Cognition 178 (2018) 26–36

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

Mechanisms of value-learning in the guidance of spatial attention

Brian A. Anderson*, Haena Kim

Texas A&M University, United States

ARTICLE INFO

Keywords:
Selective attention
Spatial attention
Attentional capture
Reward learning
Real world scenes

ABSTRACT

The role of associative reward learning in the guidance of feature-based attention is well established. The extent to which reward learning can modulate spatial attention has been much more controversial. At least one demonstration of a persistent spatial attention bias following space-based associative reward learning has been reported. At the same time, multiple other experiments have been published failing to demonstrate enduring attentional biases towards locations at which a target, if found, yields high reward. This is in spite of evidence that participants use reward structures to inform their decisions where to search, leading some to suggest that, unlike feature-based attention, spatial attention may be impervious to the influence of learning from reward structures. Here, we demonstrate a robust bias towards regions of a scene that participants were previously rewarded for selecting. This spatial bias relies on representations that are anchored to the configuration of objects within a scene. The observed bias appears to be driven specifically by reinforcement learning, and can be observed with equal strength following non-reward corrective feedback. The time course of the bias is consistent with a transient shift of attention, rather than a strategic search pattern, and is evident in eye movement patterns during free viewing. Taken together, our findings reconcile previously conflicting reports and offer an integrative account of how learning from feedback shapes the spatial attention system.

1. Introduction

The role of an observer's goals (*top-down* factors) and the physical salience of objects (*bottom-up* factors) in the control of attention have been well established and serve as the foundation for prominent models of selective attention (e.g., Corbetta & Shulman, 2002; Desimone & Duncan, 1995; Theeuwes, 2010; Wolfe, Cave, & Franzel, 1989). More recently, it has been argued that this dichotomy cannot explain the role of selection history in the control of attention, which appears to be both non-strategic and independent of the physical salience of stimuli (Awh, Belopolsky, & Theeuwes, 2012). In this context, an important component of selection history has been argued to reflect associative reward learning, with objects previously associated with reward automatically capturing visual attention (Anderson, 2013).

The role of associative reward learning in the control of attention was initially demonstrated using stimuli defined by shape, with results showing that stimulus competition was biased for or against different shapes based on whether observers were rewarded for selecting or ignoring them, respectively (Della Libera & Chelazzi, 2009). This bias carried over into extinction, suggesting that it was non-strategic. Powerful evidence for the unique role of associative reward learning in the control of attention was provided by a study in which task-irrelevant distractors were rendered in a color that had been predictive of reward during a prior training phase. These distractors were not physically salient (less so than the target), and the color of stimuli was known by participants to be completely irrelevant to the task. Nevertheless, attention was automatically captured by the previously reward-associated colors, suggesting a distinct mechanism of attentional control that has been referred to as *value-driven attention* (Anderson, Laurent, & Yantis, 2011).

Many subsequent studies have adopted this approach of associating stimulus features (often color) with reward and then presenting the previously reward-associated features as distractors, replicating and extending the phenomenon of value-driven attention (e.g., Anderson, 2016a, 2016c; Anderson, Folk, Garrison, & Rogers, 2016; Anderson, Laurent, & Yantis, 2012; Anderson & Yantis, 2012, 2013; Anderson, Kuwabara, et al., 2016; Failing & Theeuwes, 2014; Le Pelley, Pearson, Griffiths, & Beesley, 2015; Mine & Saiki, 2015; Moher, Anderson, & Song, 2015; Pool, Brosch, Delplanque, & Sander, 2014; see Anderson, 2016b, for a recent review). Attention has been successfully trained to favor a variety of stimulus features, ranging from specific colors (e.g., Anderson et al., 2011) and orientations (Laurent, Hall, Anderson, & Yantis, 2015; Lee & Shomstein, 2014) to shapes (Della Libera & Chelazzi, 2009) and even object categories (Hickey & Peelen, 2015). In addition to feature-based attention, object-based attention also appears to be strongly modulated by associative reward learning (Lee &

* Corresponding author at: Texas A&M University, Department of Psychology, 4235 TAMU, College Station, TX 77843-4235, United States.
E-mail address: brian.anderson@tamu.edu (B.A. Anderson).

https://doi.org/10.1016/j.cognition.2018.05.005
Received 13 October 2017; Received in revised form 24 April 2018; Accepted 5 May 2018

Shomstein, 2013; Shomstein & Johnson, 2013).

Few studies have probed the influence of associative reward learning on the control of space-based attention, where reward is not predicted by a particular stimulus feature but rather by where in space attention needs to be directed in order to receive high reward. Value-driven attentional biases for a stimulus feature can be modulated by spatial information, such that the bias is specific to when that feature appears in a region of space in which it was rewarded (Anderson, 2015a). More purely space-based attentional biases have also been shown to be modulated by reward. When a high reward is received for identifying a target appearing in a given location, this location is prioritized on the subsequent trial (Hickey, Chelazzi, & Theeuwes, 2014), extending earlier evidence for reward-mediated priming of stimulus color (Hickey, Chelazzi, & Theeuwes, 2010). A more enduring bias towards a previously rewarded location was demonstrated following a multi-day training protocol in which participants performed a difficult visual search for alphanumeric among nonalphanumeric characters. Which of eight possible stimulus positions a searched-for character appeared in predicted the amount of money earned for reporting that target on a given trial. During extinction, participants were more likely to report a target appearing in a previously high-value location, specifically when two targets were simultaneously presented that competed for attention (Chelazzi et al., 2014).

At least two cases have been reported in which a spatial reward manipulation failed to produce any evidence for an enduring attentional bias, or even an attentional bias during the period in which the reward structure was currently in place. In each study, multiple experiments were conducted in which participants searched for a "T" among offset "L" distractors (Jiang, Li, & Remington, 2015; Won & Leber, 2016). When the target appeared anywhere within a particular quadrant of the screen, it was much more likely to yield a high reward if correctly identified. Under a variety of conditions, including conditions of time pressure in which participants should be highly motivated to preferentially search the high-value quadrant in order to maximize rewards, no measurable spatial attention bias was observed (Jiang et al., 2015; Won & Leber, 2016). Furthermore, under similar conditions in which participants were asked to instead choose a particular stimulus rather than perform visual search, robust spatial preferences for highly rewarded locations were observed. These results suggest that spatial reward can readily influence choice behavior (Won & Leber, 2016), but seemingly not the allocation of attention during visual search (Jiang et al., 2015; Won & Leber, 2016).

It is important to note that space-based attentional biases are robustly influenced by a different source of selection history using the same experimental paradigm that failed to show value-driven biases. Specifically, participants are much faster to report targets in locations that more frequently contained targets in prior trials, even well after such biased probabilities are no longer in place (e.g., Jiang & Swallow, 2013; Jiang, Swallow, Rosenbaum, & Herzig, 2013; Jiang et al., 2015; Won & Leber, 2016). Similarly, targets can be found more quickly when the position of the target is consistently predicted by the spatial configuration of non-targets (e.g., Chun & Jiang, 1998, 2003), a phenomenon termed *contextual cueing*. That a different form of selection history can so robustly bias spatial attention in this paradigm argues against a general insensitivity of the paradigm to the ability to detect a learned spatial attention bias.

These repeated failures to observe reliable effects of reward history on the allocation of spatial attention during visual search have naturally led to skepticism concerning whether principles of value-driven attention extend to the spatial domain (Jiang et al., 2015; Won & Leber, 2016). Indeed, it has been suggested that evolutionary pressures imposed by naturally occurring reward structures might strongly favor feature-reward pairings over space-reward pairings, rendering influences of the reward system on spatial attention phylogenetically implausible (Won & Leber, 2016). We would argue that this is a fair point, in the context of how space is defined in these studies.

In traditional visual search paradigms, including those used by Jiang et al. (2015) and Won and Leber (2016), space is defined in a highly abstract manner: a region of a blank computer screen. In fact, there is no clear anchor point for defining where one region would end and another begin, apart from the borders imposed by the edges of the monitor. Such a highly abstract notion of space is unlikely to engage the spatial representations one might use to guide search for a valued item based on learning history, such as where ice cream tends to be stored in the freezer. In this case, the valued location is defined in the context of the spatial arrangement of objects in the scene (e.g., the position of the freezer relative to other objects in the room, and which section of which shelf when looking inside the freezer). The spatial information provided by real-world scenes can serve as the basis for contextual cueing of target position (e.g., Brockmole & Henderson, 2006a, 2006b), suggesting a rich source of spatial guidance, although contextual cueing is also evident with the more abstract stimulus displays that have failed to produce evidence of value-based attentional biases (Jiang et al., 2015; Won & Leber, 2016).

Perhaps information pertaining to the spatial layout and arrangement of objects in a scene is useful for guiding spatial attention on the basis of reward history, which might help explain the apparent discrepancy between Chelazzi et al. (2014) on the one hand, and Jiang et al. (2015) and Won and Leber (2016) on the other hand. Chelazzi et al. (2014) only found an effect of reward when two targets simultaneously competed for attention, where reward was not only tied to the absolute spatial location of the targets but also to their relative positions. With the aim of reconciling these conflicting reports, in the present study, we examined the role of value learning in the context of real-world scenes, both scenes containing a rich array of objects with a consistent spatial arrangement and scenes containing no objects (textures).

2. Experiment 1

In Experiment 1, participants were first trained to associate a specific region of multiple different scenes with monetary reward. On each trial, one of eight scenes remained on the screen until participants clicked on a pixel within the scene using the mouse cursor. Participants were instructed that they would be rewarded for each click, and that the amount of reward received depended on where they clicked. Unbeknownst to the participant, for each of the eight scenes, clicking in one quadrant would always yield more reward than clicking in any other quadrant, and clicking in the center of that quadrant was associated with the best possible payout. Each quadrant served as the high-value quadrant equally-often across scenes, requiring that participants' memory for high-value locations be context-specific, rather than reflect a global bias towards one particular region of the computer screen. In the feature domain, value-driven attentional capture can exhibit contextual specificity (Anderson, 2015b). In the present study, the scene context manipulation demanded that participants take into account the unique spatial layout of each scene.

In a subsequent test phase, participants performed visual search for a side-ways "T" among three upright or upside down "T" distractors, with one search item appearing in the center of each quadrant on the screen. The previously presented scenes were used as the background and were irrelevant to the task, and participants were informed that they could neither earn nor lose money in this task. To the degree that spatial attention is automatically oriented towards previously high-value locations within a scene, participants should be significantly faster to report the target when it appears within a previously high-value location, which would be reflected in a robust validity effect.

2.1. Methods

2.1.1. Participants

Thirty-six participants were recruited from the Texas A&M