. *Annual Review of Neuroscience*. 2003; 26:133–179.
- Gottlieb J, Balan P. Attention as a decision in information space. *Trends in Cognitive Sciences*. 2010; 14:240–248. [PubMed: 20399701]
- Hickey C, Chelazzi L, Theeuwes J. Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*. 2010a; 30:11096–11103. [PubMed: 20720117]
- Hickey C, Chelazzi L, Theeuwes J. Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *PLoS ONE*. 2010b; 5:e14087. [PubMed: 21124893]
- Kamin, LJ. Punishment and Aversive Behavior. Campbell, BA.; Church, RM., editors. New York: Appleton-Century-Crofts; 1969. p. 279-296.
- Kincade JM, Abrams RA, Astafiev SV, Shulman GL, Corbetta M. An event-related functional magnetic resonance imaging study of voluntary and stimulus-driven orienting of attention. *Journal of Neuroscience*. 2005; 25:4593–4604. [PubMed: 15872107]
- Kiss M, Driver J, Eimer M. Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science*. 2009; 20:245–251. [PubMed: 19175756]
- Kristjansson A, Sigurjonsdottir O, Driver J. Fortune and reversals of fortune in visual search: Reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Attention, Perception, & Psychophysics*. 2010; 72:1229–1236.
- Kyllingsbaek S, Schneider WX, Bundesen C. Automatic attraction of attention to former targets in visual displays of letters. *Perception & Psychophysics*. 2001; 12:763–764.
- Louie K, Gratton LE, Glimcher PW. Reward value-based gain control: Divisive normalization in parietal cortex. *Journal of Neuroscience*. 2011; 31:10627–10639. [PubMed: 21775606]
- Mack, A.; Rock, I. In *attentional Blindness*. Cambridge: MIT Press; 1998.
- Miranda AT, Palmer EM. Intrinsic motivation and attentional capture from game like features in a visual search task. *Behavioral Research Methods*. 2013; 46:159–172.
- Pearce JM, Hall G. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*. 1980; 87:532–552. [PubMed: 7443916]
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J. Reward modulates attention independently of action value in posterior parietal cortex. *Journal of Neuroscience*. 2009; 29:11182–11191. [PubMed: 19741125]
- Pessoa L, Engelmann JB. Embedding reward signals into perception and cognition. *Frontiers Neuroscience*. 2010; 4(17):1–8.
- Pool E, Brosch T, Delplanque S, Sander D. Where is the chocolate? Rapid spatial orienting toward stimuli associated with primary rewards. *Cognition*. 2014; 130:348–359. [PubMed: 24389311]
- Qi S, Zeng Q, Ding C, Li H. Neural correlates of reward-driven attentional capture in visual search. *Brain Research*. 2013; 1532:32–43. [PubMed: 23916733]

- Raymond JE, O'Brien JL. Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science*. 2009; 20:981–988. [PubMed: 19549080]
- Rensink RA, O'Regan K, Clark JJ. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*. 1997; 8:368–373.
- Roelfsema PR, van Ooyen A, Watanabe T. Perceptual learning rules based on reinforcers and attention. *Trends in Cognitive Sciences*. 2010; 14:64–71. [PubMed: 20060771]
- Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G. Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin and Review*. 2009; 16:225–237. [PubMed: 19293088]
- Serences JT, Shomstein S, Leber AB, Golay X, Egeth HE, Yantis S. Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychological Science*. 2005; 16:114–122. [PubMed: 15686577]
- Serences JT, Yantis S. Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*. 2006; 10:38–45. [PubMed: 16318922]
- Serences JT, Yantis S. Spatially selective representations of voluntary and stimulus-driven attentional priority in human occipital, parietal, and frontal cortex. *Cerebral Cortex*. 2007; 17:284–293. [PubMed: 16514108]
- Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
- Shiffrin RM, Schneider W. Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*. 1977; 84:127–190.
- Shulman GL, Astafiev SV, Franke D, Pope DLW, Snyder AZ, McAvoy MP, Corbetta M. Interaction of stimulus-driven reorienting and expectation in ventral and dorsal frontoparietal and basal ganglia-cortical networks. *Journal of Neuroscience*. 2009; 29:4392–4407. [PubMed: 19357267]
- Siep N, Roefs A, Roebroek A, Havermans R, Bonte ML, Jansen A. Hunger is the best spice: An fMRI study of the effects of hunger on attention, hunger and calorie content on food reward processing in the amygdala and orbitofrontal cortex. *Behavioural Brain Research*. 2009; 198:149–158. [PubMed: 19028527]
- Theeuwes J, Belopolsky AV. Reward grabs the eye: Oculomotor capture by rewarding stimuli. *Vision Research*. 2012; 74:80–85. [PubMed: 22902641]
- Wang L, Yu H, Zhou X. Interaction between value and perceptual salience in value-driven attentional capture. *Journal of Vision*. 2013; 13(3):1–13. 5.
- Waelti P, Dickinson A, Schultz W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature*. 2001; 412:43–48. [PubMed: 11452299]
- Watanabe T, Sasaki Y, Nanez J. Perceptual learning without perception. *Nature*. 2001; 413:844–847. [PubMed: 11677607]
- Yantis S, Egeth HE. On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*. 1999; 3:661–676. [PubMed: 10385983]
- Yantis S, Johnston JC. On the locus of visual selection: Evidence from focused attention tasks. *Journal of Experimental Psychology: Human Perception and Performance*. 1990; 16:135–149. [PubMed: 2137515]

Figure 1.

A) Training phase paradigm. In Experiments 1 and 3, participants searched for either a red circle or a green circle on every trial and reported the orientation of the corresponding line segment. In Experiments 2 and 4, target items during the training phase could be either a red or a green circle on a given trial. Participants earned a monetary reward for each correct response made during the training phase, although how the magnitude of reward was determined varied across experiments. B) Test phase paradigm. Participants reported the

orientation of a line segment within the unique shape in each array. On half of the trials, a previously rewarded distractor appeared in the search array.

Figure 2.

Behavioral results for Experiment 1. Response time (A) and accuracy (B) for distractor present and distractor absent trials across both halves of the experiment. Error bars denote 1 within-subject SEM.

Figure 3.

Behavioral results for Experiment 2. Response time (A) and accuracy (B) for distractor present and distractor absent trials across both halves of the experiment. Error bars denote 1 within-subject SEM.

Figure 4.

Behavioral results for Experiment 3. Response time (A) and accuracy (B) for distractor present and distractor absent trials across both halves of the experiment. Error bars denote 1 within-subject SEM.

Figure 5.

Behavioral results for Experiment 4. Response time (A) and accuracy (B) for distractor present and distractor absent trials across both halves of the experiment. Error bars denote 1 within-subject SEM.