How do magnitude and frequency of monetary reward guide visual search?

Bo-Yeong Won[1] · Andrew B. Leber[1]

Abstract How does reward guide spatial attention during visual search? In the present study, we examine whether and how two types of reward information—magnitude and frequency—guide search behavior. Observers were asked to find a target among distractors in a search display to earn points. We manipulated multiple levels of value across the search display quadrants in two ways: For reward magnitude, targets appeared equally often in each quadrant, and the value of each quadrant was determined by the average points earned per target; for reward frequency, we varied how often the target appeared in each quadrant but held the average points earned per target constant across the quadrants. In Experiment 1, we found that observers were highly sensitive to the reward frequency information, and prioritized their search accordingly, whereas we did not find much prioritization based on magnitude information. In Experiment 2, we found that magnitude information for a nonspatial feature (color) could bias search performance, showing that the relative insensitivity to magnitude information during visual search is not generalized across all types of information. In Experiment 3, we replicated the negligible use of spatial magnitude information even when we used limited-exposure displays to incentivize the expression of learning. In Experiment 4, we found participants used the spatial magnitude information during a modified choice task—but again not during search. Taken together, these findings suggest that the visual search apparatus does not equally exploit all potential sources of spatial value information; instead, it favors spatial reward frequency information over spatial reward magnitude information.

Interaction with our spatial environment is inherently goal driven, aimed at maximizing our behavioral outcomes. Consider a fisherman choosing between two equidistant locations to travel to and lower his lines. By a *frequency maximization* principle, he will pick the location that promises more fish caught in a given interval (all other things being equal). By a *magnitude maximization* principle, he will pick the location that promises individual fish that are more valuable to him (e.g., larger or tastier). Research has long examined the role of reward frequency and magnitude in learning (Crespi, 1942; Herrnstein, 1961). One classic demonstration of how these factors influence behavioral choice comes from Spear and Pavlik (1966), who used a T-maze procedure. Rats had to run through the "stem" of the T toward a choice point, where they had to go left versus right to obtain a potential food reward at the end of the chosen arm. The frequency manipulation placed a reward more frequently in one arm than the other; the magnitude manipulation placed rewards equally often but altered the number of food pellets in each arm. Results showed that the rats took advantage of both frequency and magnitude information to maximize their overall gains. Humans also incorporate these two sources of information in their decision-making behavior (see Ernst & Paulus, 2005).

How do these principles extend beyond classic scenarios such as behavioral decision making and spatial navigation? Here, we bring the question to the domain of human visual

✉ Bo-Yeong Won
    bywon@ucdavis.edu

[1]   Department of Psychology, Ohio State University, 1835 Neil Avenue, Columbus, OH 43210, USA

search. From an ecological standpoint, it is sensible to prioritize search to locations containing high-frequency and high-magnitude rewards. However, the cognitive machinery mediating visual search might not always be penetrable by the kinds of information used in higher level behavioral choice. For instance, it has been argued—albeit hotly debated—that memory for recently inspected locations does not strongly guide search (Horowitz & Wolfe, 1998; Wolfe, 2003). Additionally, random visual search is more efficient than volitionally constrained search (Wolfe, Alvarez, & Horowitz, 2000), suggesting that it might not be desirable or advantageous to integrate all sources of information into a spatial search task.

A review of the literature does show strong evidence for the use of frequency information. In studies of probability cueing, in which targets are presented more often in one display quadrant than the others, participants learn to prioritize their search to the high-frequency locations (e.g., Geng & Behrman, 2002; Jiang, Swallow, Rosenbaum, & Herzig, 2013; Miller, 1988). However, evidence for using magnitude information has been mixed. Just one study has reported that individuals learn to bias spatial attention to locations containing more valuable targets (Chelazzi et al., 2014). Note that, overall, the number of studies of attention and reward have surged in recent years, most of which have shown robust prioritization of rewarded nonspatial features (e.g., Anderson, Laurent, & Yantis, 2011; Della Libera & Chelazzi, 2006; Hickey, Chelazzi, & Theeuwes, 2010; Kiss, Driver, & Eimer, 2009; Navalpakkam, Koch, Rangel, & Perona, 2010). Given the traditionally strong focus of attention research on spatial attention (Bisley & Goldberg, 2010; Fecteau & Munoz, 2006), one might expect more published studies demonstrating learning of spatial reward magnitude. Moreover, one recent study has reported a failure to observe such learning (Jiang, Li, & Remington, 2015).

In the present study, we closely evaluate whether and how individuals learn to prioritize frequency and magnitude information within a visual search task. To advance our work beyond previous studies, we focused on a direct comparison of the usage of these two types of information. Specifically, when each source of information is manipulated in such a way that it returns a similar payoff schedule of monetary gains, will they drive behavior similarly? Our basic approach was to manipulate different levels of value across the display quadrants in two ways: In the frequency manipulation, targets appeared more often in some quadrants than others; in the magnitude manipulation, targets were more valuable in some quadrants than others. We contrast two predictions. On the one hand, the visual search apparatus could be guided simply by overall expected value, whether it be frequency or magnitude. On the other hand, some indicators of expected value could be exploited more than others (e.g., we could see greater exploitation of frequency than magnitude information).

Experiment 1A and 1B

This study is not the first to test the use of both frequency and magnitude information in a visual search task. Jiang et al. (2015) recently reported robust effects the former and negligible effects of the latter, although the tasks were not directly equated. One key difference was that Jiang et al.'s magnitude manipulation used monetary incentives whereas the frequency manipulation did not. The presence of reward could increase anxiety, influence attentional deployment, and/or affect decision/response processes (see Mathews & MacLeod, 2005, for a review). It is thus possible that participants learned the reward contingencies but did not express their learning; for instance, participants could have responded more conservatively to targets in expected high-reward locations to ensure success on these high-stakes trials. This strategy could mask RT evidence of spatial biasing to the high-reward quadrant. Moreover, because the incentive structure differed across the experiments, it is difficult to objectively compare the strength of the two manipulations. What if the reward manipulation was weaker than the frequency manipulation?

In the present study, we placed both reward magnitude (Experiment 1A) and frequency (Experiment 1B) manipulations in a similar monetary incentive context. The goal was for both sets of participants to earn equivalent amounts and be similarly motivated. To accomplish this, we closely matched the display quadrants in both tasks for expected value (EV). In each experiment, we included high-, low-, and neutral-EV quadrants. Our inclusion of neutral-EV quadrants allowed us a baseline for which to measure biasing toward high-EV quadrants and biasing away from low-EV quadrants.

For reward magnitude, targets appeared equally often in each quadrant, and the value of each quadrant was determined by the average points earned per target. For reward frequency, we varied how often the target was presented in each quadrant but held the average points earned per target constant across the quadrants. If participants are sensitive to the magnitude and/or frequency information, then we will find significant RT facilitation in high-EV quadrants and/or RT slowing in low-EV quadrants (compared to neutral-EV quadrants) in the respective experiments. We will directly compare the results of these two experiments to determine the relative differences in how participants use magnitude and frequency information.

Method

Participants Twelve individuals participated in Experiment 1A (seven females; mean age = 22.5 years), and 13 participated in Experiment 1B (six females; mean age = 20.8 years). One participant was excluded from Experiment 1B for low accuracy (>3 SD below the group mean). All participants reported normal or corrected-to-normal visual acuity and normal

hearing. The Ohio State University IRB approved this protocol. They were initially informed that their compensation would range between $15 to $20 for the 1.5 hour session, based on performance; however, upon completion, we paid everyone the full $20.

Materials Participants were tested in a dimly lit room. Stimuli were presented on a 24-in. LCD monitor and generated using MATLAB (www.mathworks.com), with PsychToolbox extensions (Brainard, 1997; Pelli, 1997).

Stimuli Search displays contained one target (a white *T* rotated 90° clockwise or counterclockwise) and 15 distractors (white *L*s rotated 0°, 90°, 180°, or 270°), on a gray background. Targets and distractors subtended 1.02° × 1.02° (all visual angles assume a typical viewing distance of 60 cm). Item locations were chosen randomly from a 10 × 10 invisible matrix (15.28° × 15.28°), with four items appearing in each quadrant. Target and distractor orientations were all selected randomly with replacement on each trial. The number 20 or 1 (font size: .92°), indicating reward points for a given trial, was displayed at the target location, in green for correct responses or in red for errors. Auditory feedback was either a three "chirp" sequence lasting 300 ms for 20-point correct responses, a single high-pitched 100-ms tone for a 1-point correct responses, or a 200-ms "buzz" for incorrect responses. One-s blank displays followed errors to discourage low accuracy.

Design In Experiments 1A and 1B, correct responses earned 20 points on 50 % of trials and 1 point on the remaining trials, yielding a mean EV of 10.5 points per trial. Had we distributed EV equally across the display, each quadrant would be worth 2.625 points per trial (i.e., each trial's EV equals the sum of EVs across all quadrants). Thus, we set 2.625 as the neutral EV.

To create low and high EV across quadrants in Experiment 1A (Magnitude), we held target frequency constant across four quadrants but unevenly distributed reward magnitude.

Specifically, targets in the *high-EV* quadrant earned 20 points on 90 % of trials and 1 point on 10 % of trials. Targets in the *low-EV* quadrant earned the opposite: 1 point on 90 % of trials and 20 points on 10 % of trials. *Neutral-EV* quadrants earned 1 and 20 points equally often. These contingencies yielded EVs of 4.525, 0.725, and 2.625 for high-, low-, and neutral-EV quadrants, respectively. We held target frequency and reward magnitude constant across four colors.

In Experiment 1B (Frequency), we maintained constant reward magnitude across all quadrants but manipulated target frequency. That is, targets were 20 points and 1 point equally often in all quadrants. However, targets appeared in the *high-EV* quadrant on 41.7 % of the trials, in the *low-EV* quadrant on 8.3 % of the trials, and in each of the *neutral-EV* quadrants on 25 % of the trials. These contingencies yielded EVs of 4.375, 0.875, and 2.625 for high-EV, low-EV, and neutral-EV quadrants, respectively, which are closely matched to Experiment 1A. When not identical, we erred on the side of more extreme values for low- and high-EV quadrants in the spatial reward experiment, to ensure as strong or stronger of a manipulation than spatial probability, which we expected to be robust from previous research (see Fig. 1a).

Procedure Participants initiated each trial by clicking on a small white square (.51° × .51°), which appeared near the screen center (jittered by .77°). After the click and a 500-ms delay, the search display appeared. Participants were instructed to press the left or right arrow for clockwise or counterclockwise targets, respectively. Upon response, the search array was removed, and the point value earned was displayed, along with the auditory feedback. Next, the cumulative total points were displayed at the screen center for 200 ms (see Fig. 1b).

Participants completed 12 practice trials before advancing to the main task, which consisted of eight blocks of 120 trials each.

Generation task After the main trials, participants were told about the reward manipulation and asked whether
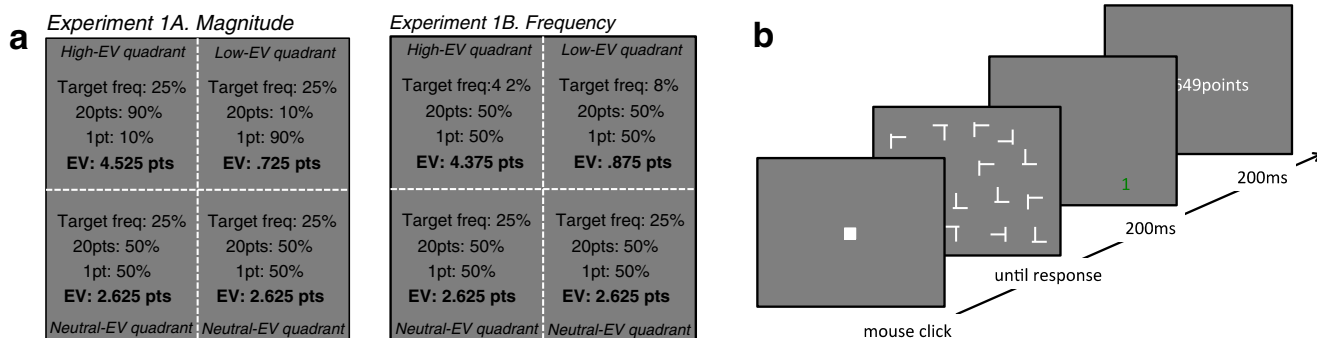
Fig. 1 **a**. Target frequencies, reward magnitudes, and associated EVs used in Experiment 1A (left) and Experiment 1B (right). Which quadrant was assigned as high-EV, low-EV, and neutral-EV quadrants were counterbalanced across participants. **b**. Sample trial events (see the Method section for additional details). (Color figure online)

they noticed any regularity. Next, they completed a 24-trial *generation task* (similar to Chun & Jiang, 2003), as follows. In Experiment 1A, participants were shown a search array and required to click on the target, which revealed two point values, 1 and 20. Participants had to then choose which of these they felt to best match the reward typically earned at that location. In Experiment 1B, participants were shown a search array of 16 *L*s, and they had to click on one *L* that they felt most likely would contain a target. No feedback was provided.

Results

Experiment 1A (magnitude)

Accuracy Accuracy in high-EV, low-EV, and neutral-EV quadrants was 99.1 %, 99.0 %, 99.0 %, respectively. An analysis of variance (ANOVA), including quadrant type (high-EV, low-EV, and neutral-EV) and blocks (1–8) as within-subject factors, showed neither significant main effects nor the two-way interaction, all $Fs \leq 1.0$.

RT Correct responses were analyzed after removing trials with RTs slower than 3 standard deviation above the mean, separately for each quadrant type; 1.7 % of trials were removed and remaining mean RTs are plotted in Fig. 2a. A Quadrant Type × Block ANOVA revealed a main effect of block, as RT became faster as the experiment progressed, $F(7, 77) = 9.19$, $p < .001$, $\eta_p^2 = .46$. However, there was

neither a main effect of quadrant type nor a two-way interaction, $Fs < 1$.

Generation task Although 6 of 12 participants reported noticing the regularity, the overall group did not reliably vary in their selection of 20 points versus 1 point across the high-EV, neutral-EV, and low-EV quadrants, $F(2, 22) = 2.28$, $p > .1$.
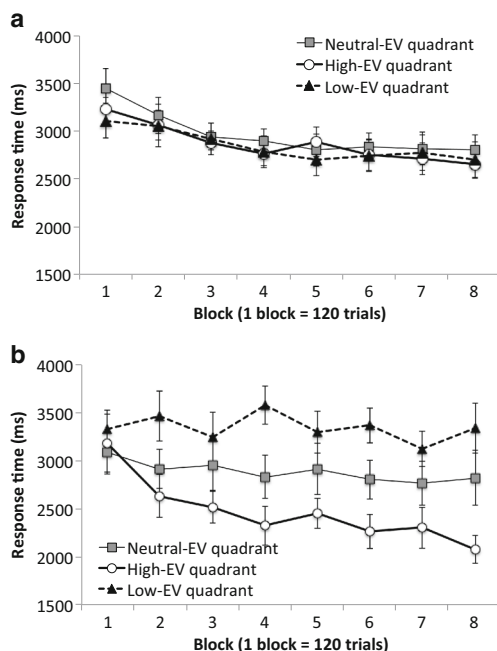
Experiment 1B (frequency)

Accuracy Accuracy in high-EV, low-EV, and neutral-EV quadrants was 99.5 %, 99.8 %, and 99.4 %, respectively. A Quadrant Type × Block ANOVA showed neither main effects nor an interaction, $ps > .05$.

RT In Fig. 2b, 1.7 % of trials were trimmed and remaining mean RTs are plotted in. As in Experiment 1A, the Quadrant Type Block ANOVA again showed a significant main effect of block, $F(7, 77) = 5.94$, $p < .001$, $\eta_p^2 = .35$. It also now showed a significant main effect of quadrant type, $F(2, 22) = 25.46$, $p < .001$, $\eta_p^2 = .70$, reflecting faster RTs to targets in more frequent quadrants. Additionally, the two-way interaction was significant, $F(14, 154) = 3.19$, $p < .001$, $\eta_p^2 = .23$, as the quadrant learning effect increased over time. These results replicate previous studies (Geng & Behrmann, 2005; Jiang et al., 2013).

Generation task Nine of 13 participants reported noticing regularity. However, the overall group did not reliably vary in their selection of one *L* that they felt most likely would contain a target across the high-EV, neutral-EV, and low-EV quadrants, $F(2, 22) = 2.23$, $p > .1$

Experiment 1A versus Experiment 1B

To directly compare RT across experiments, we conducted an Experiment (2 levels, between-subject) × Quadrant Type (3 levels, within-subject) × Block (8 levels, within-subject) ANOVA. Main effects of block and quadrant were significant, $F(2, 44) = 15.41$, $p < .001$, $\eta_p^2 = .41$; $F(7, 154) = 14.56$, $p < .001$, $\eta_p^2 = .40$, but that of experiment was not, $F < 1$. Critically, the Experiment × Quadrant Type interaction was significant, $F(2, 44) = 16.94$, $p < .001$, $\eta_p^2 = .44$, showing that the robust frequency effect of Experiment 1B was significantly greater than the negligible magnitude effect of Experiment 1A. Moreover, the three-way interaction was significant, $F(14, 308) = 1.78$, $p < .05$, $\eta_p^2 = .08$, as spatial attention became more biased toward the higher value quadrants over time in Experiment 1B but did not change in Experiment 1A.

Fig. 2 Results from Experiment 1A and 1B (panels A and B, respectively), showing RT as a function of quadrant type, across blocks. Error bars show ±1 *SE* of the mean

Experiment 2

Compared to the robust frequency learning expressed in Experiment 1B, we saw little evidence for magnitude learning in Experiment 1A. One possibility is that our visual search task is broadly insensitive not just to spatial reward magnitude learning but to many forms of magnitude learning. As mentioned in the introduction, several researchers have reported robust prioritization of *features* that are associated with greater monetary reward (e.g., Anderson, et al., 2011; Kiss et al., 2009; Navalpakkam et al., 2010). In Experiment 2, we test whether the visual search task used in Experiment 1 produces learning of feature magnitude. Instead of manipulating EV quadrants, we used high-, low-, and neutral-EV colors. The EV of each color was determined by the average points earned per target in each color. If the visual search task in Experiment 1 only favors frequency learning of any kind, not magnitude learning, then we will fail to find any RT facilitation in high-EV color in Experiment 2. If the negligible magnitude learning was specific to the biasing of spatial attention, we should observe significant feature magnitude learning.

Method

Participants Twelve individuals participated in Experiment 2 (eight females; mean age = 20 years).

Materials, stimuli, design, and procedure Methods were identical to those in Experiment 1A, except the following changes. All of the items in each quadrant were now assigned a unique color–red ([RGB]; [255 0 0]), green ([0 255 0]), blue ([0 0 255]), yellow ([255 255 0]). While all items in each quadrant had the same color, the colors changed randomly from trial to trial, such that color information was not associated with location information. The number 20 or 1 was displayed in white for correct responses, or in black for errors. The reward values were now assigned across different colors instead of quadrants. Specifically, targets in one color were deemed *high-EV*, earning 20 points on 90 % of trials and 1 point on 10 % of trials. Targets in another color were deemed *low-EV,* earning 1 point on 90 % of trials and 20 points on 10 % of trials. The remaining two colors were *neutral-EV*, earning 1 and 20 points equally often. These contingencies yielded the same EVs as Experiment 1A. The specific color–EV assignments were randomized across participants, and we held color frequency, target frequency, and reward magnitude constant across four quadrants. Participants completed 12 practice trials before advancing to the main task, which consisted of six blocks of 160 trials each.

Generation task The generation task was modified from Experiment 1A. Here, participants were told about the color reward manipulation and asked whether they noticed any regularity. Next, they completed a 32-trial *generation task*, in which they clicked on the target and had to choose which of the two point values best matched the reward typically earned for that color.

Results

Accuracy Accuracy in high-EV, low-EV, and neutral-EV colors was 99.4 %, 99.4 %, 99.5 %, respectively. A Color Type × Block (1–6) did not show any significant main effects or the two-way interaction, all $ps > .3$.

RT In Fig. 3, 1.8 % of trials were trimmed. Mean RTs for the three color types across blocks are plotted in Fig. 3. A Color Type × Block ANOVA revealed a significant main effect of block, as RT became faster as the experiment progressed, $F(5, 55) = 12.68$, $p < .001$, $\eta_p^2 = .54$, but there was no two-way interaction, $F(10, 110) = 1.26$, $p > .2$. It is important that we found a significant main effect of a color type, $F(2, 22) = 5.59$, $p = .01$, $\eta_p^2 = .34$, confirming that feature reward learning influenced target search.

Generation task Eight of 12 participants reported noticing the regularity. Additionally, color choices for the whole group significantly varied across EV-types (high, neutral, and low), $F(2, 22) = 12.22$, $p < .001$, $\eta_p^2 = .53$, yielding evidence for explicit knowledge of the color–reward relationship. However, further analysis revealed no correlation between generation task accuracy (i.e., how accurately they chose 20 pts for the high-EV color target and 1 pt for the low-EV color target) and magnitude of learning from the main experiment (i.e., RT for low-EV target color minus RT for high-EV target colors), $r = -.06$. Therefore, although explicit knowledge was present, it did not seem predictive of the presence of feature–reward learning.
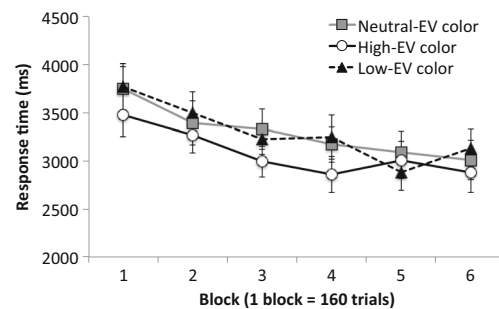
Fig. 3 Results from Experiment 2, showing RT as a function of color type, across blocks. Error bars show ±1 *SE* of the mean

Experiment 3

Experiment 2 demonstrates that visual search is sensitive to at least some form of magnitude learning, although not *spatial* reward magnitude. This brings us back to our initial question following Experiment 1. Why is learning and/or expression of spatial magnitude information so much weaker than for frequency information? One likely possibility is that exploiting frequency information confers a clearer behavioral benefit than does exploiting magnitude information. Specifically, by shifting attention first to locations that contain targets more frequently, mean RTs become faster. However, in the magnitude manipulation, targets appear equally often at each location, so magnitude information cannot be used by participants to predict target locations and complete the search faster. Thus, the incentive to use magnitude information is arguably weaker than for frequency information.

Note that participants are not motivated only to act on information that will speed search. Classic findings from the animal literature clearly underscore this point. In one such study (Goldstein & Spence, 1963), researchers alternately placed rats at the starting point in one of two 60-inch alleys on each trial. The opposite end of the alleys contained a food reward, which had consistently greater magnitude in one alley compared to the other. Results showed the rats ran faster down the alley associated with the higher magnitude reward (see also Davenport, 1962; Spear, 1964). Even though the rats could not choose the alley on each trial and could not use their behavior to exert any control over the outcome of the trial, they demonstrated greater motivation to navigate to the location with the greater expected value. The finding is not directly analogous to the current study, but it suggests that participants in the magnitude manipulation should be motivated to first check the high-EV quadrant, even though it does not affect the trial's outcome.

This speculation aside, we ran Experiment 3 to increase participants' incentive to prioritize the high-EV location in the magnitude task. Jiang et al. (2015) previously addressed this issue by only rewarding responses on trials whose RTs were faster than the participants' median RTs. The logic goes that participants only have enough time to search part of the display while the reward is still available. Thus, all things being equal, they would earn more by beginning their search at the highest magnitude quadrant. Despite this manipulation, Jiang et al. still failed to observe magnitude learning. However, one limitation to their approach is that less frequent rewards could have reduced the possibility of learning.

Here, we modified our Experiment 1A (magnitude) as follows. We provided unlimited search time for the entire first half of the experiment, allowing greater potential learning of the reward contingencies; in the second half,

we manipulated exposure duration, which reduced accuracy and thus incentivized a spatial bias toward the high-EV quadrant. If participants express learning of the reward contingencies, they should show a behavioral bias toward the high-EV quadrant and/or away from the low-EV quadrant in the limited search displays.

Method

Participants Twelve individuals (seven female; mean age 22.2 years) were included.

Materials, stimuli, and procedure Phase 1, consisting of four blocks (120 trials each), was identical to Experiment 1A. Phase 2 also consisted of four blocks (120 trials each), but the displays were now limited; exposures of 0.5 s, 1.25 s, and 2 s were each presented equally often and in random order. Twelve practice trials preceded Phase 1; before Phase 2, participants were instructed about the limited exposures and shown three sample trials before proceeding.
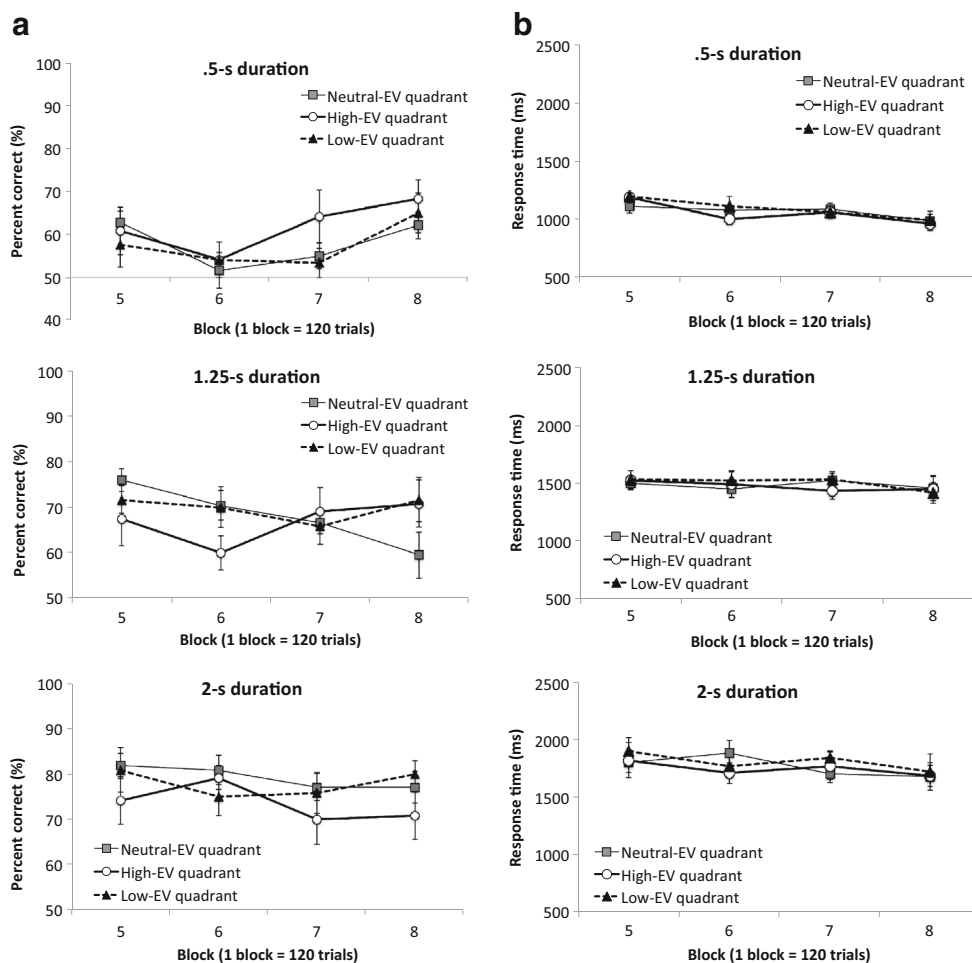
Results

Phase 1 (unlimited exposure)

Accuracy Accuracy across quadrant types and blocks is shown in Fig. 4a. A Quadrant Type × Block (1–4) ANOVA yielded neither main effects nor an interaction, $ps > .2$.
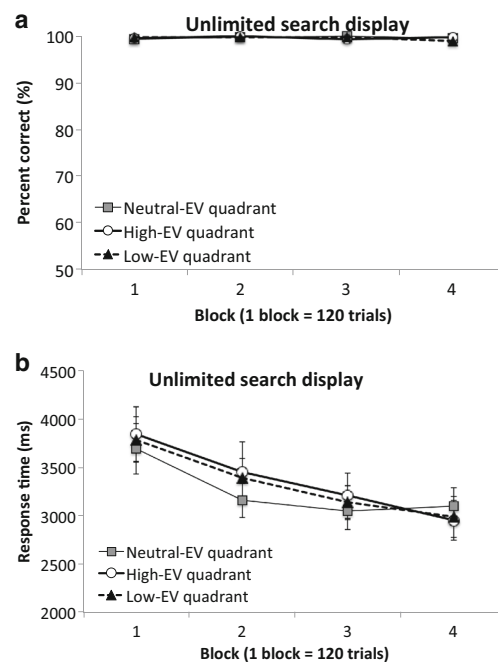
Fig. 4 Accuracy (**a**) and RT (**b**) results from Experiment 3, Phase 1. Error bars indicate ±1 *SE* of the mean

Fig. 5 Accuracy (**a**) and RT (**b**) results from Experiment 3, Phase 2, plotted separately by exposure duration. Error bars show ±1 *SE* of the mean

RT Trimming removed 1.7 % of trials in this phase. Figure 4b plots mean RT in this phase. The Quadrant Type × Block (1–4) ANOVA revealed a significant main effect of block, $F(3, 33) = 17.11$, $p < .001$, $\eta_p^2 = .61$, reflecting faster RTs over time. Neither the main effect of quadrant type nor interaction was significant, Fs < 1.

Phase 2 (limited exposure)

Accuracy Accuracy by quadrant type and block is plotted separately for each exposure duration in Fig. 5a. A Quadrant Type × Block (5–8) × duration ANOVA yielded only a main effect of duration, $F(2, 22) = 58.54$, $p < .001$, $\eta_p^2 = .84$, and a significant interaction between duration and block, $F(6, 66) = 2.80$, $p < .05$, $\eta_p^2 = .20$. Further analyses showed that the performance increased as the experiment progressed in the .5 s duration condition compared with both the 2-s duration, $F(3, 33) = 3.29$, p < .5, $\eta_p^2 = .23$, and the 1.25-s duration conditions, $F(3, 33) = 4.06$, p < .5, $\eta_p^2 = .27$. It is important to note, however, that quadrant type showed neither a main effect nor interaction with other factors, ps > .1.

RT Trimming removed 0.2 % of trials.[1] Mean RTs are shown in Fig. 5b. The Quadrant Type × Block (5–8) × Duration ANOVA revealed only a significant main effect of duration, $F(2, 22) = 265.38$, $p < .001$, $\eta_p^2 = .96$. No other main effects or interactions were significant, ps > .2.

Generation task Similar to Experiment 1A, although six participants reported a regularity, the overall group did not reliably vary in their selection of 20 points versus 1 point across the high-EV, neutral-EV, and low-EV quadrants, $F < 1$.

Experiment 4A and 4B

Experiment 3 provides clear evidence that spatial magnitude information is not prioritized, even when using such prioritization would increase overall earnings. In Experiment 4A and

---

[1] Here, we reported the RT data from only correct trials. Since the overall accuracy scores were much lower compared with those of Experiment 1 and Experiment 2, we further analyzed RTs from the complete data set. Regardless, we did not find any differences between the two analyses.

4B, we attempt to further demonstrate that the negligible spatial magnitude learning is specific to visual search. Here, we used a visual choice task in which observers did not search for a target but instead were asked to click any one of the items on each trial to reveal a reward. The stimulus displays were virtually identical to previous experiments, except all stimuli were now *L*s, and the participants were asked to click on just one of these items, at which point they received either a low or high reward. Also, like the previous magnitude experiments, we manipulated the reward value to be higher in some display locations than others. If participants are generally insensitive to the spatial distribution of value in the displays, then they will fail to demonstrate a bias toward the high-EV locations. However, if participants are only insensitive to the spatial magnitude during a search task, then we should find robust learning during visual choice.

For completeness, we directly compare spatial magnitude learning in visual search (Experiment 4A) to that of visual choice (Experiment 4B). The value contingencies were similar to the previous experiments, except we simplified the displays to include one high-EV and three low-EV quadrants. Also, to keep the two experiments comparable to one another, we modified the visual search task to consist of a click response on the target *T* (instead of a forced choice discrimination on the target orientation). Note that we only accepted a click response when the mouse hovered over the target; thus we did not measure response errors in this task.

Method

Participants Twelve individuals participated in Experiment 4A (nine females; mean age = 23.3 years), and 12 participated in Experiment 4B (10 females; mean age = 22.9 years).

Materials, stimuli, and procedure Both experiments consisted of 24 blocks, and each block consisted of 30 trials.

Experiment 4A (visual search) All aspects were identical with Experiment 1A except the following changes: The search display had only two types of quadrants, one *high-EV quadrant* and three *low-EV quadrants*. In Blocks 1–12, targets in the *high-EV quadrant* earned 20 points on 75 % of trials and 1 point on 25 % of trials (quadrant EV = 3.894); targets in the *low-EV* quadrant earned 1 point on 75 % of trials and 20 points on 25 % of trials (quadrant EVs = 1.438). In the second half of experiment (Blocks 13–24), targets in all quadrants earned 20 points on 50 % of trials and 1 point on 50 % of trials (all quadrant EVs = 2.635). Last, participants clicked on the target instead of pressing a key.

Experiment 4B (visual choice) All aspects were identical with Experiment 4A, except for the following changes: There were 16 *L*s in each display and no *T*s. Participants were instructed to click one *L* to receive a reward of either 1 or 20 points. We emphasized that participants should try to maximize their earnings. For the first half of experiment (Blocks 1–12), the *high-EV quadrant* contained three *L*s that when clicked were rewarded with 20 points and one *L* that when clicked was rewarded with 1 point. These values were matched to those of Experiment 4A. In the second half of experiment (Blocks 13–24), all quadrants had two 20-point *L*s and two 1-point *L*s.

Generation task The generation task for Experiment 4 was identical with that for Experiment 1A, except that in Experiment 4B, participants clicked a circled *L* instead of *T*.

Results

Experiment 4A (visual search)

Analysis focused on RT, which was measured as the time to click on the target. Note that, unlike Experiments 1 through 3, we do not present accuracy results here; recall that click responses were only accepted when the mouse hovered over the target, so errors were not possible. RT trimming removed 1.7 % of trials. Mean RTs across quadrant types and blocks are shown in Fig. 6a. A Quadrant Type × Block ANOVA revealed one main effect of block, $F(23, 253) = 3.59$, $p < .001$, $\eta_p^2 = .25$, meaning the overall search RT became faster as the experiment progressed. However, replicating the previous experiments, there was no RT difference between the two quadrant types, $F(1, 11) = 2.15$, $p > .1$. Quadrant type and block did not interact, $F(23, 253) = 1.18$, $p > .2$.

Generation task Although five of 12 participants reported noticing the regularity, the overall group did not reliably vary in their selection of 20 points versus 1 point across the high-EV and low-EV quadrants, $F < 1$.

Experiment 4B (visual choice)

Choice proportion across quadrant types and blocks is shown in Fig. 6b. First, participants chose an *L* in the high-EV quadrant significantly more frequently than in each of the low-EV quadrants (52.9 % vs. 15.7 % of trials, respectively; chance = 25 % per quadrant), $t(11) = 3.80$, $p < .005$. Additionally, the difference between high-EV and low-EV quadrant choices varied as a function of block, $F(23, 253) = 2.55$, $p < .001$, $\eta_p^2 = .19$, reflecting that the bias toward the high-EV quadrant grew over roughly the first 10 blocks.

Generation task Nine of 12 participants reported noticing the regularity. However, the overall group still did not reliably vary in their selection of 20 points versus 1 point across the high-EV and low-EV quadrants, $F(1, 11) = 1.87$, $p > .1$.

Fig. 6 a. Search RTs from Experiment 4A. **b**. Choice proportion in Experiment 4B, representing the total proportion of choice a given quadrant was selected. Low-EV data are averaged across three quadrants. Error bars show ±1 *SE* of the mean
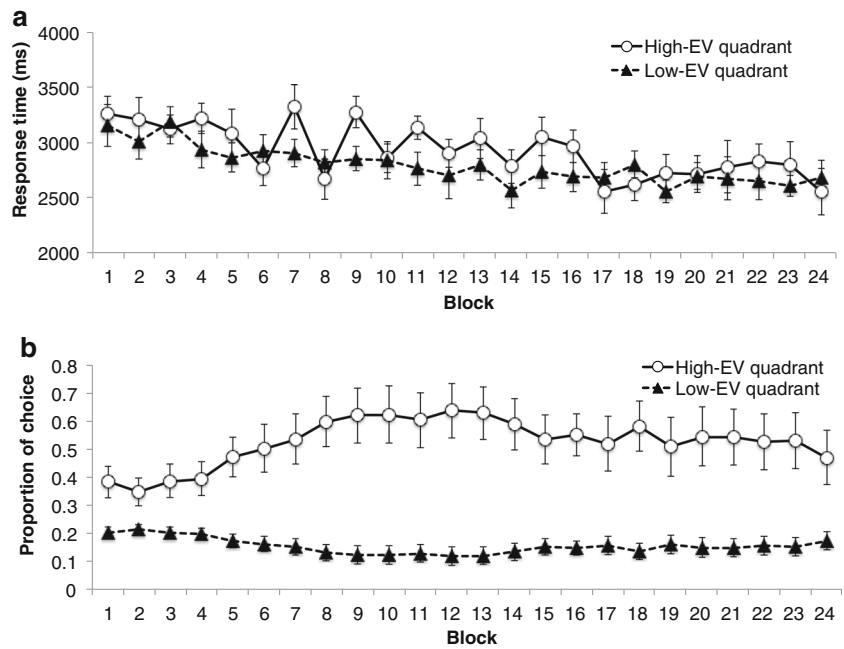
Experiment 4A vs. Experiment 4B

To compare results from Experiments 4A and 4B, we converted the respective dependent measures of interest (RT and choice frequency) to arbitrary but comparable *learning efficiency units* for each block (see Fig. 7). Specifically, in visual search (Experiment 4A), we subtracted the mean RT in high-EV quadrant from that in low-EV quadrants, and then divided by the sum of the two mean RTs from high-EV quadrant and low-EV quadrant (i.e., low-EV quadrant's RT – high-EV quadrant's RT) / (low-EV quadrant's RT + high-EV quadrant's RT)). For visual choice (Experiment 4B), we subtracted the mean choice frequency in the low-EV quadrants from that in the high-EV quadrant and then divided the result by the sum of the two types of choice frequencies (i.e., high-EV quadrant's frequency – low-EV quadrant's frequency) / (high-EV quadrant's frequency + low-EV quadrant's frequency)).

After computing the learning efficiency unit measures, we conducted an Experiment (between-subject) × Block (within-subject) ANOVA. We found a significant main effect of experiment, $F(1, 22) = 14.69$, $p = .001$, $\eta_p^2 = .40$, but there was no effect of block, $F(23, 506) = 1.40$, $p > .1$. Moreover, we

found a significant interaction between experiment and block, $F(23, 506) = 1.76$, $p < .05$, $\eta_p^2 = .074$, reflecting the increasing learning in visual choice over time while visual search remained flat.

General discussion

The present results clearly show that people do not treat all sources of spatial value information equally. Specifically, in Experiment 1, when we carefully equated frequency and magnitude manipulations, the exploitation of the former was quite robust, but results for the latter were much weaker—even negligible. We also showed, in Experiment 2, that the lack of spatial magnitude learning in visual search did not generalize to feature magnitude learning, as we found significant prioritization of high-EV colors. Even when we implemented limited-exposure displays to incentivize the use of magnitude information in Experiment 3, we still failed to find spatial biasing. Experiment 4 further dissociated the failure of magnitude learning in visual search with robust learning in a choice task.

Fig. 7 Learning efficiency from Experiment 4A (visual search) and from Experiment 4B (visual choice). Error bars show ±1 *SE* of the mean
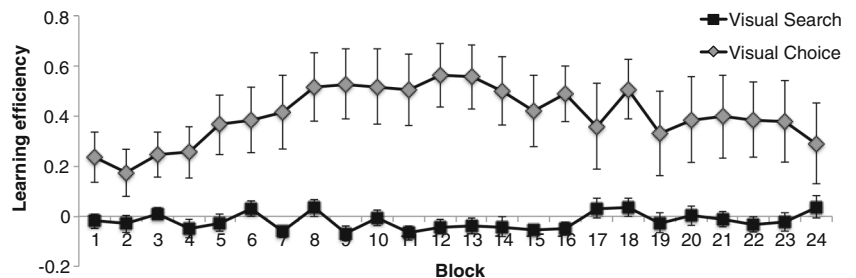
We feel these results fit parsimoniously with the larger literature on spatial learning and attention. Although there have been many demonstrations of frequency learning or probability cueing (Geng & Behrmann, 2002, 2005; Jiang et al., 2013; Miller 1988; Won & Jiang, 2015), only a few reports of spatial reward magnitude learning have emerged. Camara, Manohar, and Husein (2013) and Hickey, Chelazzi, and Theeuwes (2014) reported that spatial reward on one trial influenced attentional allocation on subsequent trials, although such effects were on a short timescale and did not reflect long-term spatial biasing. Both Shomstein and Johnson (2013) and Lee and Shomstein (2013) used an object-based cueing procedure, modified from Egly, Driver, and Rafal (1994), in which invalidly cued targets could appear at equidistant locations either within the cued object or in an uncued object. When invalid same-object locations were rewarded more than invalid different-object locations, spatial attention was biased to the *same* object; when this contingency was reversed, the spatial attention effects also reversed, demonstrating that reward magnitude determined attentional allocation. This effect was not one of absolute space but rather was defined with respect to the cued object location and thus may reveal properties of allocation in object-centered coordinates; future work should more closely compare effects of reward within object-centered versus purely space-based coordinates.

Of greatest direct relevance is the study by Chelazzi et al. (2014). Their task included a training procedure requiring a difficult visual discrimination on a target appearing in one of eight possible locations (with the remaining locations populated with distractors). Targets at each location were paired with an 80 %, 50 %, or 20 % high-reward schedule. After training, participants searched the same eight locations, now for one or two letter targets on each trial among a variety of non-alphanumeric distractors (e.g., #, %), for very brief exposures (70 ms). Performance was well below ceiling. In the two-target condition, the researchers found significantly greater accuracy for targets at the high-rewarded locations and significantly worse accuracy for those at the low-rewarded locations. Similar effects were not seen in the single-target trials, which the researchers interpreted to be due to lower competitive demands on spatial attention.

The Chelazzi et al. (2014) results demonstrate that spatial reward magnitude can be learned, although perhaps only via a sophisticated and highly demanding procedure. This is contrasted with the failure to find such learning with a larger scale, higher accuracy visual search task used by us and Jiang et al. (2015). Most important for the present purposes, the Chelazzi study did not compare the reward magnitude manipulation to a frequency manipulation. We do not intend for our present results to rule out the possible existence of magnitude learning; rather, we demonstrate that spatial magnitude learning is comparatively weaker than frequency learning, feature magnitude learning, and visual choice learning.

Given that we matched expected value across multiple manipulations, such that they were equally informative, why would participants not prioritize all information types equally? Although we cannot definitively answer this question, we do offer some speculation. First, a number of neurophysiological and imaging studies have implicated distinct neural substrates in the processing of magnitude versus frequency information during behavioral choice. For instance, subcortical structures including nucleus accumbens have been linked to the former and cortical prefrontal structures have been linked to the latter (Knutson, Taylor, Kaufman, Peterson, & Glover, 2005; see also Smith et al., 2009; Yacubian et al., 2007). Perhaps the mechanisms mediating learning and/or expression of learning in visual search interface in distinct ways with the unique substrates representing magnitude and frequency information. Second, as we mentioned in the introduction to Experiment 3, the exploitation of frequency information always produces a clear behavioral advantage, whereas exploiting magnitude information need not (especially when accuracy is at ceiling). We devised an experiment to artificially introduce time pressure to incentivize the use of magnitude information, but it is possible that the visual system did not evolve under such circumstances. That is, in the real world, exploitation of magnitude information might not usually improve behavioral outcomes; given this possibility, our visual search mechanisms may have developed without a significant evolutionary pressure to exploit spatial magnitude information. Thus, while such a pressure can be contrived in the laboratory, the visual system could be relatively insensitive to it.

Of course, any broad distinction between how magnitude and frequency information are used during search does not explain the successful learning observed in the color-based manipulation of Experiment 2. It could be that specific object properties (e.g., shape, color) acquire value in a way that spatial locations do not. This could make sense from an ecological standpoint: Objects inherently maintain their value over time, whereas spatial locations could be more variable (e.g., berries found in a new location are just as rewarding as the berries previously harvested from a now-exhausted location).

One further question that cannot be addressed by our current data is whether our participants either failed to (incidentally) learn the magnitude contingencies during visual search or express their learning. We aim to address this important question in future research.

In conclusion, we offer the provocative finding that despite its great complexity and sophistication, the visual system does not equally exploit all potential sources of spatial value information. While we compared only two such types of information, further types of potential spatial prioritization remain open for future investigation. This broader research venture will contribute vital insights into the intersection between reward learning and spatial attention.

References

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences, USA, 108,* 10367–10371. doi:10.1073/pnas.1104047108

Bisley, J. W., & Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe. *Annual Review Neuroscience, 33,* 1–21. doi:10.1146/annurev-neuro-060909-152823

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

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. *The Journal of Neuroscience, 34*(25), 8594–8604. doi:10.1523/JNEUROSCI.0277-14.2014

Chun, M. M. & Jiang, Y. (2003). Implicit, long-term spatial contextual memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*(2), 224–234. doi:10.1037/0278-7393.29.2.224

Crespi, L. P. (1942). Quantitative variation in incentive and performance in the white rat. *American Journal of Psychology, 55,* 467–517.

Davenport, J. W. (1962). The interaction of magnitude and delay of reinforcement in spatial discrimination. *Journal of Comparative and Physiological Psychology, 55,* 267–273. doi:10.1037/h0043603

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

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General, 123*(2), 161–177.

Ernst, M., & Paulus, M. P. (2005). Neurobiology of decision making: A selective review from a neurocognitive and clinical perspective. *Biology Psychiatry, 58*(8), 597–604. doi:10.1016/j.biopsych.2005.06.004

Fecteau, J. H., & Munoz, D. P. (2006). Salience, 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

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. doi:10.1111/1467-9280.00491

Geng, J. J., & Behrmann, M. (2005). Spatial probability as an attentional cue in visual search. *Perception & Psychophysics, 67*(7), 1252–1268. doi:10.3758/BF0319355710.1016/j.tics.2012.10.003

Goldstein, H., & Spence, K. W. (1963). Performance in differential conditioning as a function of variation in magnitude of reward. *Journal of Experimental Psychology, 65,* 86–93.

Herrnstein, R. J. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour, 4,* 267–272.

Hickey, C., Chelazzi, L., & Theeuwes, J. (2010). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience, 30*(33), 11096–11103. doi:10.1523/JNEUROSCI.1026-10.2010

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. eCollection 2014

Horowitz, T. S., & Wolfe, J. M. (1998). Visual search has no memory. *Nature, 394*(6693), 575–577. doi:10.1038/29068

Jiang, V. Y., Li, Z. S., & Remington, R. W. (2015). Modulation of spatial attention by goals, statistical learning, and monetary reward. *Attention, Perception, and Psychophysics, 77*(7), 2189–2206. doi:10.3758/s13414-015-0952-z

Jiang, Y. V., Swallow, K. M., Rosenbaum, G. M., & Herzig, C. (2013). Rapid acquisition but slow extinction of an attentional bias in space. *Journal of Experimental Psychology: Human Perception and Performance, 39*(1), 87–99. doi:10.1037/a0027611

Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science, 20*(2), 245–251. doi:10.1111/j.1467-9280.2009.02281.x

Knutson, B., Taylor, J., Kaufman, M., Peterson, R., & Glover, G. (2005). Distributed neural representation of expected value. *Journal of Neuroscience, 25*(19), 4806–4812. doi:10.1523/JNEUROSCI.0642-05.2005

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

Mathews, A., & MacLeod, C. (2005). Cognitive vulnerability to emotional disorders. *Annual Review of Clinical Psychology, 1,* 167–195. doi:10.1146/annurev.clinpsy.1.102803.143916

Miller, J. (1988). Components of the location probability effect in visual search tasks. *Journal of Experimental Psychology: Human Perception & Performance, 14,* 453–471. doi:10.1037/0096-1523.14.3.453

Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences, 107*(11), 5232–5237. doi:10.1073/pnas.0911972107

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442.

Shomstein, S., & Johnson, J. (2013). Shaping attention with reward: Effects of reward on space- and object-based selection. *Psychological Science, 24*(12), 2369–2378. doi:10.1177/0956797613490743

Smith, B. W., Mitchell, D. G., Hardin, M. G., Jazbec, S., Fridberg, D., Blair, R. J. R., & Ernst, M. (2009). Neural substrates of reward magnitude, probability, and risk during a wheel of fortune decision-making task. *NeuroImage, 44*(2), 600–609. doi:10.1016/j.neuroimage.2008.08.016

Spear, N. E. (1964). Choice between magnitude and percentage of reinforcement. *Journal of Experimental Psychology, 68*(1), 44–52.

Spear, N. E., & Pavlik, W. B. (1966). Percentage of reinforcement and reward magnitude effects in a T MAZE: Between and within subjects. *Journal of Experimental Psychology, 71*(4), 521–528.

Wolfe, J. M. (2003). Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Sciences, 7*(2), 70–76. doi:10.1016/S1364-6613(02)00024-4

Wolfe, J. M., Alvarez, G. A., & Horowitz, T. S. (2000). Attention is fast but volition is slow. *Nature, 406*(6797), 691. doi:10.1038/35021132

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, & Cognition, 41*(3), 787–806. doi:10.1037/xlm0000040

Yacubian, J., Sommer, T., Schroeder, K., Gläscher, J., Braus, D. F., & Büchel, C. (2007). Subregions of the ventral striatum show preferential coding of reward magnitude and probability. *NeuroImage, 38*(3), 557–563. doi:10.1016/j.neuroimage.2007.08.007