Modulation of spatial attention by goals, statistical learning, and monetary reward

Yuhong V. Jiang¹ · Li Z. Sha¹ · Roger W. Remington^{1,2}

Published online: 24 June 2015
© The Psychonomic Society, Inc. 2015

Abstract This study documented the relative strength of task goals, visual statistical learning, and monetary reward in guiding spatial attention. Using a difficult T-among-L search task, we cued spatial attention to one visual quadrant by (i) instructing people to prioritize it (goal-driven attention), (ii) placing the target frequently there (location probability learning), or (iii) associating that quadrant with greater monetary gain (reward-based attention). Results showed that successful goal-driven attention exerted the strongest influence on search RT. Incidental location probability learning yielded a smaller though still robust effect. Incidental reward learning produced negligible guidance for spatial attention. The 95 % confidence intervals of the three effects were largely nonoverlapping. To understand these results, we simulated the role of location repetition priming in probability cuing and reward learning. Repetition priming underestimated the strength of location probability cuing, suggesting that probability cuing involved long-term statistical learning of how to shift attention. Repetition priming provided a reasonable account for the negligible effect of reward on spatial attention. We propose a multiple-systems view of spatial attention that includes task goals, search habit, and priming as primary drivers of top-down attention.

Keywords Spatial attention · Reward-based attention · Goal-driven attention · Probability cuing · Visual search · Visual statistical learning · Repetition priming

✉ Yuhong V. Jiang
jiang166@umn.edu

¹ Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA

² University of Queensland, Brisbane, Australia

Introduction

Human errors are a major source of disasters and accidents in driving, aircraft-control, and other human-machine interfaces. These errors largely arise from a fundamental limit in the number of tasks and stimuli that people can attend to simultaneously (Pashler, 1998). A critical mechanism for reducing information overload is selective attention, which allows humans to prioritize relevant or salient information for further processing (Johnston & Dark, 1986). Decades of psychological and neuroscience research has revealed two major drivers of selective attention: top-down goals and bottom-up salience (Bisley & Goldberg, 2010; Desimone & Duncan, 1995; Egeth & Yantis, 1997; Fecteau & Munoz, 2006; Wolfe, 2007). Information that facilitates one's task goals or that is perceptually salient receives greater priority, ensuring that perception and action are governed by significant, rather than random, input. In recent years, exciting new research has uncovered incidental visual statistical learning (VSL) and monetary reward as additional drivers of selective attention (Anderson, 2013; Chelazzi, Perlato, Santandrea, & Della Libera, 2013; Chun & Jiang, 1998; Hutchinson & Turk-Browne, 2012; Raymond & O'Brien, 2009; Turk-Browne, 2012). Features or locations frequently associated with a search target are prioritized, even when participants are unaware of the statistical regularities (Chun & Jiang, 1998, 1999; Hon & Tan, 2013). Similarly, features or locations associated with higher monetary reward are prioritized, often without explicit awareness (Anderson, Laurent, & Yantis, 2011; Chelazzi et al., 2014; Le Pelley, Pearson, Griffiths, & Beesley, 2014). These findings led to new theories of attention that postulate goals, salience, and previous experience (including statistical and reward learning) as major drivers of attention (Awh, Belopolsky, & Theeuwes, 2012; Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014). To date, research on different sources of attention has

used different approaches and experimental paradigms, making it difficult to compare their relative utility in directing attention. The present study documents the impact of statistical regularities, monetary reward, and top-down goals on spatial attention. Unlike previous research, we will test all three sources using the same visual search task, facilitating their comparison. Following this, the empirical data will be used to construct a quantitative account of monetary reward and statistical regularity.

Experimental approach

We asked participants to search for a letter “T” target among multiple “L” distractors presented in random locations on the screen (Fig. 1). The large number of distractors and the high similarity between the target and distractors make the T-among-L task among the most inefficient of visual search tasks (Wolfe, 1998). Bottom-up featural differences between the target and distractors cannot effectively guide attention to the target, minimizing the role of salience-driven attention. These properties make the T-among-L task an ideal choice for investigating top-down effects on spatial attention. Although the T-among-L task was frequently used in previous studies on goal-driven attention and visual statistical learning (Chun & Jiang, 1998), other tasks and paradigms (such as a two-color search task) were more commonly used in research on reward-based attention (Anderson, 2013). One reason for the difference in paradigms is that the majority of existing studies on reward-based attention have associated monetary reward with visual features (Anderson, 2013). Our work focuses on spatial attention, motivating us to select search tasks that minimize effects of bottom-up salience (similar rationale can be found in Chelazzi et al., 2014).

To investigate effects of statistical regularities on spatial attention, we used the location probability learning paradigm

to induce an attentional bias toward one quadrant of the display (Geng & Behrmann, 2002; Jiang, Swallow, Rosenbaum, & Herzog, 2013; Miller, 1988). Unbeknownst to the participants, across multiple trials the target appeared in one visual quadrant three-times more often than in any of the other visual quadrants. Monetary reward was not used. We examined how quickly participants started to prioritize the high-probability quadrant, the magnitude of the response time (RT) advantage for targets in the high-probability relative to the low-probability quadrants, and the persistence of this effect after the target’s location became completely random.

To investigate effects of monetary reward on spatial attention, we implemented analogous reward structure as that used in previous research on value-driven, featural attention (Anderson, 2013; Chelazzi et al., 2013), or reward-based spatial attention (Camara, Manohar, & Husain, 2013; Chelazzi et al., 2014; Hickey, Chelazzi, & Theeuwes, 2014). The target appeared equally often in the four visual quadrants. However, unbeknownst to participants, finding a target in one visual quadrant yielded three times more monetary reward than finding a target in any of the other visual quadrants (the magnitude of reward was comparable to that used in previous research, see Stankevich & Geng, 2014). Similar to previous research, reward was given probabilistically such that the high-reward quadrant was associated with a higher probability of gaining a high-reward than were the low-reward quadrants. The probabilistic allocation of reward reduced the likelihood that participants would explicitly associate reward with spatial locations (Stadler & Frensch, 1998). We examined the pace of learning, the magnitude of the RT advantage in the high-reward quadrant relative to the low-reward quadrants, and the persistence of this effect after all quadrants were rewarded equally.

Finally, to investigate goal-driven attention, we modified the reward experiment by adding explicit instructions about the reward structure. Participants were told that one quadrant

Fig. 1 Sample display and trial sequences used in the current study. Left: A sample T-among-L search display. Right: Schematic illustration of the two trials used in Experiment 2. Participants searched for the letter T to uncover either a treasure box (high reward) or old boots (no reward)

was associated with higher monetary reward and were encouraged to discover that quadrant and prioritize search in that quadrant. The use of explicit instructions about reward contingency distinguished this experiment from the incidental learning experiments, and from the majority of previous research on reward-based attention (Anderson et al., 2011; Chelazzi et al., 2014; Hickey et al., 2014; Le Pelley et al., 2014). We examined the RT advantage in the high-reward quadrant relative to the other quadrants. We separated participants who accurately identified the high-reward quadrant (and therefore were successful in setting up goal-driven attention) from those who failed to identify the high-reward quadrant. Similar to the other experiments, we examined the pace, magnitude, and persistence of attentional preference for the instructed quadrant. Setting up a task-goal to prioritize one visual quadrant can be done in many ways, including verbal instructions or central cues (Jiang, Swallow, & Rosenbaum, 2013; Won & Jiang, 2015). We chose to add explicit instructions to the reward experiment to facilitate comparisons with experiments on incidental reward learning. Table 1 summarizes the three types of attentional manipulations used in the present study.

Theoretical development

Besides documenting effects of statistical learning, monetary reward, and goals on spatial attention, our second goal was to examine the mechanisms underlying the three sources of attention. Among them, goal-driven attention is least controversial, because it has been well described by classic models of attention (Bisley & Goldberg, 2010; Desimone & Duncan, 1995; Egeth & Yantis, 1997; Fecteau & Munoz, 2006; Wolfe, 2007). In contrast, because VSL and reward learning are incidental, it is unclear how they influence attention (for controversies, see Asgeirsson & Kristjánsson, 2014; Kunar, Flusberg, Horowitz, & Wolfe, 2007).

One candidate mechanism that contributes to both VSL and reward-based attention is repetition priming. Walthew and Gilchrist (2006) noted that when a target appeared in some locations more often than in others, the high-probability target locations were more likely to repeat on consecutive trials. Because RT is faster when the target's location is the same as that of the preceding trial (Maljkovic & Nakayama, 1996), location probability cuing may originate

from increased repetition priming in the high-probability locations.

Research on reward-based attention has also implicated repetition priming as a potential mechanism for reward modulation. Della Libera and Chelazzi (2006) evaluated the size of priming across consecutive trials. Identification of the prime stimulus was followed by either a high reward (e.g., 10 cents) or a low reward (e.g., 1 cent). Repetition of the prime stimulus' features in the probe stimulus influenced probe identification RT, demonstrating priming. Crucially, the size of reward given to the prime stimulus modulated priming. High reward increased priming, whereas low reward reduced it. Similar modulatory effects of reward have been found on priming of pop-out (Kristjánsson, Sigurjónsdóttir, & Driver, 2010) and on spatial priming (Hickey et al., 2014). These studies converged on the idea that repetition priming may be an important mechanism underlying reward-based attention. However, because effects of priming on attention have generally being short-lived (Kristjánsson & Campana, 2010; Maljkovic & Nakayama, 1994, 1996), previous studies have not used priming to account for changes in attention to locations that receive high reward consistently. Our study will use mathematical simulations to quantify the impact of repetition priming on location probability cuing and reward-based attention. In turn, it helps us develop a multiple-systems view of attention that link VSL and reward with the different subsystems of top-down attention.

Overview

We present data from four experiments on VSL-based attention (Experiment 1), reward-based attention (Experiments 2–3), and goal-driven attention (Experiment 4). In all experiments, we focused on the pace of acquiring a spatial attentional bias toward the more “important” quadrant, the size of the RT advantage in the “important” quadrant relative to the other quadrants, and the persistence of this bias after the removal of statistical regularities or differential monetary reward. Based on our previous researching using the T-among-L search task, we calculated repetition priming over successive trials. We then simulated location probability learning and reward modulation of priming. The empirical data and simulation led us to propose the multiple-systems view of top-down spatial attention.

Table 1 Summary of experimental manipulations for VSL-based, reward-based, and goal-driven spatial attention

	Manipulation	Nature of the more “important” visual quadrant	Nature of learning
VSL	Target's location probability	Three times more likely than the other quadrants to contain the target	Incidental
Reward	Target's reward probability	Associated with three times more monetary reward than the other quadrants	Incidental
Goal-driven	Instructions	Experimenter instructions to prioritize one quadrant for more monetary gain	Intentional

Experiment 1

Experiment 1 examined spatial attention based on statistical learning of the target's location probability. In the training phase, the target was presented in one, high-probability quadrant on 50 % of the trials, and in each of the other low-probability quadrants on 16.7 % of the trials (Jiang, Swallow, Rosenbaum, et al., 2013). In the testing phase, the target appeared in all quadrants equally often (25 % probability). To increase the comparability with experiments on reward-based attention (Experiments 2–3), we used both a target localization task (Experiment 1A) and a target identification task (Experiment 1B). No monetary reward was given. We examined whether participants became faster finding the target in the high-probability quadrant relative to the low-probability quadrants and how the attentional bias persisted in the testing phase.

Method

Participants Participants in this study were students from the University of Minnesota. They were 18–45 years old and had normal or corrected-to-normal visual acuity. Each participant completed just one experiment. Participants received a standard base compensation for their time plus a cash bonus based on the reward used in the experiments.

The number of participants (16 per experiment) was predetermined to be comparable to our previous studies on probability cuing (Jiang, Swallow, Rosenbaum, et al., 2013). The sample size was doubled in Experiment 4 due to the presence of subgroups of participants.

Thirty-two participants completed Experiment 1, half of them in Experiment 1A (12 females and 4 males with a mean age of 20.2 years) and the other half in Experiment 1B (11 females and 5 males with a mean age of 19.0 years). The two versions of the experiment differed in how participants responded to the target.

Equipment Participants were tested individually in a room with normal interior lighting. Stimuli were generated using Psychtoolbox (Brainard, 1997; Pelli, 1997) implemented in MATLAB (www.mathworks.com) and displayed on a 17" CRT monitor (1024 × 768 pixels; 75 Hz). Viewing distance was unconstrained but was approximately 40 cm.

Stimuli Before each trial a white fixation point ($0.4^\circ \times 0.4^\circ$) was placed at a random location within the central 1.5° of the monitor. The mouse cursor was set to the center of the display. Each search display contained 12 items (one T and 11 Ls) placed in randomly selected locations in an invisible 10×10 matrix ($25^\circ \times 25^\circ$; Fig. 1). The items were white presented against a black background. Each visual quadrant contained 3 items. The target letter T was rotated 90° either clockwise or

counterclockwise, randomly determined on each trial. The distractor Ls had a random orientation of 0° , 90° , 180° , or 270° . Each search item subtended $1.6^\circ \times 1.6^\circ$.

Procedure Participants first completed one block of practice trials involving randomly placed target. They were then tested in 11 blocks of experimental trials. The first 7 blocks contained 36 trials per block and constituted the training phase. The last 4 blocks contained 32 trials per block and constituted the testing phase. Participants clicked on the fixation square to initiate a trial. The click required eye-hand coordination and ensured that the eyes returned to the approximate center of the display. After a brief delay of 200 ms the search display appeared and remained in view until participants made a response. Participants in Experiment 1A were asked to click on the location of the letter T using the mouse cursor. The target localization task was used to increase the salience of the target's location. Participants in Experiment 1B were asked to indicate the orientation of the target by pressing the left or right arrow keys. When an incorrect response was made, the computer voice spoke the sentence "That was wrong. Please try to be accurate." When the search RT was longer than a cutoff RT, the message "too slow" was displayed. Similar to later experiments on reward-based attention, the cutoff RT was the median RT of the preceding block of trials. For the first block, the cutoff RT was the median RT in the practice block. A demonstration of a few trials can be found at the following website: <http://jianglab.psych.umn.edu/reward/reward.html>.

Trials were self-paced. Participants were asked to respond as accurately and as quickly as they could. The computer reminded people to take a short break at the end of each block. We did not tell participants where the target was likely to be found, nor did we inform them of a potential change in the target's location probability.

Design In the training phase the target appeared in one quadrant, the high-probability quadrant, in 50 % of the trials. It appeared in any of the other three quadrants on 17 % of the trials. The high-probability quadrant was randomly selected for each participant and counterbalanced across participants. It remained the same throughout the training phase. In the testing phase, the target appeared in each quadrant 25 % of the trials.

Recognition test At the completion of the experiment, we first asked participants to tell us whether they thought the target was equally likely to appear anywhere or whether it was more often in some locations than others. Regardless of their response, participants were told that the target was not randomly placed. They were asked to provide a forced-choice response by clicking on the quadrant where they thought the target was most often found.

Results

Accuracy in this experiment and all subsequent experiments were higher than 97 % and was unaffected by any experimental conditions, all p s > 0.10. We therefore report only RT data, excluding incorrect trials and trials with RTs longer than 10 seconds (<0.5 % in all experiments). Figure 2 shows mean search RT across trial blocks, separately for trials in which the target appeared in the high-probability and the low-probability quadrants.

Search RT In the training phase (Fig. 2), participants were significantly faster when the target was in the high-probability quadrant than in the other quadrants, $F(1, 15) = 20.14$, $p < 0.001$, $\eta_p^2 = 0.57$ in Experiment 1A (localization task), $F(1, 15) = 31.38$, $p < 0.001$, $\eta_p^2 = 0.68$ in Experiment 1B (identification task). Probability cuing emerged early, resulting in a lack of significant interaction between target quadrant and block, $F(6, 90) = 1.91$, $p < 0.09$, $\eta_p^2 = 0.11$ in Experiment 1A, and $F < 1$ in Experiment 1B. The spatial bias toward the high-probability quadrant persisted in the testing phase (Blocks 8–11). RT was significantly faster when the target was in the high-probability quadrant, $F(1, 15) = 19.86$, $p < 0.001$, $\eta_p^2 = 0.57$ in Experiment 1A, $F(1, 15) = 18.10$, $p < 0.001$, $\eta_p^2 = 0.55$ in Experiment 1B. It was unaffected by experimental block or block by target quadrant interaction, F s < 1. A direct comparison between the two versions showed that RT was significantly slower in Experiment 1A (localization task) than in Experiment 1B (identification task), $p < 0.001$. However, experimental version did not interact with other factors, all p s > 0.10. Probability cuing did not become greater in the localization experiment, which had slower RT. This may be because the increase in RT was due to greater motor complexity, a factor unaffected by location probability learning. Similar to previous studies, participants demonstrated rapid acquisition and slow extinction of an

attentional bias toward the high-probability quadrant (Jiang, Swallow, Rosenbaum, et al., 2013).

Recognition Six of the 16 participants in Experiment 1A, and 6 of the 16 participants in Experiment 1B correctly identified the high-probability quadrant. The recognition rates were not significantly above chance, $\chi^2(1) = 2.67$, $p > 0.10$. Because recognition was administered after the testing blocks (with random target locations), it may have reduced participants' ability to recognize the high-probability quadrant. Nonetheless, the lack of significant awareness was inconsistent with the RT data, which showed a strong probability cuing effect even toward the end of the testing phase. The difference between the high-probability and low-probability quadrants remained significant at the end of testing (Block 4), $F(1, 30) = 6.72$, $p < 0.02$, $\eta_p^2 = 0.18$. To further examine whether recognition had any bearings on search performance, we separated participants who correctly identified the high-probability quadrant from those who failed. We then examined visual search RT as a function of recognition accuracy. This analysis showed that recognition accuracy did not interact with any experimental factors in either version of the experiment, all p -values > 0.15. All participants, including those who were unable to identify the high-probability quadrant, demonstrated probability cuing.

Discussion

Experiment 1 found that spatial attention was sensitive to the uneven location probability of a visual search target. This sensitivity was revealed under incidental learning conditions and was strong even in participants who failed to identify the high-probability quadrant. The effect generalized to identification and localization tasks. The rapid onset of learning implied short-term repetition priming as a potential mechanism for location probability learning. However, the long-term persistence in the testing phase indicated that long-term learning had occurred. To fully capture the relative impact of short-term priming and long-term learning, it would be necessary to simulate the learning data by modeling repetition priming. We will pursue this approach after presenting experiments on reward-based and goal-driven attention.

Experiment 2

Experiment 2 investigated whether incidental learning of the association between monetary reward and the target's locations produced an attentional bias toward the high-reward locations. To integrate reward as part of the search task, we introduced a foraging component to standard visual search tasks by asking participants to look for a treasure box hidden behind the letter T. Participants therefore searched for the

Fig. 2 Results from Experiment 1. The first 7 blocks (36 trials per block) involved unequal target location probability across quadrants. The last 4 blocks (32 trials per block) involved equal target location probability. Error bars show ± 1 standard error of the mean

letter T and clicked on its location, uncovering either a treasure box or old boots at the target's location (Fig. 1). Each treasure box yielded 4 points, but old boots yielded no points. Points accumulated across trials and were converted into monetary payment at the completion of the experiment. The reward values used here were the same as those used in a previously published study that had revealed successful reward learning (Stankevich & Geng, 2014).

To produce uneven reward across different locations, we manipulated the odds of finding treasure in different visual quadrants. Targets in the high-reward quadrant frequently led to treasure (75 % treasure, 25 % boots), whereas targets in the low-reward quadrants frequently led to old boots (25 % treasure, 75 % boots). Reward was delivered probabilistically rather than deterministically (e.g., 100 % treasure or 100 % boots) to reduce the likelihood that participants would explicitly associate reward with locations (Stadler & Frensch, 1998). After 8 blocks of training, we changed the reward structure to contain equal reward in all quadrants. Specifically, in the testing phase (4 blocks) locating the target yielded a treasure on 50 % of the trials, regardless of where the target was. Similar to Experiment 1, the testing phase was included to examine the persistence of the attentional bias toward the high-reward quadrant should such a bias emerge during training. The inclusion of the testing phase also is justified, because several previous studies showed that reward-based effects were more noticeable in the testing phase than in the training phase (Anderson et al., 2011; Chelazzi et al., 2014).

Experiment 2 preserved key features of most previous reward-based attention studies. First, although participants were explicitly asked to uncover reward by clicking on the target T, they were not informed of the contingency between monetary reward and the target's locations (Anderson et al., 2011; Chelazzi et al., 2014; Le Pelley et al., 2014). Second, reward had no functional significance for the search task (Anderson, 2013; Chelazzi et al., 2014). Because the target was equally likely to appear in the high-reward and low-reward locations, reward was not predictive of the target's location. In addition, reward was delivered regardless of search RT. Assuming that participants do not give up search until the target is found (a valid assumption when the target appears on every trial, see Chun & Jiang, 1998), searching the high-reward locations first would not yield greater monetary earnings than random search. Under these conditions, reward could not serve as an incentive for finding the target or for gaining more money. These conditions were stringent, yet they were standard practice in the existing literature and ensured that observed effects on attention cannot be attributed to the strategic use of reward. If participants prioritized the high-reward quadrant under these conditions, it would provide compelling evidence for the automatic guidance of spatial attention by monetary reward.

Method

Participants Sixteen new participants completed Experiment 2. There were 15 females and 1 male with a mean age of 20.2 years.

Stimuli Stimuli used in Experiment 2 were similar to those of Experiment 1 except for the following additions. Before each trial reward point total (e.g., “Point total: 200”, 5.7° below fixation) and cash bonus total (e.g., “Cash bonus total: \$1.00”, 8.9° below fixation) were displayed (font size: 40), along with a fixation point. After the target was found, the T turned into either a picture of a treasure chest (1.6° × 1.6°) or old boots (1.6° × 1.6°). To increase the salience of the reward, a large picture (31° × 31°) of four stacked gold coins indicated reward on trials when participants gained reward, followed by the text “+4” shown in large font (font size 150) 7.8° above fixation (Fig. 1).

Procedure Participants first completed one block of practice trials involving randomly placed reward. They then completed 384 trials of the main experiment. Participants clicked on the fixation square to initiate a trial. After a brief delay of 200 ms, the search display appeared and remained in view until participants used the mouse cursor to click on the location of the letter T. We used a target localization task to increase the salience of the target's location. Similar to Experiment 1, participants heard the computer voice speaking the sentence “That was wrong. Please try to be accurate.” for incorrect trials. If participants correctly clicked on the T, the T would immediately turn into either a gold treasure box or black boots while the Ls would remain in their original locations and a chirp would follow. This display lasted 500 ms. Uncovering the treasure box would lead to a large picture of stacked gold coins for 1 s, whereas uncovering old boots would lead to a blank display for 1 s. The fixation display before the next trial then showed cumulative points total and cash bonus total. A large text “+4” was at the top of the fixation display on trials when a reward was gained (Fig. 1). A demonstration of a few trials can be found at the following website: <http://jianglab.psych.umn.edu/reward/reward.html>.

We did not tell participants where the treasure might be, nor did we inform them of a potential change in the reward locations.

Design The first 8 blocks comprised the training phase and the last 4 blocks comprised the testing phase (32 trials per block¹).

¹ The training phase had a total of 256 trials, comparable to the 252 trials used in Experiment 1. The mismatch was incurred because trial number must be divisible by 6 in Experiment 1's training phase, and divisible by 16 in Experiment 2's training phase.

The target T was randomly placed, with a 25 % probability of appearing in each of the four visual quadrants. However, the reward hidden behind the T varied across the visual quadrant. In the training phase, one quadrant was designated as the *high-reward quadrant*. In this quadrant, the treasure box appeared on 75 % of the trials, whereas the old boots appeared on 25 % of the trials. The other three quadrants were the *low-reward quadrants*, which contained the treasure box 25 % of the time and the old boots 75 % of the time. Each time when a treasure box was discovered participants were awarded 4 points, which were converted to a cash bonus of 2 cents. Each time old boots were discovered participants received 0 points. Which quadrant was the high-reward quadrant was randomly selected and counterbalanced across participants. However, it was fixed for a given participant throughout the training phase. In the testing phase, all quadrants had equal reward probability: 50 % probability of treasure and 50 % probability of old boots.

Recognition test At the completion of the experiment, we asked participants whether they thought the treasure chest was randomly located on the screen. Regardless of their response, we told them that the treasure more often was placed in one quadrant than in the others and asked to indicate the high-reward quadrant using a mouse click.

Results

Search RT We conducted an ANOVA using target quadrant and block as within-subject factors on visual search RT (Fig. 3). In the training phase (Blocks 1–8), monetary reward was three times higher in the high-reward quadrant than in any of the low-reward quadrants. However, participants were not significantly faster finding the target in the high-reward quadrant than in the other quadrants. The main effect of target quadrant was not significant, $F < 1$, neither did target quadrant interact with block, $F < 1$. RT became faster as the experiment progressed, resulting in a significant main effect of block, $F(7, 105) = 2.98$, $p < 0.007$, $\eta_p^2 = 0.16$.

Fig. 3 Results from Experiment 2. Error bars show ± 1 standard error of the mean. The gap in the graph indicated a transition in reward structure after block 8. Monetary reward was higher in the high-reward quadrant than the low-reward quadrants in Blocks 1–8 but was equal across all quadrants in Blocks 9–12

In the testing phase (Blocks 9–12), targets were rewarded equally often in all quadrants. RT was unaffected by target quadrant, block, or their interactions, all F s < 1 .

Recognition of the reward structure Six of the 16 participants correctly identified the high-reward quadrant, which was not statistically higher than chance, $\chi^2(1) = 1.33$, $p > 0.20$. To examine how reward learning might have differed among participants who may have different degrees of awareness, we conducted a further analysis on training phase search RT that grouped participants based on the forced-choice accuracy. Recognition accuracy did not interact with any experimental factors, all $p > 0.10$. No effects of reward were found when we restricted the analysis to the six participants who made the correct forced-choice, $p > 0.10$.

Discussion

Even though monetary reward was higher for targets in one visual quadrant than the other quadrants, participants did not develop a strong attentional bias toward the high-reward quadrant. These results showed that monetary reward used in our study was not a strong driver of spatial attention. However, as indicated earlier, Experiment 2 adopted stringent conditions for testing reward-based attention. Not only were the high-reward locations no more likely to contain the search target but also were there no monetary advantage for prioritizing the high-reward quadrant. In addition, participants were not informed of the contingency between reward and locations. The lack of incentives for detecting targets or for monetary earning and the lack of explicit instructions about reward contingency were used in the majority of previous studies in support of the automatic nature of reward-based attention (Anderson, 2013; Failing & Theeuwes, 2014). However, these previous studies tested primarily feature-based attention rather than spatial attention (but see Experiment 3 of Anderson et al., 2011; Chelazzi et al., 2014). Because we did not induce strong reward-based prioritization of spatial attention, it is important to test reward-based attention under less stringent conditions.

Experiment 3

Experiment 3 tested participants under conditions in which prioritizing the high-reward locations can yield more money. We made one change to the reward structure of Experiment 2. Participants gained money only when their RT was faster than the median RT of the preceding block. This change resulted in higher monetary earnings when the high-reward quadrant was prioritized (scenario #1) than when participants did not prioritize the high-reward quadrant (scenario #2). The use of median RT as the reward cutoff means that approximately 50 % of the trials would be rewarded. For scenario #1, all trials in

which the target was in the high-reward quadrant (this was 25 % of all trials) would be fast enough to meet the cutoff. Thus, for trials when participants were fast enough to gain reward, half of these would come from the high-reward quadrant (expected reward: 3 points), whereas the other half would come from the low-reward quadrants (expected reward: 1 point). The mean expected reward would be 2 points. For scenario #2 in which participants searched randomly, trials that were fast enough to gain a reward would be equally likely to come from any quadrant. That is, a quarter of these would come from the high-reward quadrant and three-quarter would come from the low-reward quadrants. The expected reward under scenario #2 would be 1.5 points. Thus, prioritizing the high-reward quadrant would yield more reward than search at random.²

Because monetary earnings increased if participants prioritized the high-reward quadrant, Experiment 3 contained contingent reinforcement that was absent in Experiment 2. Even though the target's location was still random and unpredicted by reward, prioritizing the high-reward quadrant would yield greater earnings. To extend the generality of the findings we tested participants in both a localization task and an identification task.

Method

Participants Sixteen participants each completed Experiments 3A and 3B. Experiment 3A included 7 females and 9 males with a mean age of 20.3 years. Experiment 3B included 13 females and 3 males with a mean age of 21.1 years. The two versions of the experiment differed only in the response participants made to the target.

Procedure The experiment was the same as Experiment 2 except for the following changes. First, the point value doubled (each point corresponded to 1 cent rather than ½ cents). In addition, reward was given only when participants had successfully uncovered the treasure box within a prespecified cutoff time. Similar to Experiment 1, the cutoff time was the median RT in the preceding block of trials.

At the beginning of each block, participants were informed that reward would be gained only if they discovered a treasure box within a specific cutoff time (e.g., 2,500 ms). Participants in Experiment 3A used the mouse cursor to click on the T. Participants in Experiment 3B pressed the arrow keys on the keyboard to identify the T's orientation (left or right). Upon their response, participants heard either the chirps (for correct localization) or the computer voice saying “That was wrong...” (for incorrect localization).

² A formal proof of the difference in monetary earning under more complex cutoff conditions can be found at: <http://jianglab.psych.umn.edu/reward/reward.html>.

Like Experiment 2, for correct trials the T immediately turned into either a treasure box or old boots. For trials in which RT was faster than the cutoff, the trial sequence was identical to that of Experiment 2. For trials in which RT was slower than the cutoff, instead of seeing a picture of the stacked gold coins or a blank display, the text “too slow” (font size 20) was displayed at the center of fixation. At the end of each block the computer displayed the new cutoff time for the next block of trials. Sample trials can be viewed at: <http://jianglab.psych.umn.edu/reward/reward.html>. Similar to Experiment 2, we did not provide explicit instructions about the reward structure used in the study, so reward-based learning must occur incidentally.

Results

Search RT Results (Fig. 4) were similar to those of Experiment 2. We performed ANOVA using task (localization or identification) as a between-subject factor and target quadrant and block as within-subject factors. In the training phase (Blocks 1–8), RT became faster in later blocks than earlier blocks, $F(7, 210) = 10.98$, $p < 0.001$, $\eta_p^2 = 0.27$. However, there was no main effect of target quadrant, $F(1, 30) = 2.13$, $p > 0.15$, or a quadrant by block interaction, $F(7, 210) = 1.26$, $p > 0.20$. Similarly, the testing phase (Blocks 9–12) showed no effects of target quadrant, block, or their interaction, all $p > 0.50$. Participants who made the mouse click response were significantly slower than those who made keyboard responses, $F(1, 30) = 68.97$, $p < 0.001$, $\eta_p^2 = 0.70$. However, response type did not interact with any of the experimental factors, all $ps > 0.15$.

Relationship between attentional prioritization and earnings

An important assumption that we made in Experiment 3 was that participants could earn more money when they prioritized the rich quadrant than when they did not. To validate this assumption, we examined individual differences in monetary gain (Fig. 4B). Some participants happened to be faster searching in the high-reward quadrant than the low-reward quadrants, whereas others did not. If attentional prioritization of the rich quadrant was associated with greater earnings, then those who showed a larger RT advantage in the high-reward quadrant relative to the low-reward quadrants should earn more money in the training phase. This was indeed the case. The correlation between the RT advantage in the high-reward quadrant and monetary earnings in the training phase was significant, Pearson's $r = 0.55$, $p < 0.03$ in Experiment 3A, Pearson's $r = 0.84$, $p < 0.001$ in Experiment 3B. These data validated the assumption that an attentional bias toward the rich quadrant could yield higher earnings. Despite of this feature, monetary reward did not substantially influence spatial attention in Experiment 3.

Fig. 4 A. Search RT results from Experiment 3. Error bars show ± 1 standard error of the mean. B. Individual differences in earnings in the training phase correlated with the RT advantage in the rich quadrant. Each dot represented a single participant

Recognition of the reward structure Seven of the 16 participants in Experiment 3A were able to identify the high-reward quadrant in the recognition test, which was marginally higher than expected by chance, $\chi^2(1) = 3.00, p = 0.08$. Six of the 16 participants in Experiment 3B successfully identified the high-reward quadrant, $\chi^2(1) = 1.33, p > 0.20$. To examine whether explicit recognition correlated with search performance, we split participants based on whether they made a correct forced-choice recognition response. In the search data, recognition accuracy did not interact with any of the experimental factors, all $p > 0.10$.

Discussion

We extended the findings from Experiment 2 to show that under incidental learning conditions, locations associated with higher monetary reward were not searched significantly faster than were other locations. This was found even though prioritizing the high-reward quadrant would have yielded greater monetary earnings. The earning difference between the prioritization strategy and random search was moderate, amounting to approximately 33 % greater monetary gain. This may be one reason why effects of reward were insignificant. However, the moderate difference in earnings did not explain why significant effects of reward were observed in previous research that rewarded people equally regardless of search strategy (Anderson, 2013; Chelazzi et al., 2013; Della Libera & Chelazzi, 2006; Hickey et al., 2014). An important challenge for any theoretical account is to provide a rational explanation for the negligible effects of reward. In the *General Discussion*, we will use mathematical simulation of repetition priming to account for the findings observed in Experiments 2–3.

Experiment 4

Like the majority of the existing research on reward-based attention (Anderson et al., 2011; Chelazzi et al., 2014; Le Pelley et al., 2014), our investigation so far focused on

incidental learning of a spatial attentional bias. We found that whereas the target's location probability exerted a strong influence on spatial attention, its reward probability did not. Experiment 4 aimed to document effects of goal-driven attention in the same search task as that used in Experiments 1–3. Our main purpose was to show that reward can be a powerful mediator of spatial attention if it is incorporated into the task goal. To this end, in Experiment 4 we told participants that reward would be more often placed in one quadrant than in any of the other quadrants. However, we did not tell participants where the high-reward quadrant would be but encouraged them to discover it in the experiment. To assess explicit awareness, we asked them to identify the high-reward quadrant at the beginning of each block. We then split participants into those who successfully identified the reward structure from those who failed and separately examined the development of an attentional bias toward the high-reward quadrant.

By providing explicit instructions about the reward-contingency, Experiment 4 incorporated reward into the task goals. This feature distinguished Experiment 4 from the majority of previous research on reward-based attention, which typically required participants to learn the reward contingency through incidental learning (Anderson et al., 2011; Chelazzi et al., 2014; Le Pelley et al., 2014; but see Lee & Shomstein, 2013; Stankevich & Geng, 2014 for the use of explicit instructions).

Method

Thirty-two new participants completed Experiment 4 (21 females and 11 males; mean age 20.8 years).

This experiment was identical to Experiment 3 except for an additional instruction given after practice but before the first testing block. Participants were told that: “In this experiment, the treasure box will be placed in one visual quadrant three times more often than in any of the other visual quadrants. Although we will not tell you where it is, we would like you to try to identify the rich quadrant. The location of this quadrant will be constant throughout the experiment.”

Following this instruction participants were asked to report where the rich quadrant was. They were encouraged to prioritize the high-reward quadrant to maximize earnings. Because participants had not been exposed to the uneven reward structure, recognition before the first block was based on random guesses. The experiment then proceeded the same way as in Experiment 3, except that participants were asked to click on the high-reward quadrant after each experimental block.

Results

Recognition of the reward structure We separated participants into three groups based on their recognition performance at the completion of Blocks 1–8 (Fig. 5). This calculation did not include recognition data before Block 1, because participants had not been exposed to the reward structure. Nor did it include recognition that happened after the change of the reward structure (note that the reward was randomly distributed in Blocks 9–12). The “low-recognition” participants were those whose mean recognition rate was at or below chance (25 %). The “high-recognition” participants were those whose mean recognition rate was 2 standard deviations (SD) above chance (56 % or above, calculated by assuming a binomial distribution for recognition performance across the 8 blocks). The “intermediate-recognition” participants had recognition rates between 25 % and 56 %.

Search RT We first analyzed search RT data (Fig. 6) in the training phase using an ANOVA that included recognition group as a between-subject factor, and target’s quadrant and block (1–8) as within-subject factors. This analysis showed a significant main effect of the target’s quadrant, $F(1, 31) = 26.24$, $p < 0.001$, $\eta_p^2 = 0.46$. RT was faster when the target was in the high-reward quadrant than in the low-reward quadrants. However, this effect interacted significantly with recognition group, $F(2, 31) = 17.56$, $p < 0.001$, $\eta_p^2 = 0.53$. Follow-up tests showed that only participants in the high-recognition group showed faster search in the high-reward quadrant, $F(1, 16) = 93.80$, $p < 0.001$, $\eta_p^2 = 0.85$, whereas the other two groups did not, $p > 0.10$. A significant main effect of block

showed improved RT over time, $F(7, 203) = 8.69$, $p < 0.001$, $\eta_p^2 = 0.23$. Although block did not interact with recognition group, $F(14, 203) = 1.24$, $p > 0.10$, the three-way interaction between recognition group, block, and target’s quadrant was significant, $F(14, 217) = 1.81$, $p < 0.05$, $\eta_p^2 = 0.11$. For the high-recognition group, target’s quadrant interacted with block, $F(7, 112) = 4.63$, $p < 0.001$, $\eta_p^2 = 0.22$, demonstrating an increase of spatial attentional bias over time. For the other groups, the interaction between target’s quadrant and block was not significant, $F_s < 1$. Finally, the main effect of group was significant, $F(2, 31) = 5.09$, $p < 0.05$, $\eta_p^2 = 0.25$, driven by faster RT in the high-recognition group than the other groups.

The reward structure changed in the testing phase: reward was randomly distributed in all quadrants. As shown in the recognition data (Fig. 5), the high-recognition group waived in their conviction that the high-reward quadrant was what they had selected before: the reduction from 100 % recognition rates after Block 8 to 76 % recognition rates after Block 9 was significant, $t(16) = 2.22$, $p < 0.05$. This should correspond to a reduction in goal-driven attention. However, perhaps because the instructions had mentioned that the high-reward quadrant would remain the same, most people continued to select the previously high-reward quadrant as where the treasure was most often found. The bias in the recognition choice was also reflected in the search RT data (Fig. 6). ANOVA on recognition group, target’s quadrant, and block (9–12) showed a significant main effect of target’s quadrant, $F(1, 29) = 7.49$, $p < 0.01$, $\eta_p^2 = 0.21$, which interacted with recognition group, $F(2, 29) = 8.61$, $p < 0.001$, $\eta_p^2 = 0.37$. None of the other main effects or interaction effects were significant, all $p > 0.10$. Follow-up tests restricted to each group confirmed that only the high-recognition group persisted in prioritizing the high-reward quadrant, $F(1, 16) = 32.87$, $p < 0.001$, $\eta_p^2 = 0.67$. Among the high-recognition participants, the RT difference between the high-reward and low-reward quadrants declined significantly from the last 4 blocks of training to the testing phase, $F(1, 16) = 18.23$, $p < 0.001$, $\eta_p^2 = 0.53$. Thus, a reduction in goal-driven attention had occurred when participants became less certain about which quadrant was associated with high reward.

Comparisons across experiments

Explicit instructions were important for the attentional bias toward the high-reward quadrant. This was demonstrated in an ANOVA that directly compared the training phase of Experiment 4 (the high-recognition group only) with either Experiment 2 or Experiment 3A. This analysis revealed a significant interaction between experiment and quadrant condition, $F(1, 31) = 42.52$, $p < 0.001$, $\eta_p^2 = 0.58$ when comparing Experiments 2 and 4, $F(1, 31) = 65.58$, $p < 0.001$, $\eta_p^2 = 0.68$ when comparing Experiments 3A and 4. Thus,

Fig. 5 Explicit recognition rates from Experiment 4. Error bars show ± 1 standard error of the mean

Fig. 6 Visual search RT from Experiment 4, separately for participants who had high, intermediate, or low recognition rates. Error bars show ± 1 standard error of the mean

intentional use of reward produced significantly greater attentional biases relative to incidental learning of reward.

Further evidence that reward itself did not increase the attentional effect toward the instructed quadrant came from a comparison between Experiment 4 and a previous study. In one experiment of Won and Jiang (2015) participants were instructed to search one visual quadrant before the other quadrants, but the experiment did not use monetary reward. A direct comparison between Experiment 4's high-recognition group and the comparable condition from the previous study (Won & Jiang, 2015, Experiment 5, set size 12 under no working memory load) revealed no interaction between target quadrant (instructed quadrant vs. other quadrants) and experiment, $F < 1$. The RT advantage in the instructed quadrant relative to the other quadrants was approximately 750 ms in the previous study without reward, which was similar to the 762-ms RT advantage in the training phase of the current study. This finding supports our characterization of Experiment 4 as an example of goal-driven attention.

Discussion

Explicit knowledge of the reward-location contingency resulted in the adoption of a task goal that prioritized the processing of locations in the high-reward quadrant. Thus, much like task instructions and other top-down factors, reward can be a powerful source of goal-driven attention. On the other hand, participants who could not identify the high-reward quadrant at above-chance levels failed to set up the proper task goal. Their failure to adjust spatial attention showed that for reward to strongly influence goal-driven attention, participants would need to have specific knowledge about which locations were more often rewarded.

Experiment 4 used reward to set up attentional priority in one quadrant. As noted in the introduction, this was but one method to manipulate goal-driven attention. Nonetheless, the strength of the manipulation appeared comparable to previous studies that used other approaches. Using the same T-among-L search task, one previous study presented participants with a central arrow whose direction changed from trial to trial. The quadrant indicated by the arrow was more likely to contain the

search target for the trial. The RT advantage in the validly cued quadrant, relative to the invalidly cued quadrant, was approximately 600 ms (Jiang, Swallow, & Rosenbaum, 2013). Another study provided verbal instructions for participants to prioritize one visual quadrant throughout an experimental block, a situation that's analogous to Experiment 4. The RT advantage in the instructed quadrant, relative to the other quadrants, was approximately 750 ms (Experiment 5, Won & Jiang, 2015). These findings showed that reward could be used to motivate a task goal.

Even when explicitly informed that reward was associated with the target's location, not all participants in Experiment 4 were able to learn which quadrant was the high-reward quadrant. This difficulty may be attributed to the probabilistic nature of the reward association, which is known to reduce explicit learning (Stadler & Frensch, 1998). The difficulty of acquiring explicit learning, however, does not necessarily preclude implicit learning of the reward association. Studies using other implicit learning paradigms (e.g., contextual cuing or artificial grammar learning) have shown that participants frequently acquire implicit learning of complex probabilistic structures even though they fail to acquire explicit learning (Chun & Jiang, 2003; Stadler & Frensch, 1998). The negligible learning shown in Experiments 2 and 3, therefore, cannot be attributed just to the complexity of the underlying statistical association.

General discussion

Summary of results

This study aimed to document the size of attentional effects based on visual statistical learning (particularly location probability cuing), monetary reward, and explicit goals. Using the same T-among-L search task, we found statistically significant effects of location probability learning and goal-driven attention, but not differential monetary reward. To synthesize the results, we calculated the mean RT difference between the "important" quadrant (e.g., the high-probability or high-reward quadrant) and the other quadrants in each experiment tested in the current study. Table 2 shows the RT difference,

Table 2 Summary of the experimental results for the last four training blocks and the 4 testing blocks. Mean RT difference between the more “important” quadrant and the other quadrants was calculated (inmilliseconds and as a percentage saving relative to the mean RT), along with effect size (Cohen’s *d*) and the 95 % confidence interval (CI)

		Last 4 training blocks (learning)			Testing blocks (persistence)		
Exp.	Manipulation	Mean (% saving)	Cohen’s <i>d</i>	95 % CI (ms)	Mean (% saving)	Cohen’s <i>d</i>	95 % CI (ms)
1A	VSL; mouse response	408 ms (18 %)	3.32	[272,543]	349 ms (14 %)	2.30	[182,516]
1B	VSL; button press	472 ms (27 %)	2.80	[289,654]	331 ms (19 %)	2.20	[165,498]
2	Reward; no deadline	9 ms (0 %)	0.05	[−192,210]	−36 ms (−0.1 %)	0.22	[−216,144]
3A	Reward; mouse response	−42 ms (−0.8 %)	0.24	[−234,150]	−41 ms (−0.6 %)	0.23	[−241,158]
3B	Reward; button press	7 ms (1.5 %)	0.06	[−120,134]	−108 ms (−5 %)	0.55	[−309,104]
4: high recogn.	Goals; successful	963 ms (41 %)	6.12	[796, 1130]	635 ms (27 %)	2.87	[400,870]
4: others	Goals; unsuccessful	30 ms (1 %)	0.13	[−242,301]	13 ms (1 %)	0.07	[−194,220]

the corresponding effect size (Cohen’s *d*), and the 95 % confidence interval for the last four blocks of the training phase and the testing phase. Because different experiments and participants had different mean RT, we also normalized the RT data, expressing the RT difference between the “important” quadrant and other quadrants as a percentage saving relative to an individual’s mean RT (e.g., [low-probability quadrant – high-probability quadrant]/mean RT).

Figure 7 plots data from the training phase, displaying the mean RT facilitation by the different attention manipulations used in the current study. Although not displayed here, the same pattern of results was seen when RT was normalized relative to an individual’s mean RT. The largest effect came from successful goal-driven attention, followed by location probability cuing. Monetary reward of the type used in the

current study produced the smallest effect. The 95 % confidence intervals of the three types of manipulations (successful goals, VSL, and reward) did not overlap, demonstrating substantial differences in the efficacy of these manipulations on spatial attention. The 95 % confidence intervals did not include 0 for successful goal-driven attention and for location probability learning but included 0 for unsuccessful goals and incidental reward. If one is to ensure that one region of space is prioritized relative to other regions, the most effective method is to explicitly cue spatial attention toward that region. In the absence of explicit cues, incidental learning of the target’s location probability can induce a large, persistent attentional bias toward the high-probability locations. In contrast, incidental monetary reward has a negligible effect.

Fig. 7 Mean RT difference between the more “important” quadrant and the other quadrants. The height of each line showed 95 % confidence interval. Data were obtained from the last 4 blocks in the training phase. From left to right were data from Experiment 4’s high-

recognition group, Experiment 4’s intermediate/low recognition groups, Experiment 1A, Experiment 1B, Experiment 2, Experiment 3A, and Experiment 3B

Conceptual issues

A long list of reasons may be cited to explain why, under incidental learning conditions, reward-based attentional effects were negligible in our study. Chief among them are (1) the association of money with a visual quadrant rather than with a specific location, (2) the relatively small amount of money, and (3) the probing of reward-based attention in the training task rather than in a transfer task. Although the contribution of these factors to reward-based attention is ultimately an empirical issue, we find them unsatisfactory as a full explanation. First, quadrant manipulation did not prevent participants from acquiring strong location probability cuing. Second, the small amount of money did not prevent reward-based effects from emerging in previous studies (Stankevich & Geng, 2014), and it was sufficient to drive goal-driven attention (Experiment 4). Finally, although reward-based attention was frequently assessed in transfer tasks (Anderson et al., 2011; Chelazzi et al., 2014), significant effects have been observed in the training task (Failing & Theeuwes, 2014; Jiang, Swallow, Won, Cistera, & Rosenbaum, 2014; Lee & Shomstein, 2014; Le Pelley et al., 2014). Thus, although it is important to explore effects of reward size and spatial precision in future research, we believe that alternative approaches are needed for a deeper understanding of our findings.

The approach we take next is to relate our findings to theories of location probability learning and reward-based attention. In the introduction we focused on short-term location repetition priming as a potential common mechanism for probability cuing and reward-based attention. Priming can contribute to probability cuing because quadrant repetition on consecutive trials is more frequent in the high-probability quadrant than the low-probability quadrants. In addition, previous studies showed that reward modulates location repetition priming (Hickey et al., 2014). However, that study only examined situations in which the rewarded locations changed randomly. Our study involved a long-term manipulation in which the same locations were more highly rewarded throughout training. The question is whether location repetition priming can explain the resultant attentional bias (or lack thereof) toward the high-reward or high-probability locations.

Repetition priming in location probability learning

We took two approaches to examine the contribution of repetition priming to location probability cuing. First, we examined whether participants in Experiment 1 were sensitive to the repetition of the target's visual quadrant. If so, their RT should be faster if the current trial's target was in the same, rather than a different, quadrant as that of the preceding trial. This was indeed the case. Participants in Experiment 1 were faster on repeat trials than on nonrepeat trials, both when the

target was in the high-probability quadrant (a difference of 120 ms in the training phase), and when it was in the low-probability quadrants (a difference of 91 ms in the training phase), yielding a significant main effect of quadrant repetition, $F(1, 31) = 9.80$, $p < 0.004$, $\eta_p^2 = 0.24$ but no interaction between the target's quadrant and repetition, $F < 1$. This analysis provided evidence that location repetition priming influenced search RT in our task.

Because repetition occurred more often in the high-probability quadrant, the attentional bias toward the high-probability quadrant may reflect, in part, greater accumulation of repetition priming in that quadrant. However, the analysis performed above is inadequate in quantifying the role of repetition priming. This is because priming estimation is confounded with the manipulation of location probability (i.e., repetition occurred more frequently in the more probable quadrant). It is therefore important to take a second approach that measures repetition priming in an independent dataset, and using that measurement to simulate RT advantage based on repetition priming.

To this end, we measured quadrant repetition priming using an independent dataset that did not include a manipulation of location probability or reward. The baseline dataset came from a previously published study that used the same T-among-L button-press task, but the target's location was random and there was no reward (Experiment 2 in Jiang, Swallow, Rosenbaum, et al., 2013, the “even” phase before an introduction of location probability manipulation). Figure 8 displays the baseline repetition priming data. Consistent with previous research (Maljkovic & Nakayama, 1996), finding a target in one quadrant facilitated search in the same quadrant on subsequent trials. The effect was largest for the immediately next trial (trial N+1) but also was noticeable for trial N+2. Priming appeared to dissipate thereafter, fluctuating around zero in longer lags (N+3 through N+8). Following previous research on repetition priming in visual search, we can model the repetition priming data in Fig. 8 as a power function with an asymptote of zero (Gupta & Cohen, 2002; Logan, 1990): $\text{Priming} = 358.5 \times j^{-3.805}$, where j was the trial lag between the current trial and the preceding trial. The R-square for the goodness of fit was 0.71. The following analysis used the power function to estimate repetition priming. The same pattern of results was observed when simulation was based on Fig. 8's raw priming measures.

How much could repetition priming contribute to location probability cuing? To address this question, we took a simulation approach. Specifically, we obtained the trial sequences that participants in Experiment 1B experienced and used the priming function from Fig. 8 to estimate the contribution of repetition priming for each trial. For example, suppose that A, B, C, and D represented the four visual quadrants, and that the target's quadrant in a sequence of 12 trials was *BCAADABAACDA*. The target was in quadrant B on the first

Fig. 8 Left: RT difference between quadrant repetition and quadrant nonrepetition as a function of the trial lag between the current trial and the preceding 1–8 trials. The observed data came from Jiang, Swallow, Rosenbaum, et al. (2013). Middle: Simulated repetition priming effects in the low-probability and high-probability quadrants in Experiment 1B.

trial, and by being the first trial, it was associated with no repetition priming. The target was in quadrant C on the second trial and therefore did not benefit from repetition priming. The target's quadrant on the third trial (quadrant A) differed from all preceding trials, yielding no repetition priming. On trial 4, however, the target was in quadrant A again. It received priming from the preceding trial but not from the earlier trials. Trial 5 (quadrant D) received no repetition priming, but trial 6 (quadrant A) received priming from two trials back and three trials back. Using this procedure, we can calculate repetition priming for each trial based on the history of the target's location. We can then average across trials in which the target was in the high-probability quadrant and ask whether these trials were associated with greater repetition priming than the other trials. Note that we did not model the search RT data from Experiment 1B; all that was done was to take trial sequences from this experiment to estimate repetition priming for each trial.

As shown in Fig. 8 (middle), estimated quadrant priming was substantially stronger on trials when the target was in the high-probability quadrant rather than the low-probability quadrants in the training phase, $F(1, 15) = 4515.76$, $p < 0.001$, $\eta_p^2 = 0.997$. The difference between the these quadrants rapidly disappeared in the testing phase, $F < 1$. To examine the degree to which repetition priming could account for location probability cuing, we compared the simulated priming difference between the high- and low-probability quadrants, with the observed RT difference between the high- and low-probability quadrants (Fig. 8, right). An ANOVA on the type of data (simulated or observed), target quadrant, and block revealed a significant main effect of the type of data, $F(1, 15) = 22.85$, $p < 0.001$, $\eta_p^2 = 0.60$. Repetition priming significantly underestimated the RT advantage in the high-probability quadrant in the training phase. It also failed to capture the persistence of the attentional bias in the testing phase. These findings were confirmed in a second simulation attempt that relied on another independent dataset to estimate repetition priming. Thus,

Right: Simulated priming difference between the high- and low-probability quadrants was smaller than the observed RT difference between these conditions. Error bars show ± 1 standard error of the mean. Some error bars may be too small to see

although repetition priming contributed to the RT difference between the high and low probability quadrants in the training phase, it was unable to account for the much larger and persistent effect of location probability learning. A full account for location probability cuing must include an additional component of long-term learning.

In our previous work on location probability learning, we theorized that the long-term component relates to changes in “procedural attention” or *how* people shift spatial attention during the search process (Jiang, Swallow, & Capistrano, 2013). When the target is frequently found in one quadrant, participants develop a strong habit to shift location in the direction of the high-probability quadrant. For simplicity of description, we will refer to this component as “search habit.” The simulation data provided strong evidence that location repetition priming contributed to location probability learning, but it was not the only mechanism behind it. Changes in search habit are likely the second component of location probability learning.

Repetition priming and reward

Our study indicated that reward could effectively guide goal-driven attention if it was successfully integrated with the task goal (Experiment 4). Under conditions of incidental learning, however, reward did not produce statistically significant effects (Experiments 2–3). Like the majority of reward-based studies in the existing literature, reward used in Experiments 2–3 can be characterized as being “retrospective.” Reward was delivered after task completion and carried no predictive value for the future. This is unlike location probability learning, in which the target's location probability is predictive of the target's future location. It is therefore puzzling how reward could influence attention in previous research. Priming is a strong logical explanation for this puzzle because priming is a memory process that influences subsequent performance (Gabrieli, 1998; Schacter & Buckner, 1998). If reward works by modulating priming then it can have prospective impact on

subsequent allocation of attention. In fact, previous studies that directly measured priming have shown a modulatory role of reward on priming (Della Libera & Chelazzi, 2006; Hickey et al., 2014; Kristjánsson et al., 2010). Our study found some support for the idea that reward modulated repetition priming. Across all participants tested in Experiments 2 and 3, after receiving a high reward, RT on the next trial was 2,660 ms if the target's quadrant switched and 2,524 ms if the target's quadrant repeated, a 136-ms repetition effect. After receiving a low reward, RT on the next trial was 2,634 ms if the target quadrant switched and 2,551 ms if the target quadrant repeated, a 83-ms repetition effect. The numerical trend was consistent with the idea that repetition priming was greater following a high reward than following a low reward, although this difference failed to reach statistical significance, $F < 1$. To adequately assess effects of reward on repetition priming, however, it is necessary to estimate priming from an independent dataset (Fig. 8) and to simulate priming from the preceding 8 trials (rather than the preceding 1 trial).

Can a repetition priming account of reward-based attention help understand results from Experiments 2–3? Although quadrant repetition did not differ systematically between the high and low reward quadrants, high reward may increase priming and low reward may decrease priming (Della Libera & Chelazzi, 2006; Kristjánsson et al., 2010). For example, in Hickey et al. (2014), location repetition priming was 26 ms following a low-reward and 50 ms following a high-reward, a twofold increase. If we assume that the twofold difference in priming also applies to our study, then we can estimate the size of repetition priming for a sequence of trials. The twofold difference in priming may be implemented as an increase of priming by a factor of $\sqrt{2}$ following a high reward, relative to the baseline, and a decrease of priming by a factor of $\sqrt{2}$ following a low reward, relative to the baseline (a factor of $\sqrt{2}$ makes priming following a high-reward twice as large as priming following a low-reward). For example, suppose in a sequence of 6 trials the target's quadrant was DCCBBD and

suppose participants received the following reward after each trial: high, low, low, high, low, low. The first two trials received no repetition priming because the target's quadrant did not repeat. Trial 3's target was in the same quadrant as that of trial 2's and therefore should be primed. The size of priming, however, should be modulated by reward that participants received on trial 2, which was a low reward, decreasing the baseline repetition priming. Trial 4 included no repetition, but trial 5's target was in the same quadrant as that of trial 4's. The size of priming should be modulated by reward that participants received on trial 4, which was a high reward, increasing the baseline priming. With this logic, one can simulate priming for each trial as a function of whether it matched the target quadrant of the preceding trials, and whether the preceding trials involved a high or a low reward. We can then compare all trials in which the target was in the high-reward quadrant with trials in which it was in the other quadrants.

Figure 9 (left) displays the simulated priming for the high-reward and low-reward quadrants across the 12 experimental blocks used in Experiment 3A (the button press version). In the simulation, the high-reward quadrant showed a consistently stronger priming effect than the low-reward quadrants in the training phase, $F(1, 15) = 48.05$, $p < 0.001$, $\eta_p^2 = 0.76$, and this difference disappeared in the testing phase when reward became random, $F < 1$. Figure 9 (right) plots the simulated priming difference, along with the observed RT difference, between the high-reward and low-reward quadrants. An ANOVA on the type of data (observed or simulated), quadrant, and block revealed no main effect of the type of data, $F < 1$. The simulated priming difference of 36 ms in the training phase fell within the 95 % confidence interval of the observed data (Table 2).

If reward influences attention primarily by modulating repetition priming, then its impact on spatial attention in our paradigm would necessarily be minimal. In fact, the simulated priming difference of 36 ms may have been an overestimation. This is because although priming is spatially diffused, reward modulation of priming appears to be spatially precise (Hickey

Fig. 9 Simulating reward modulation of repetition priming. Left: size of repetition priming in the high-reward and low-reward quadrants, assuming that repetition priming increased by a factor of $\sqrt{2}$ following a high reward and decreased by the same factor following a low reward.

Right: Differences in simulated priming between the high-reward and low-reward quadrants were in line with differences in the observed RT difference between these conditions. Error bars show ± 1 standard error of the mean. Some error bars may be too small to see

et al., 2014). In our study the target could appear in 100 possible locations, rendering exact location repetition extremely unlikely (just 1 %). If reward had modulated priming on only the 1 % of trials in which an exact location repetition had occurred, then the priming difference between the high- and low-reward quadrants should have been even smaller than the 36-ms effect simulated here.

Our simulation also indicates that stronger effects of reward may be observed by increasing either baseline priming or the reward modulator. Baseline priming may increase when the target repeats its exact location rather than its quadrant. Reward modulation may increase if we had substituted monetary reward with primary reward (such as water for thirsty participants, Seitz, Kim, & Watanabe, 2009) or if the magnitude of the reward had substantially increased.

Multiple-systems view of spatial attention

Consistent with previous theories of spatial attention, our study suggests that multiple sources can drive selective attention to prioritize some locations for processing. However, our study highlighted the substantial differences in the strength of spatial attentional guidance from different sources. Goal-driven attention was associated with a large RT facilitation, a large effect size (Cohen's $d > 6$), and a 95 % confidence interval that was largely nonoverlapping with that produced by statistical learning or incidental reward learning. Our findings therefore provide strong support for theories that emphasize goal-driven, top-down control of attention (Folk, Remington, & Johnston, 1992; Jiang, Swallow, & Rosenbaum, 2013; Kunar et al., 2007; Rosenbaum & Jiang, 2013).

Unlike some previous research that has cast doubt on the implicit guidance of attention (Kunar et al., 2007), our study shows that incidental learning can yield strong and durable spatial attentional biases. Location probability learning produced an attentional effect that was large in overall mean and strong in effect size (Cohen's $d > 2$). The effect was about half of that induced by successful task goals, but it did not depend on explicit learning. Location probability cuing therefore provides a complementary mechanism to goal-driven attention. Given that probability cuing does not depend on awareness (this study; see also Jiang & Swallow, 2014; Rosenbaum & Jiang, 2013) or on working memory resources (Won & Jiang, 2015), it may be particularly useful when goal-driven attention is absent or compromised (e.g., as in patients with hemifield neglect, Geng & Behrmann, 2002). The simulation approach taken here allows us to partition location probability learning into a short-term, repetition priming effect and a long-term learning effect, uniting seemingly conflicting previous findings on location probability cuing (Jones & Kaschak, 2012; Walthew & Gilchrist, 2006).

The weakest effect of spatial attentional guidance in our study came from incidental learning of monetary reward.

The RT advantage for targets in the high-reward quadrant relative to the low-reward quadrants was near zero, although the 95 % confidence interval was wide. Assuming that (retrospective) monetary reward shapes attention primarily by modulating the size of repetition priming, we simulated the differences in priming between the high- and low-reward quadrants and found that it was small. Our priming simulation also suggested that reward-based effects could be stronger if one strengthened repetition priming (e.g., by using exact location repetitions) or if stronger reward was used.

Effect sizes differed dramatically between the incidental learning (Experiments 2 and 3) and goal-driven experiment (Experiment 4), suggesting that these experiments may have tapped into the two extremes of reward-based spatial attention. Although we did not explore the intermediate cases, our study makes predictions about what those cases may be. Specifically, intermediate effects may be observed when (i) instructions are less informative or more variable across trials, reducing the efficacy of goal-driven attention, and (ii) reward modulation of location repetition priming is strengthened with stronger reward or fewer locations. These predictions should be examined in future research.

By testing three major sources of top-down attention, our study allowed us to propose the multiple-systems view of top-down attention. This view extends our previous “dual-system view” of attention to include not just task goals and search habit (Jiang, Swallow, & Capistrano, 2013) but also repetition priming as a major component of top-down attention (Johnston & Dark, 1986; Kristjánsson & Campana, 2010; Treisman, 1964). Figure 10 shows a schematic illustration of the multiple-systems view. According to this view, top-down attention may be directed on the basis of three sources: goals, priming, and search habit. Task goals influence declarative attention and are the most effective driver of spatial attention. Repetition priming is an automatic mechanism that draws attention to the previous target's location (Kristjánsson & Campana, 2010; Maljkovic & Nakayama, 1996). The effect is relatively short-lived but allows the consequence of one's action (such as monetary reward) to modulate spatial attention. Finally, people may develop a search habit by habitually shifting spatial attention in certain directions (the premotor attention; Rizzolatti, Riggio, Dascola, & Umiltà, 1987).

The multiple-systems view also includes a tentative mapping between experimental manipulations and the three

Fig. 10 Multiple-systems view of top-down attention

subsystems of top-down attention. Monetary reward can drive attention by changing task goals or by modulating repetition priming. Even relatively low reward can mobilize goal-driven attention. However, this mechanism is contingent on participants' ability to successfully integrate reward with task goals. In the absence of explicit guidance, reward can shape attention by modulating repetition priming. This effect is smaller and may become significant only when repetition priming is strong (e.g., with exact location repetition, and with the use of a small number of locations) and likely depends on the presence of high-valued reward and on exact location repetitions. Similarly, visual statistical learning can drive attention by changing each subsystem. High-degrees of explicit awareness, along with an intention to prioritize the high-probability locations, may allow VSL to mobilize goal-driven attention. In the absence of explicit knowledge, VSL can change attention by modulating repetition priming or by changing search habit.

The three drivers of attention have complementary roles in everyday tasks. Goal-driven attention is powerful when the goal is successfully implemented but is less useful when goals fail or when working memory resources are taxed (Won & Jiang, 2015). Location probability learning is constrained by the presence of statistical regularities. It is not useful if the target's location is random. Repetition priming is widespread, although the effect may be short-lived. The integration of all three drivers of spatial attention may be the reason why, despite severe capacity limitations in attention and working memory, humans are largely successful in daily activities.

Acknowledgments Research reported in this publication was supported in part by NIH R03 MH102583. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. Roger Remington was supported by the Australian Research Council Discovery Grant DP130101001 and the University of Queensland Vice-Chancellor's Senior Research Fellowship 004807-3558. We thank Windy Torgerud, Hyejin Lee, Michaela DeBolt, Gracelynn Goh, Anthony Asaad, Xinyi Wang, and Yang Yang for help with data collection. Thanks also to Windy Torgerud, Wilma Koutstaal, and Steve Engel for discussions.

References

- Anderson, B. A. (2013). A value-driven mechanism of attentional selection. *Journal of Vision*, 13(3). doi:10.1167/13.3.7
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences of the United States of America*, 108(25), 10367–10371. doi:10.1073/pnas.1104047108
- Asgeirsson, A. G., & Kristjánsson, A. (2014). Random reward priming is task-contingent: The robustness of the 1-trial reward priming effect. *Frontiers in Psychology*, 5, 309. doi:10.3389/fpsyg.2014.00309
- Awah, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437–443. doi:10.1016/j.tics.2012.06.010
- Bisley, J. W., & Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review of Neuroscience*, 33, 1–21. doi:10.1146/annurev-neuro-060909-152823
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436.
- Camara, E., Manohar, S., & Husain, M. (2013). Past rewards capture spatial attention and action choices. *Experimental Brain Research*, 230(3), 291–300. doi:10.1007/s00221-013-3654-6
- Chelazzi, L., Eštočinová, J., Calletti, R., Lo Gerfo, E., Sani, I., Della Libera, C., & Santandrea, E. (2014). Altering spatial priority maps via reward-based learning. *Journal of Neuroscience*, 34(25), 8594–8604. doi:10.1523/JNEUROSCI.0277-14.2014
- Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research*, 85, 58–72. doi:10.1016/j.visres.2012.12.005
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, 36(1), 28–71. doi:10.1006/cogp.1998.0681
- Chun, M. M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, 10(4), 360–365. doi:10.1111/1467-9280.00168
- Chun, M. M., & Jiang, Y. (2003). Implicit, long-term spatial contextual memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(2), 224–234.
- Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science*, 17(3), 222–227. doi:10.1111/j.1467-9280.2006.01689.x
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222. doi:10.1146/annurev.ne.18.030195.001205
- Egeth, H. E., & Yantis, S. (1997). Visual attention: Control, representation, and time course. *Annual Review of Psychology*, 48, 269–297. doi:10.1146/annurev.psych.48.1.269
- Failing, M. F., & Theeuwes, J. (2014). Exogenous visual orienting by reward. *Journal of Vision*, 14(5), 6. doi:10.1167/14.5.6
- Fecteau, J. H., & Munoz, D. P. (2006). Saliency, relevance, and firing: A priority map for target selection. *Trends in Cognitive Sciences*, 10(8), 382–390. doi:10.1016/j.tics.2006.06.011
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, 18(4), 1030–1044.
- Gabrieli, J. D. (1998). Cognitive neuroscience of human memory. *Annual Review of Psychology*, 49, 87–115. doi:10.1146/annurev.psych.49.1.87
- Geng, J. J., & Behrmann, M. (2002). Probability cuing of target location facilitates visual search implicitly in normal participants and patients with hemispatial neglect. *Psychological Science*, 13(6), 520–525.
- Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward, and information seeking. *Journal of Neuroscience*, 34(46), 15497–15504. doi:10.1523/JNEUROSCI.3270-14.2014
- Gupta, P., & Cohen, N. J. (2002). Theoretical and computational analysis of skill learning, repetition priming, and procedural memory. *Psychological Review*, 109(2), 401–448.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2014). Reward-priming of location in visual search. *PloS One*, 9(7), e103372. doi:10.1371/journal.pone.0103372
- Hon, N., & Tan, C.-H. (2013). Why rare targets are slow: Evidence that the target probability effect has an attentional locus. *Attention, Perception, & Psychophysics*, 75(3), 388–393. doi:10.3758/s13414-013-0434-0
- Hutchinson, J. B., & Turk-Browne, N. B. (2012). Memory-guided attention: Control from multiple memory systems. *Trends in Cognitive Sciences*, 16(12), 576–579. doi:10.1016/j.tics.2012.10.003

- Jiang, Y. V., & Swallow, K. M. (2014). Changing viewer perspectives reveals constraints to implicit visual statistical learning. *Journal of Vision*, 14(12).
- Jiang, Y. V., Swallow, K. M., Rosenbaum, G. M., & Herzig, C. (2013a). Rapid acquisition but slow extinction of an attentional bias in space. *Journal of Experimental Psychology: Human Perception and Performance*, 39(1), 87–99. doi:10.1037/a0027611
- Jiang, Y. V., Swallow, K. M., Won, B.-Y., Cistera, J. D., & Rosenbaum, G. M. (2014). Task specificity of attention training: The case of probability cuing. *Attention, Perception, & Psychophysics*. doi:10.3758/s13414-014-0747-7
- Jiang, Y. V., Swallow, K. M., & Capistrano, C. G. (2013b). Visual search and location probability learning from variable perspectives. *Journal of Vision*, 13(6), 13. doi:10.1167/13.6.13
- Jiang, Y. V., Swallow, K. M., & Rosenbaum, G. M. (2013c). Guidance of spatial attention by incidental learning and endogenous cuing. *Journal of Experimental Psychology: Human Perception and Performance*, 39(1), 285–297. doi:10.1037/a0028022
- Johnston, W. A., & Dark, V. J. (1986). Selective attention. *Annual Review of Psychology*, 37(1), 43–75. doi:10.1146/annurev.ps.37.020186.000355
- Jones, J. L., & Kaschak, M. P. (2012). Global statistical learning in a visual search task. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1), 152–160. doi:10.1037/a0026233
- Kristjánsson, A., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics*, 72(1), 5–18. doi:10.3758/APP.72.1.5
- Kristjánsson, A., Sigurjónsdóttir, O., & Driver, J. (2010). Fortune and reversals of fortune in visual search: Reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Attention, Perception, & Psychophysics*, 72(5), 1229–1236. doi:10.3758/APP.72.5.1229
- Kunar, M., Flusberg, S., Horowitz, T., & Wolfe, J. (2007). Does contextual cuing guide the deployment of attention? *Journal of Experimental Psychology: Human Perception & Performance*, 33(4), 816–828. doi:10.1037/0096-1523.33.4.816
- Le Pelley, M. E., Pearson, D., Griffiths, O., & Beesley, T. (2014). When goals conflict with values: counterproductive attentional and oculomotor capture by reward-related stimuli. *Journal of Experimental Psychology: General*. doi:10.1037/xge0000037
- Lee, J., & Shomstein, S. (2013). The differential effects of reward on space- and object-based attentional allocation. *Journal of Neuroscience*, 33(26), 10625–10633. doi:10.1523/JNEUROSCI.5575-12.2013
- Lee, J., & Shomstein, S. (2014). Reward-based transfer from bottom-up to top-down search tasks. *Psychological Science*, 25(2), 466–475. doi:10.1177/0956797613509284
- Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, 22(1), 1–35. doi:10.1016/0010-0285(90)90002-L
- Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, 22(6), 657–672.
- Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Perception & Psychophysics*, 58(7), 977–991.
- Miller, J. (1988). Components of the location probability effect in visual search tasks. *Journal of Experimental Psychology: Human Perception and Performance*, 14(3), 453–471.
- Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation the consequences of value learning in an attentional blink task. *Psychological Science*, 20(8), 981–988. doi:10.1111/j.1467-9280.2009.02391.x
- Rizzolatti, G., Riggio, L., Dascola, I., & Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia*, 25(1A), 31–40.
- Rosenbaum, G. M., & Jiang, Y. V. (2013). Interaction between scene-based and array-based contextual cuing. *Attention, Perception, & Psychophysics*, 75(5), 888–899. doi:10.3758/s13414-013-0446-9
- Schacter, D. L., & Buckner, R. L. (1998). Priming and the brain. *Neuron*, 20(2), 185–195.
- Seitz, A. R., Kim, D., & Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*, 61(5), 700–707. doi:10.1016/j.neuron.2009.01.016
- Stadler, M. A., & Frensch, P. A. (1998). *Handbook of implicit learning*. Thousand Oaks: Sage.
- Stankevich, B. A., & Geng, J. J. (2014). Reward associations and spatial probabilities produce additive effects on attentional selection. *Attention, Perception, & Psychophysics*. doi:10.3758/s13414-014-0720-5
- Treisman, A. (1964). Monitoring and storage of irrelevant messages in selective attention. *Journal of Verbal Learning and Verbal Behavior*, 3(6), 449–459. doi:10.1016/S0022-5371(64)80015-3
- Turk-Browne, N. B. (2012). Statistical learning and its consequences. *Nebraska Symposium on Motivation*, 59, 117–146.
- Walthew, C., & Gilchrist, I. D. (2006). Target location probability effects in visual search: An effect of sequential dependencies. *Journal of Experimental Psychology: Human Perception and Performance*, 32(5), 1294–1301. doi:10.1037/0096-1523.32.5.1294
- Wolfe, J. (1998). What can 1 million trials tell us about visual search? *Psychological Science*, 9(1), 33–39. doi:10.1111/1467-9280.00006
- Wolfe, J. (2007). Guided Search 4.0: Current progress with a model of visual search. In *Integrated models of cognitive systems* (pp. 99–119). New York: Oxford University Press.
- Won, B.-Y., & Jiang, Y. V. (2015). Spatial working memory interferes with explicit, but not probabilistic cuing of spatial attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(3), 787–806. doi:10.1037/xlm0000040