Published in final edited form as:

Psychon Bull Rev. 2015 June ; 22(3): 750–756. doi:10.3758/s13423-014-0724-0.

Value-Driven Attentional Priority is Context Specific

Brian A. Anderson

Johns Hopkins University

Abstract

Attention is automatically drawn to stimulus features previously associated with reward, a phenomenon referred to as *value-driven attentional capture*. To date, value-driven attentional capture has been studied exclusively by manipulating stimulus–reward contingencies in an experimental setting. Although practical and intuitively appealing, this approach poses theoretical challenges to understanding the broader impact of reward on attention in everyday life. These challenges arise from the fact that associative learning between a given visual feature and reward is not limited to the context of an experiment, yet such extra-experimental learning is completely ignored in studies of value-driven attention. How is it, then, that experimentally established reward associations even influence attention, seemingly overshadowing any prior learning about particular features and reward? And how do the effects of this experimental learning persist over long periods of time, in spite of all the intervening experiences outside of the lab that might interfere with this learning? One potential answer to these questions is that value-driven attention is context specific, such that different contexts evoke different value priors that the attention system uses to assign priority. In the present study, I directly test this hypothesis. The results show that the same stimulus feature either does or does not capture attention, depending on whether it has been rewarded specifically in the context within which it appears. The findings provide insight into how multiple reward structures can efficiently guide attention with minimal interference.

Keywords

selective attention; reward learning; contextual learning

Attention determines which among multiple stimuli encountered in an environment are accessible to capacity-limited cognitive processes such as those governing decision-making and action selection. Therefore, in order to promote survival and well-being, it is important that organisms attend to stimuli that are associated with reward (Anderson, 2013). Consistent with this, reward-related stimuli are preferentially selected and processed (e.g., Della Libera & Chelazzi, 2009; Kiss, Driver, & Eimer, 2009; Krebs, Boehler, & Woldorff, 2010; Lee & Shomstein, 2014; Muhle-Karbe & Krebs, 2012; Raymond & O'Brien, 2009; Rutherford, O'Brien, & Raymond, 2010; Serences, 2008). Although attention can be strongly influenced by current goals (e.g., Folk, Remington, & Johnston, 1992), and goals often reflect what we expect is currently of value (e.g., Maunsell, 2004), reward-associated stimuli may appear unexpectedly or during the performance of an unrelated task. By

Author Manuscript

automatically orienting attention to stimuli previously associated with reward, potentially rewarding opportunities would be given a consistently high priority relative to other competing sources of information. Recent findings support such an automatic role for reward in driving attention. The experience of reward strongly primes attention to the rewarded stimulus (e.g., Della Libera & Chelazzi, 2006; Hickey, Chelazzi, & Theeuwes, 2010), and stimulus features previously associated with reward involuntarily capture attention when presented as task-irrelevant distractors (e.g., Anderson, Laurent, & Yantis, 2011a, 2011b; Anderson & Yantis, 2012, 2013; Qi, Zeng, Ding, & Li, 2013).

Author Manuscript

The idea that organisms would develop a bias to automatically attend to reward-associated visual features has intuitive appeal, and experimental support for this claim is strong. Attention either is or is not biased to select a particular feature based on whether it was associated with reward in an earlier part of an experiment (e.g., Anderson et al., 2011a, 2011b; Qi et al., 2013), and stimulus features previously associated with high reward capture attention more strongly than stimulus features previously associated with comparatively lower reward (e.g., Anderson et al., 2011a, 2012; Anderson & Yantis, 2013; Failing & Theeuwes, 2014; Theeuwes & Belopolsky, 2012). As straightforward as these demonstrations of reward's influence on attention are, how these experimental effects of reward generalize to situations beyond the laboratory setting is both unclear and controversial.

Author Manuscript

Experimentally manipulating task demands (e.g., Folk et al., 1992) or the physical conspicuity of objects (e.g., Theeuwes, 1992) is intended to provide insight into how attention selects goal-relevant and salient objects in the real world, respectively. In the same vein, experimental manipulations of a reward structure are intended to serve as an analog for how reward influences attention in everyday encounters with stimuli. However, the history component of reward learning creates an important complexity in making this generalization, as associative reward learning is not limited to the context of a laboratory setting. Participants constantly experience the same visual features that are experimentally manipulated both before and after this manipulation is introduced, and these experiences all involve the presence or absence of different rewarding outcomes. A complete theory of value-based attention needs to account for how experimentally created reward structures influence attention in spite of competing reward structures between the same visual features and reward.

Author Manuscript

More specifically, two general findings together pose an important theoretical puzzle when considering experimental effects of learned value on attention. A particular stimulus feature will capture attention in an experiment based on the reward it has been associated with in that experiment (e.g., Anderson et al., 2011a, 2011b, 2012; Anderson & Yantis, 2013; Failing & Theeuwes, 2014; Theeuwes & Belopolsky, 2012; Qi et al., 2013). Yet, such reward learning only accounts for a tiny fraction of the experience that each participant has with those particular features. If the attention system simply aggregated the value associated with each visual feature across life experiences, then experimental associations between features and reward should be completely overshadowed by prior learning that had occurred before the experiment even began. This is clearly not the case.

One way of accounting for experimental effects of reward on attention is to hypothesize that the attention system strongly prioritizes the most recent experiences with reward in assigning priority. By this account, the dominance of the experimental reward contingencies over prior learning outside of the experiment simply reflects the recency of the experimental learning. This account does not stand up to other sources of evidence, however. Value-driven attentional priority is slow to extinguish in the absence of reward (e.g., Anderson et al., 2011b, 2012; Anderson & Yantis, 2012; Failing & Theeuwes, 2014; Theeuwes & Belopolsky, 2012). Most prominently, attentional capture by stimulus features previously associated with reward can be observed days and even months after the experimental learning has occurred (Anderson et al., 2011b; Anderson & Yantis, 2013), in spite of all the visual experiences that occur with these features between experimental testing sessions.

How can value-driven attentional priorities be simultaneously persistent yet highly robust to other competing sources of learning occurring outside of the experiment? Is there something very powerful about how reward is being experimentally manipulated (e.g., over several hundred trials in a brief timeframe) that outcompetes everyday learning? Is the attention system very slow at updating established biases, or do these biases otherwise reflect a habit that is resistant to contrary evidence from reward feedback? Or is there rather something highly specific about when and how experiences with reward modulate attention? In the present study, I examine the hypothesis that multiple sets of value priors can guide attention, depending on the current context. More specifically, I test the idea that a context evokes a set of learned stimulus–reward associations that have been experienced specifically within that context, and it is these contextually specific value associations that bias attention.

Although a non-strategic influence of context on value-driven attention would provide an elegant solution to this puzzle, there is currently no evidence that the influence of learned stimulus–reward associations on attention can be so specific. In fact, prior studies of value-driven attention have tended to support the opposite principle: generalization of learning. Attentional biases for a stimulus feature learned in one task (visual search task) can affect attentional processing in a different task (flankers task) even though the specific objects, task set, and response requirements are different (Anderson et al., 2012). In the typical paradigm used to examine value-driven attention, the color distractor can appear as a different shape than that of the previously reward-associated target (diamond vs circle: e.g., Anderson et al., 2011a, 2011b; Anderson & Yantis, 2012). In all of these prior studies, however, such contextual information never predicted anything about available reward, which was always associated with a single feature dimension (color).

In the reported experiment, participants experienced two different reward structures in which the same visual feature either was or was not associated with reward, depending on the context in which it was presented. For example, red stimuli were rewarded in context A but not context B, while the opposite was true of green stimuli. Context was manipulated through the presentation of a background scene upon which the search array was rendered. Once these contextually dependent reward structures had been experienced in a training phase, irrelevant distractors possessing the previously rewarded features were presented in an unrewarded test phase. Of interest was whether value-driven attentional capture by a

given feature would only occur when that feature was presented in the context within which it was previously rewarded.

Methods

Participants

Thirty participants (18-33 years of age, mean = 22.1y, 21 female) were recruited from the Johns Hopkins University community. All reported normal or corrected-to-normal visual acuity and normal color vision.

Apparatus

A Mac Mini equipped with Matlab software and Psychophysics Toolbox extensions (Brainard, 1997) was used to present the stimuli on an Asus VE247 monitor. The participants viewed the monitor from a distance of approximately 70 cm in a dimly lit room. Manual responses were entered using a standard keyboard.

Training Phase

Stimuli—Each trial consisted of the presentation of a context scene upon which a fixation display and a search array were subsequently presented, followed by a feedback display (Figure 1A). The context scene consisted of a black-and-white picture of a forest or city street (as in, e.g., Cosman & Vecera, 2013), which remained on screen throughout the fixation display and search array. The fixation display contained a white fixation cross ($0.7^\circ \times 0.7^\circ$ visual angle) presented in the center of the screen, and the search array consisted of the fixation cross surrounded by six colored circles (each $2.5^\circ \times 2.5^\circ$) presented along an imaginary circle with a radius of 5° . The fixation cross and each colored circle was presented within a black box to increase contrast with the background scene. All stimuli were presented on a grey background.

The target was defined as the red or green circle, exactly one of which was presented on each trial. The color of each nontarget circle was drawn from the set {blue, cyan, pink, orange, yellow, white} without replacement. Inside the target circle, a white bar was oriented either vertically or horizontally, and inside each of the nontarget circles, a white bar was tilted at 45° to the left or to the right (randomly determined for each nontarget). The feedback display indicated the amount of monetary reward earned on the current trial, as well as the total accumulated reward.

Design—The target appeared in each of the six possible locations equally-often. For one background scene (counterbalanced across participants), red would always be followed by a 10¢ reward on correct trials while green would always be followed by a 0¢ reward, while for the other background scene it was green that was rewarded 10¢. The background scene was the forest on half of the trials and the city street on the other half. Trials were presented in a random order.

Procedure—The training phase consisted of 480 trials, which were preceded by 48 practice trials during which no background scenes were presented. Each trial began with the

presentation of the background scene for 1500 ms, after which the fixation cross appeared and remained on screen for randomly varying interval of 400, 500, or 600 ms. The search array then appeared and remained on screen until a response was made or 1000 ms had elapsed, after which the trial timed out. The search array was followed by just the background scene again for another 500 ms, which then disappeared to reveal a blank grey screen for 500 ms. The trial concluded with the reward feedback display for 1500 ms, which was followed by a blank 1000 ms inter-trial interval (ITI).

Participants made a forced-choice target identification by pressing the “z” and the “m” keys for the vertically- and horizontally-orientated bars within the targets, respectively. Correct responses were followed by monetary reward feedback in which either 10¢ or 0¢ was added to the participant’s total earnings, depending on the relationship between target color and background context. Incorrect responses were followed by feedback in which the word “Incorrect” was presented in place of the monetary increment, and responses that were too slow (i.e., no response before the trial timed out) were followed by a 500 ms 1000 Hz tone and no monetary increment (i.e., just the total earnings were presented in the feedback display).

Test Phase

Stimuli—Each trial consisted of the presentation of a context scene upon which a fixation display and a search array were subsequently presented (Figure 1B). The context scene was presented prior to the search array in order to ensure adequate processing of the scene was possible before localization of the target. The six shapes comprising the search array now consisted of either a diamond among circles or a circle among diamonds, and the target was defined as the unique shape. On a subset of the trials, one of the nontarget shapes was rendered in the color of a formerly reward-associated target from the training phase (referred to as the *valuable distractor*); the target shape was never the color of a target from the training phase.

Design—Target identity, target location, distractor identity, and distractor location were fully crossed and counterbalanced separately within each context, and trials were presented in a random order. The context was the forest and city street equally-often. Valuable distractors were presented on 50% of the trials within each context, half of which were red and half of which were green.

Procedure—Participants were instructed to ignore the color of the shapes and to focus on identifying the oriented bar within the unique shape using the same orientation-to-response mapping. The test phase consisted of 240 trials, which were preceded by 32 practice (distractor absent) trials that did not include the background scenes. The search array was followed immediately by non-reward feedback (the word “Incorrect”) for 1000 ms in the event of an incorrect response (this display was omitted following a correct response) and then by a 1000 ms ITI; no monetary rewards were given in the test phase. Trials timed out after 1500 ms. As in the training phase, if the trial timed out, the computer emitted a 500 ms 1000 Hz tone. Upon completion of the experiment, participants were paid the cumulative reward they had earned in the training phase.

Exit Question

At the conclusion of the test phase, participants were asked to select which of several statements they believed best described the reward contingencies in the training phase. Participants were given six options, one of which reflected the actual relationship between target color, background scene, and reward (see Appendix).

Data Analysis

Only correct responses were included in all analyses of RT, and RTs more than three SDs above or below the mean of their respective condition for each participant were trimmed. In the training phase, trials were divided based on whether the target color appeared within the context in which it was associated with reward. The same distinction was made for distractor-present trials in the test phase with respect to reward history from training. Note that both target/distractor colors and both contexts were represented in each target/distractor condition; what differed was the specific pairings of color and context based on the reward structure.

Results

Training Phase

Participants were no faster, $t(29) = 0.79$, $p = .434$, or more accurate, $t(29) = 0.60$, $p = .554$, to report a target when its color was associated with reward in the current context compared to when its color was not associated with reward in that context (551 and 90.9% vs 553 ms and 90.6%, respectively). This suggests that participants searched for the two target colors with roughly equal priority across contexts. There was also no switch cost in target identification associated with a change in the background scene across consecutive trials, mean switch cost = -1 ms, $t(29) = -0.64$, $p = .526$, suggesting that top-down goals did not change with the context.

Test Phase

A repeated-measures ANOVA on mean RT with distractor condition (absent, present in unrewarded context, present in rewarded context) as a factor revealed a main effect, $F(2,58) = 3.32$, $p = .043$, $\eta^2 = .103$ (see Figure 2). A planned comparison revealed that when a color distractor appeared in a context within which it was previously rewarded, RTs were slower than when it was absent or appeared in a context in which it was never rewarded, $t(29) = 2.21$, $p = .035$, $d = .40$. By contrast, a distractor had no influence on RT when it appeared in the context within which it was never rewarded compared to distractor-absent trials, $t(29) = -.047$, $p = .643$. As in the training phase, there was no evidence of a switch cost associated with a change in context, mean switch cost = -6 ms, $t(29) = -1.08$, $p = .291$. Accuracy did not differ by distractor condition, $F(2,58) = 0.41$, $p = .668$ (90.4%, 90.2%, and 89.6% across the absent, unrewarded context, and rewarded context distractor conditions, respectively).

Exit Question

Six of the thirty participants selected the correct reward contingency, which was only one more than what would be expected from random guessing (1/6). This suggests that the

learning of the color–context contingencies governing reward outcome in the present study was largely implicit.

Discussion

The present study demonstrates that the stimulus–reward associations that bias attention are context specific. Participants were equally rewarded for identifying red and green targets over the course of the entire training phase, but the identity of the background scene predicted whether red or green would be rewarded on a given trial. One possibility is that participants would assign value to red and green equally in this situation, taking into account only the target features while either ignoring or generalizing across contextual information. The results, however, tell a different story. Instead, the very same feature captured attention in one context but not another based on its contextually specific reward history.

My findings provide a mechanism for how reward learning can automatically yet efficiently guide attention across a broad range of diverse visual environments. The experience of a particular context evokes its own unique set of value priors, which can be independently updated with learning and automatically bias attention when activated. In this way, the attention system can benefit from past learning in one situation with minimal interference from important but potentially unrelated learning occurring in a different situation.

Although the idea that the attention system can make use of contextual information in guiding selection is itself unsurprising, one possibility is that such contextual modulation operates solely by providing a cue to voluntarily update task-specific goals and expectations. In the present study, I provide evidence for a much more automatic and implicit influence of contextually-specific representations on attention. There was no evidence for switch costs tied to a change in context, which suggests that changes in context were not accompanied by changes in task-specific goals or voluntary search strategies. Participants were also near chance in reporting the contextually-dependent reward contingencies in a forced-choice assessment. Importantly, in the test phase, contextual information was completely task-irrelevant and rewards were no longer available. Collectively, my findings suggest that contextual modulation of value-based attentional priority is itself a fairly automatic cognitive process that does not require strategic cognitive control.

Prior studies of value-driven attentional capture have demonstrated that attentional biases for reward-associated features are capable of generalizing across stimuli and contexts in certain situations (e.g., Anderson et al., 2011a, 2011b, 2012). One potentially important difference between these prior studies and the present study is that in the prior studies, the high-value color was always the same on every trial during training. One means of reconciling the outcomes of these studies with the present findings is to assume that contextual distinctions are only implemented when such distinctions themselves provide predictive information about reward. This hypothesis is supported by recent findings emphasizing the importance of reward prediction errors in the establishment of value-based attentional biases (Sali et al., in press). In this sense, the attention system defaults to generalizing value representations across context, with the potential to leverage prior learning in new situations, but draws a

distinction between contexts when existing value representations prove to be a poor predictor of reward in a particular context.

The findings from the present study also provide insight into cognitive impairments characteristic of addiction. Drug-dependent individuals show a substantially elevated magnitude of value-driven attentional capture by stimuli associated with non-drug reward (Anderson et al., 2013), suggesting that attentional biases for drug-related stimuli (see Field & Cox, 2008, for a review) might in part reflect a more general sensitivity to reward's influence on attention. It is also known that the desire to consume a substance of abuse and associated relapse can be powerfully evoked by a context in which the drug reward was previously experienced (e.g., Robinson & Berridge, 1993). That value-driven attention is similarly modulated by context further suggests that abnormality in this basic cognitive mechanism may contribute to addiction.

Previously, I argued that automatically attending to reward-associated stimuli, in spite of competing goals, could confer adaptive benefits (Anderson, 2013). While there are several reasons why value-driven attention could be conceived as adaptive, one potential weakness of this argument is that value-driven attention has the potential to consistently misguide attention whenever learned stimulus–reward associations do not reflect the reward structure of the current environment. Given the variety of visual environments we experience in everyday life, this issue requires serious consideration. The findings from the present study provide important insights into how the attention system overcomes this challenge of overgeneralization. By automatically and implicitly modulating currently activated value priors based on contextual information, the attention system flexibly leverages only the prior learning that is maximally reflective of current reward prospects when assigning stimulus priority.

Acknowledgements

This research was supported by NIH grants F31-DA033754 and R01-DA013165

Appendix

Which option do you believe best describes the part of the experiment in which you were earning money (please choose only one):

1. The red circle was generally worth more than the green circle regardless of what the background was
2. The green circle was generally worth more than the red circle regardless of what the background was
3. The two circles were worth the same overall, but one color was worth more when it appeared on the forest background and the other was worth more when it appeared on the city background
4. Both color circles were generally worth more when presented on the forest background

5. Both color circles were generally worth more when presented on the city background
6. How much money I received was random and unrelated to the background

References

- Anderson BA. A value-driven mechanism of attentional selection. *Journal of Vision*. 2013; 13(3):1–16.
- 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(11):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]
- Brainard DH. The psychophysics toolbox. *Spatial Vision*. 1997; 10:433–436. [PubMed: 9176952]
- Cosman JD, Vecera SP. Context-dependent control over attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*. 2013; 39:836–848. [PubMed: 23025581]
- Della Libera C, Chelazzi L. Visual selective attention and the effects of monetary reward. *Psychological Science*. 2006; 17:222–227. [PubMed: 16507062]
- 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]
- Failing MF, Theeuwes J. Exogenous visual orienting by reward. *Journal of Vision*. 2014; 14(5):1–9.
- Hickey C, Chelazzi L, Theeuwes J. Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*. 2010; 30:11096–11103. [PubMed: 20720117]
- 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]
- Krebs RM, Boehler CN, Woldorff MG. The influence of reward associations on conflict processing in the Stroop task. *Cognition*. 2010; 117:341–347. [PubMed: 20864094]
- Lee J, Shomstein S. Reward-based transfer from bottom-up to top-down search tasks. *Psychological Science*. 2014; 25:466–475. [PubMed: 24335604]
- Maunsell JHR. Neuronal representations of cognitive state: Reward or attention? *Trends in Cognitive Science*. 2004; 8:261–265.
- Muhle-Karbe PS, Krebs RM. On the influence of reward on action-effect binding. *Frontiers in Psychology*. 2012; 3
- 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]
- Robinson TE, Berridge KC. What is the role of dopamine in reward: hedonics, learning, or incentive salience? *Brain Research Reviews*. 1993; 18:247–291. [PubMed: 8401595]
- Rutherford HJV, O'Brien JL, Raymond JE. Value associations of irrelevant stimuli modify rapid visual orienting. *Psychonomic Bulletin and Review*. 2010; 17:536–542. [PubMed: 20702874]

- Sali AW, Anderson BA, Yantis S. The role of reward prediction in the control of attention. *Journal of Experimental Psychology: Human Perception and Performance*. in press.
- Serences JT. Value-based modulations in human visual cortex. *Neuron*. 2008; 60:1169–1181. [PubMed: 19109919]
- Theeuwes J. Perceptual selectivity for color and form. *Perception and Psychophysics*. 1992; 51:599–606. [PubMed: 1620571]
- Theeuwes J, Belopolsky AV. Reward grabs the eye: oculomotor capture by rewarding stimuli. *Vision Research*. 2012; 74:80–85. [PubMed: 22902641]

Sequence and time course of trial events. **(A)** Targets were defined as the red or green circle, and participants reported the identity of the line segment inside of the target (vertical or horizontal) with a key press. The background scene (forest or city street) predicted whether a particular target color would be rewarded for a correct response. **(B)** During the test phase, the target was defined as the unique shape. On half of the trials, one of the non-target items—the distractor—was rendered in the color of a formerly rewarded target, presented equally-often on each background scene. No reward feedback was provided.

Figure 2.
Mean response time by distractor condition in the test phase. Error bars reflect the within-subjects S.E.M.