Altering Spatial Priority Maps via Reward-Based Learning

Leonardo Chelazzi,^{1,2} Jana Eštočinová,^{1,3} Riccardo Calletti,¹ Emanuele Lo Gerfo,^{1,4} Ilaria Sani,¹ Chiara Della Libera,¹ and Elisa Santandrea¹

¹Department of Neurological and Movement Sciences, University of Verona, 37134 Verona, Italy, ²Italian Institute of Neuroscience, 37134 Verona, Italy,

³Centre of Applied Computer Science, Pavol Jozef Šafárik University, 04154 Košice, Slovakia, and ⁴Department of Economics, Management and Statistics, University of Milan Bicocca, 20126 Milan, Italy

Spatial priority maps are real-time representations of the behavioral salience of locations in the visual field, resulting from the combined influence of stimulus driven activity and top-down signals related to the current goals of the individual. They arbitrate which of a number of (potential) targets in the visual scene will win the competition for attentional resources. As a result, deployment of visual attention to a specific spatial location is determined by the current peak of activation (corresponding to the highest behavioral salience) across the map. Here we report a behavioral study performed on healthy human volunteers, where we demonstrate that spatial priority maps can be shaped via reward-based learning, reflecting long-lasting alterations (biases) in the behavioral salience of specific spatial locations. These biases exert an especially strong influence on performance under conditions where multiple potential targets compete for selection, conferring competitive advantage to targets presented in spatial locations associated with greater reward during learning relative to targets presented in locations associated with lesser reward. Such acquired biases of spatial attention are persistent, are nonstrategic in nature, and generalize across stimuli and task contexts. These results suggest that reward-based attentional learning can induce plastic changes in spatial priority maps, endowing these representations with the “intelligent” capacity to learn from experience.

Key words: cross-target competition; reward-based learning; spatial attention; spatial priority maps

Introduction

The notion of priority maps is at the core of most current theories and models of visual attention, at both the cognitive and neurobiological levels (Itti and Koch, 2001; Thompson and Bichot, 2005; Fecteau and Munoz, 2006; Gottlieb, 2007; Bisley and Goldberg, 2010; Awh et al., 2012; Ptak, 2012; Ptak and Fellrath, 2013). They dictate deployment of attention to specific spatial locations depending on the level of activation across the map, which in turn reflects the combined influence of stimulus driven activity and current goals of the individual (Thompson et al., 2005; Serences and Yantis, 2007; Ipata et al., 2009). Most importantly, they arbitrate which of a number of competing (potential) targets will receive the highest priority, enabling privileged perceptual processing and access to later stages of cognitive and motor activity (Itti and Koch, 2001; Thompson and Bichot, 2005; Fecteau and Munoz, 2006; Gottlieb, 2007; Bisley and Goldberg, 2010; Ptak, 2012).

Along with a theoretical attempt to clarify the interplay between reward and attention (Maunsell, 2004; Chelazzi et al., 2013), the past few years have witnessed a rapidly increasing effort to unravel the influence of reward on visual attention, and a

number of studies have demonstrated that the delivery of rewards in relation to specific low-level visual features, notably color and shape, determines robust effects on the attentional priority of those features, both in the short- and in the long-term (Anderson, 2013; Chelazzi et al., 2013). Despite a rapidly expanding literature addressing the influence of reward on the attentional processing of color and shape information and the resulting changes in the attentional priority of such features, thus far analogous effects of rewards in the spatial domain have never been reported (but for related findings, see Discussion). This is especially striking given that space has long been known to represent a sort of “primordial medium” for attentional deployment and forms of spatially directed attention have been far more thoroughly investigated than any other form of attention, including attention for nonspatial elemental features and integrated objects. Moreover, our understanding of the brain mechanisms controlling spatial attention has greatly developed over the past ~30 years, leading to the delineation of a coordinated network of interconnected cortical and subcortical brain structures (e.g., Nobre, 2001; Yantis and Serences, 2003; Beck and Kastner, 2009; Macaluso, 2010; Noudoost et al., 2010; Bisley, 2011; Chelazzi et al., 2011; Corbetta and Shulman, 2011; Knudsen, 2011; Macaluso and Doricchi, 2013; Squire et al., 2013). We therefore set out to fill this major gap by investigating the influence of a reward-based learning protocol on the attentional priority of locations in space. Specifically, we wished to test whether, by means of a suitable reward-based training regimen, we could produce enduring changes in priority maps that are responsible for directing spatial attention and for arbitrating selection under conditions of cross-stimulus competition.

Received Jan. 17, 2014; revised April 26, 2014; accepted May 14, 2014.

Author contributions: L.C., J.E., E.L.G., and C.D.L. designed research; J.E., R.C., E.L.G., and C.D.L. performed research; J.E., R.C., I.S., and E.S. analyzed data; L.C. and E.S. wrote the paper.

This work was supported by Fondazione Cariverona, Verona, Italy.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Leonardo Chelazzi, Department of Neurological and Movement Sciences, University of Verona, Strada Le Grazie 8, 37134 Verona, Italy. E-mail: leonardo.chelazzi@univr.it.

DOI:10.1523/JNEUROSCI.0277-14.2014

Copyright © 2014 the authors 0270-6474/14/348594-11\$15.00/0

To anticipate, we could provide clear-cut evidence that our reward-based training regimen leads to long-lasting alterations in spatial priority maps, in turn modifying the ability of our observers to locate and identify task-relevant information at the various locations in the visual field. Crucially, these effects were detected when rewards were no longer involved, and the task and stimulus material were different from those used during training. Overall, the reported findings support the notion that the observed changes in priority maps: (1) are not strategic in nature (i.e., are not aimed at maximizing the earning of reward during the experiment), (2) are long-lasting, well beyond the immediate availability of rewards, and (3) generalize across stimuli and task contexts.

Materials and Methods

Similar to our prior work (Della Libera and Chelazzi, 2009), the experiment comprised distinct phases: baseline, training, and test. Importantly, the tasks were the same between baseline and test, whereas a different task was used for the training phase. During training, participants were asked to locate a single target (present in all trials) among 7 nontargets and discriminate its internal structure (see below). Stimuli were represented by simple geometric shapes. After completion of each trial, a monetary reward was delivered, except when an error occurred. The reward could be high or low with the same overall probability. However, probability of earning high versus low reward varied across locations. For two locations, high reward was more likely than low reward, whereas for two other locations, low reward was more probable than high reward. Finally, for each of the remaining four locations, high and low rewards were equally likely (see below for details).

As already indicated, a different paradigm was used for both baseline and test sessions. For this purpose, we devised a psychophysical paradigm in which participants were to detect and report one or two targets briefly presented among an array of (7 or 6, respectively) nontargets on each trial. Targets were letters and digits, whereas nontargets were nonalphanumeric characters (see below). The logic of the task was as follows. First, repetition of the same task between baseline and test sessions enabled us to compare performance between before and after reward-based training. Furthermore, our chief interest was for assessing performance when two targets were presented. Therefore, based on pilot testing, the search array was visible for such short duration that, when two targets were presented, only one of them could be detected on most trials, which allowed us to ask the question as to why one versus the other target in the pair was given precedence on the given trial. In other words, we took advantage of a condition of limited exposure duration and cross-target competition to investigate which particular target would be prioritized on the given trial, and why. Our specific prediction was that, after training, a location associated with higher rewards during training would confer an advantage to a target presented at that location compared, for instance, to a target presented at a location associated with lower rewards. Single-target trials mainly served to assess performance across the various locations in the absence of cross-target competition, both before and after training.

Participants. Twenty-four healthy right-handed volunteers (9 males; mean $\pm$ SD age, 21.04 ± 1.9 years) took part in the experiment. They had normal or corrected-to-normal visual acuity. Most of the participants were students at the University of Verona. None of them had previously taken part in similar or related studies, and they were naive as to the purpose of the present research. All subjects gave their informed consent before participation. At the end of the experiment, participants received fixed monetary compensation for baseline and test sessions (€14). For the training sessions, they received monetary compensation that varied in a range between €30 and €50, depending on the overall accuracy of their performance during the training sessions; the amount of reward earned for each correct response was instead completely predetermined based on the specific reward schedule associated with each spatial location (see below).

Stimuli and apparatus. Stimuli for the baseline and test task were black capital letters, digits, and nonalphanumeric characters ($1.2^\circ \times 1.2^\circ$). Stimuli for the training task were simple geometric shapes, constituted of

two stacked triangular shape outlines ($1.2^\circ \times 0.7^\circ$): one filled in black and the other in white. For both tasks, stimuli were presented on a 17 inch CRT monitor (resolution: 1024×768 pixels; refresh rate: 75 Hz). The viewing distance was held constant at 57 cm by using an adjustable chin rest.

Baseline and test task. We used a variant of a visual search task (Fig. 1A) requiring participants to search for one (single target condition) or two targets (double target condition) among an array of seven or six distractors, respectively. Each trial started with a fixation display (500 ms) containing a white fixation cross ($0.3^\circ \times 0.3^\circ$) presented at the center of the screen against a black background and an iso-eccentric circular array (distance from fixation: 5°) of eight white squares ($1.54^\circ \times 1.54^\circ$), which marked the locations of the upcoming stimuli. The array was arranged so that each hemifield contained four white squares, two per quadrant. After the fixation display, eight stimuli were briefly presented (~ 70 ms) within each of the eight white squares and were immediately replaced by eight identical masking patterns (consisting of overlapping distractors), which remained visible until the participant delivered both behavioral responses (see below). The target stimuli were four capital letters (F, G, M, D) and four digits (2, 4, 7, 9). Distractor stimuli were seven nonalphanumeric characters (#, %, &, £, ¥, ©, ®). Participants were instructed to report all targets they could detect in the stimulus array. Single and double target conditions were presented randomly, and no indication was given to participants about the number of targets in the current trial. Participants were asked to deliver two responses in all trials, which could correspond to reporting two targets (double report), one target and a “null response” (single report), or no target (null report), based on the number of targets they could detect on the given trial. Specifically, they were asked to press one (or two) of eight keyboard keys corresponding to each specific target they could identify and/or to press the spacebar for a null report. Responses were nonspeeded, and performance was evaluated based only on the accuracy of report. No feedback was provided after behavioral responses. After an intertrial interval of 2500 ms, a new trial sequence started automatically.

The task comprised 192 single target trials and 448 double target trials, resulting in a total of 640 trials. For the single target condition, half of the targets were letters and half were digits, which appeared at all eight locations with the same frequency. For the double target condition, target pairing corresponded either to two letters, two digits, or one letter and one digit, again with the same frequency. Moreover, the spatial relationship between the two targets was completely balanced, so that pairs of targets appeared in any of the possible combinations of spatial locations the same number of times. To reiterate, targets appeared at all eight locations the same number of times. All experimental conditions were presented in a random order during the session.

Training session task and reward schedule. As in the baseline and test task, each trial started with a fixation display (500 ms) containing a white fixation cross ($0.3^\circ \times 0.3^\circ$) presented at the center of the screen against a black background, and a circular array of eight white squares ($1.54^\circ \times 1.54^\circ$). Then a stimulus array appeared, consisting of eight simple geometric shapes, each presented at one of the eight relevant locations (i.e., within each of the white squares), and lasted 300 ms (Fig. 1B). Participants were asked to discriminate the internal structure of the target shape. Specifically, targets were made of two stacked triangular shape outlines pointing upward, either with a white triangle (outlined in black) over a black triangle, or the reverse; the task of the participants was to report the color of the upper triangle (either black or white). The distractors were instead made of two triangles pointing downward, and in all cases the lower triangle was black whereas the upper triangle was white. In this task, participants were instructed to respond as quickly and accurately as possible by pressing one of two buttons on the computer numeric keyboard (“1” with the right index finger if the upper triangle of the target was white and “2” with the right middle finger if it was black). Correct responses were followed by a reward, which could be either high (10 points) or low (1 point). Reward feedback was indicated on the monitor for 1000 ms at the location of the previously presented target (Fig. 1B). Incorrect responses were followed by a 500 ms beep. To ensure that participants fully appreciated the reward value received after correct responses, on some trials, they were also asked to report the amount of earned reward in the current trial. This additional task was randomly

Figure 1. Experimental procedure. **A**, Illustration of the baseline and test paradigm. At the beginning of each trial, a fixation display appeared on the screen consisting of an iso-eccentric circular array of eight white squares, marking spatial locations for the upcoming stimuli. After 500 ms, the stimulus array was briefly displayed and was immediately followed by a mask. Participants were to identify one or two targets (letters or digits) among seven or six distractors (nonalphanumeric characters), respectively, by pressing the corresponding key on a standard computer keyboard. **B**, Illustration of the training paradigm. At the beginning of each trial, a fixation display identical to the one used in the baseline/test task appeared on the screen. After 500 ms, the stimulus array was presented on the screen for 300 ms. Participants were asked to discriminate the color of the upper triangle (either black or white) of the target stimulus as quickly and accurately as possible. Correct responses were followed by a reward feedback, which could be high or low, and the amount gained was indicated at the target location.

presented on ~10% of the correct trials along the experimental session, immediately after receiving the reward feedback. In this case, participants reported the amount of reward received by pressing one of two corresponding buttons on the computer numeric keyboard (button “4” for a high reward and button “5” for a low reward). A new trial sequence started after a 1000 ms intertrial interval.

Participants performed two training sessions, each comprising 800 trials. The stimulus display was designed so that both target shapes (with an upper white and an upper black triangle) appeared the same number of times and with the same probability at all spatial locations in the circular array. All conditions were presented in a random order during the session.

Although participants were told that the rewards received depended on their performance, the reward value that could be obtained by responding correctly on each trial was fully predetermined, so that gained values were completely decoupled from actual velocity and overall accuracy of responses, and balanced across all experimental conditions, being high or low with the same overall probability (50%). Crucially, in both training sessions, the schedule of reward assignment was systematically biased so that each of the eight spatial locations in the stimulus array was associated with a specific probability of receiving high versus low reward. Moreover, an overall imbalance in reward probability was associated with the two visual hemifields, whereby one hemifield (high reward) had a higher probability of leading to high reward relative to the other hemifield (low reward). For half of the participants, the high-reward hemifield was the right one; for the other half of the participants, it was the left hemifield. Within the high-reward hemifield (*Hh*), two locations led to high reward in 80% of cases (*80Hh*), and the remaining two locations led to high or low reward with equal probability (*50Hh*). Reward probability was reversed for the opposite hemifield; i.e., within the low-reward hemifield (*Lh*), two locations led to high reward only in 20% of the cases (*20Lh*), and the remaining two locations led to high or low reward with equal probability (*50Lh*). The labels we chose for each reward level convey two pieces of information: the number in the label corresponds to the probability of receiving a high reward for the correct identification of a target at the corresponding location, whereas the letters in the label indicate whether the location belongs to the high-reward hemifield (*Hh*) or to the low-reward hemifield (*Lh*). To avoid the possible confound stemming from a fixed association of the predetermined reward biases with specific spatial locations, we created different spatial configurations of reward assignments across the eight locations, so that individual participants had their own reward schedule. This allowed us to control for the possibility that any of the effects produced by the reward manipulation could result from a mere preference for some particular locations in space. An example of such an arrangement, for one participant, is shown in Figure 1C.

Procedure. Participants completed a block of 24 trials of practice before each experimental session. Each participant completed one pretraining (baseline) session and two training sessions on consecutive days and a post-training session (test), identical to the pretraining session, after a 4 d delay. Each session lasted about 1 h.

Results

Training

During the training phase, a reward feedback, consisting of either high or low reward, was delivered on each trial in turn for correct discrimination of the critical feature of a single target among distractors (see Materials and Methods). High and low rewards were delivered with different probabilities in association with

←

C, Example of reward contingencies for one participant. The probability of receiving high versus low reward was predetermined and systematically biased on the basis of the specific spatial location in the display, such that each location could be assigned to one of four reward categories: *80Hh*, *50Hh*, *50Lh*, and *20Lh*. In the specific example presented here, the probability of receiving a high reward was biased in favor of the left visual hemifield, which therefore corresponded to the high-reward hemifield for this example subject.

specific spatial locations, such that each spatial location was assigned to one of four categories (see Materials and Methods). To remind the reader, for two locations high reward was more likely than low reward (80Hh), for two locations low reward was more likely than high reward (20Lh), and for the remaining four locations high and low reward were equally probable; two of these locations belonged to the high-reward side (50Hh) and two belonged to the low-reward side (50Lh).

During training, performance of the participants started off relatively poor, with accuracy of target discrimination corresponding on average to $71.9 \pm 1.9\%$ (SEM) and RT corresponding on average to 902.7 ± 33.79 ms during the first training session, attesting to the demanding nature of the task, presumably necessitating focal attentional processing. We underscore that we devised a rather challenging task for the training phase because we wished to make less transparent to the participants their actual proficiency at the task, given the deceptive nature of the feedback. Participants then achieved an overall relatively high proficiency, with accuracy corresponding on average to $81.9 \pm 2.4\%$ and RT corresponding on average to 765.75 ± 23.43 ms during the second training session. To test for learning effects during the training phase, a two-way ANOVA, including the factors block (1–4, obtained by dividing each of the two training sessions in two subsequent, identical trial segments) and reward level (80Hh, 50Hh, 50Lh, and 20Lh), was performed on accuracy of report at the training task. As shown in Figure 2A, accuracy of report increased significantly along subsequent training blocks ($F_{(3,69)} = 44.966$, $p < 0.001$, $\eta_p^2 = 0.662$). Instead, accuracy was not influenced by reward level during training either in the form of a main effect ($F_{(3,69)} = 0.087$, $p = 0.967$) or in the form of a block by reward level interaction ($F_{(9,207)} = 0.328$, $p = 0.965$). A two-way ANOVA with the same factors was also performed on RT, leading to a fully compatible pattern of results. Although RT decreased significantly along subsequent blocks ($F_{(3,69)} = 23.492$, $p < 0.001$, $\eta_p^2 = 0.505$; Fig. 2A), there was no evidence of an effect of reward level (main effect of reward level: $F_{(3,69)} = 0.152$, $p = 0.928$; block by reward level interaction: $F_{(9,207)} = 0.657$, $p = 0.747$). Participants were nonetheless well aware of the reward feedback they received on each trial, as indicated by a high level of performance ($93 \pm 2\%$ and $95.5 \pm 1.8\%$, on average, during the first and second training session, respectively) in reporting the amount of reward earned on the current trial, when specifically queried (see Materials and Methods). Lack of reward-dependent variations in performance during training might suggest that the training protocol was unable to alter the attentional priority of the different locations in the display, which should yield unequal performance across locations as a result of the imbalanced reward schedule. However, as it will become clear below, this need not be the case.

Baseline and test

We first evaluated performance of the participants during the baseline session. In the single-target condition, on average participants correctly identified the target on $57.1 \pm 2.9\%$ of the trials. As assessed by a two-way ANOVA, including the main factors spatial location (1–8) and target type (letter vs digit), accuracy of report varied significantly across display locations ($F_{(7,161)} = 19.176$, $p < 0.001$, $\eta_p^2 = 0.455$), with better performance for most lateralized locations compared with locations closer to the vertical midline (Fig. 2B). This pattern of performance is fully consistent with previous reports of perceptual anisotropy across the visual field (e.g., Carrasco et al., 2001). No reliable difference emerged between correctly reporting a letter or a digit ($F_{(1,23)} =$

Figure 2. General performance effects. **A**, Training performance. Average reaction times (mean $\pm$ SEM; solid black line) and accuracy of report (mean $\pm$ SEM; dashed black line) are reported as a function of subsequent blocks of trials during the training sessions. Block 1 and block 2 belong to the first training session, whereas block 3 and block 4 belong to the second and last training session. **B**, Baseline performance in the single target condition. Average accuracy of report (%) in the single target condition is illustrated in the polar plot as a function of spatial location and target type. **C**, Baseline and test performance in the double target condition. The average incidence of different responses (double, single, and null report) across participants is illustrated (mean $\pm$ SEM), as assessed during the baseline (leftward stacked bar) and during the test (rightward stacked bar) sessions.

0.208 , $p = 0.652$) nor a reliable interaction between target type and spatial location ($F_{(7,161)} = 1.559$, $p = 0.151$; Fig. 2B).

Overall performance in the double-target condition is reported in Figure 2C. As evident from inspection of the leftward

stacked-column in the graph (baseline), on the majority of trials, on average $58 \pm 1.9\%$, participants correctly reported only one of two targets (single report). Participants were instead able to correctly report both targets (double report) on $25.2 \pm 3.2\%$ of the trials, and reported none of the targets (null report) on $16.8 \pm 2.3\%$ of the trials. Because preliminary analyses showed no difference between reporting two letters versus two digits, we pooled the data together for the homogeneous pairing and compared it against mixed pairing (i.e., one letter and one digit). We then performed a two-way ANOVA, including the factors hemifield (same vs opposite) and target pairing (homogeneous vs mixed) separately for the double, single, and null report instances. The probability of correctly reporting two targets (double report) was significantly higher on trials where they appeared in opposite visual hemifields (same = 0.18 ± 0.028 ; opposite = 0.306 ± 0.037 ; $F_{(1,23)} = 61.61$, $p < 0.001$, $\eta_p^2 = 0.728$), suggesting that, at least to some extent, the two hemispheres can elaborate information in parallel, in line with previous findings (e.g., Luck et al., 1989; Sereno and Kosslyn, 1991; Alvarez and Cavanagh, 2005; Kraft et al., 2005; Chakravarthi and Cavanagh, 2009). The probability of reporting two targets belonging to the same semantic category was higher than the probability of reporting two targets belonging to different semantic categories (homogeneous target pairing = 0.261 ± 0.031 ; mixed target pairing = 0.225 ± 0.033 ; $F_{(1,23)} = 14.973$, $p = 0.001$, $\eta_p^2 = 0.394$). The interaction hemifield by target pairing was nonsignificant ($F_{(1,23)} = 0.693$, $p = 0.414$). Complementary results were obtained from trials where subjects were able to identify only one target (single report): factors that increase the probability of a double report also directly decrease the probability of a single report. Performance was instead not influenced by either hemifield or target pairing on trials where none of the two targets was reported correctly, suggesting that null reports index an overall failure in attentional engagement.

There is a potential confound in the approach reported above in relation to the factor hemifield because the distance between targets presented in the array was not balanced between the two hemifield conditions. Specifically, the possibility that two targets appeared in adjacent locations was relatively more frequent when they were presented in the same hemifield, whereas a greater distance between targets was relatively more frequent when they were presented in opposite hemifields. For this reason, we repeated the same ANOVA with the factors hemifield and semantic category on a subset of trials leading to a double report, for which target distance was completely balanced across the two conditions, corresponding to combinations of spatial locations for which the two targets in the array were presented in nonadjacent locations with only one distractor between them. Also, for this subset of conditions, targets were presented with equal frequency at all spatial locations for both hemifield conditions. Results from this ANOVA fully replicated the opposite-hemifield advantage, as shown by a reliable main effect of the factor hemifield ($F_{(1,23)} = 69.969$, $p < 0.001$, $\eta_p^2 = 0.753$).

After a 4 d delay following the last training session, participants were engaged in a test session, identical to the baseline session, where reward feedbacks were not delivered (see Materials and Methods). As shown in Figure 2C (rightward stacked-column), overall performance of the participants in the double target condition improved in the test session relative to the baseline session (leftward stacked column), with a significant increase in the incidence of double reports (i.e., trials in which participants reported both targets correctly: $12.83 \pm 2.13\%$; $t_{(23)} = 6.037$, $p < 0.001$, $r = 0.783$), a significant decrease in the incidence of single reports (i.e., trials in which participants correctly

reported only one of two targets: $-7.76 \pm 2.08\%$; $t_{(23)} = -3.729$, $p = 0.001$, $r = 0.614$) and a trend toward a decrease in the incidence of null reports (i.e., trials in which none of the targets was correctly reported: $-5.08 \pm 2.66\%$; $t_{(23)} = -1.91$, $p = 0.069$, $r = 0.37$).

To obtain a global measure of performance improvement in the double target condition, we also calculated the percentage of correctly reported targets out of the total targets presented in this condition, for both the baseline and test session. This global measure was computed as the ratio between the total number of correctly reported targets (summed from single and double reports) in each session and the total number of targets presented (896 total targets, i.e., 2 targets on each of the 448 trials of the double target condition), multiplied by 100. Whereas the percentage of correctly identified targets in the double target condition corresponded to $66.51 \pm 3.19\%$ during the baseline phase, this percentage rose to $77.22 \pm 4.14\%$ during the test phase, reflecting a significant learning effect ($t_{(23)} = 4.0302$, $p < 0.001$, $r = 0.643$).

As stated in Materials and Methods, our main focus was on the double target condition where the two targets presented at given locations compete for attentional resources because this condition was exquisitely suited to measure any change in the priority of the competing spatial locations following our reward-based training. By comparing performance in the test versus baseline session, we wished to test the specific prediction that our reward-based training regimen led to long-lasting alterations in spatial priority maps, for instance, by conferring a relative advantage to targets presented in highly rewarded locations (80Hh) with respect to poorly rewarded locations (20Lh), when two targets were presented together in the double target condition. Therefore, we first concentrated on those trials where two targets were presented: one at a 80Hh spatial location and one at a 20Lh spatial location (80Hh-20Lh pairs), and only one target was correctly reported by the participants (single report). This condition is the one in which a maximal difference in priority between the two target locations should follow the imbalanced delivery of rewards during the training phase, and the singly reported target should reflect such difference in priority. Specifically, we assessed for any variation in the probability of correctly reporting the target displayed at a 80Hh location in the test phase with respect to the baseline phase. Crucially, in 80Hh-20Lh pairs, the probability that the single reported target was the one displayed at the 80Hh spatial location in the pair increased by 0.084 (a 8.4% increase) from the baseline to the test phase ($t_{(23)} = 3.084$, $p = 0.005$, $r = 0.541$; Fig. 3A); likewise, the probability that the single reported target was the one displayed at the 20Lh spatial location in the pair decreased by the same amount. In other words, targets appearing in spatial locations that were more likely associated with high reward during training gained a strong competitive advantage with respect to targets displayed in spatial locations more frequently associated with low reward during training.

The same approach was applied to trials where two targets were displayed, one at a 50Hh and one at a 50Lh spatial location in the array, again leading to a single report. Here we aimed to test whether our reward-based training regimen conferred a general competitive advantage to the high-reward hemifield with respect to the low-reward hemifield, despite the fact that the competing locations were themselves associated with equal reward probability during training. In this case, the probability that the single reported target was the one displayed at the 50Hh spatial location in 50Hh-50Lh pairs was completely unchanged from the baseline to the test phase (Δ probability = -0.0028 ; $t_{(23)} = 0.084$, $p =$

Figure 3. Reward effects. **A**, Δ probability of report for critical reward-associated spatial locations in the single report condition. The histogram represents the differences in the probability of correctly reporting only the target in a 80Hh location in 80Hh-20Lh pairs (left) and the target in a 50Hh location in 50Hh-50Lh pairs (right) during the test phase with respect to the same probability in the baseline phase. **B**, Correlation between Δ reward and Δ probability of report. On the x-axis, the differences in the reward value associated with two targets (Δ reward) in a given pair are reported, corresponding to -60 , -30 , 0 , 30 , and 60 . For illustrative purposes only, data points with a Δ reward of -30 , 0 , and 30 were relatively displaced from their actual x-axis coordinate in order for SEM bars not to overlap. The y-axis instead represents

0.934; Fig. 3A). Together, these results demonstrate that our reward-based training protocol is able to induce considerable and durable changes in the attentional priority of specific spatial locations and that these plastic changes can occur with a high degree of spatial resolution.

We also applied the same approach to those trials in which two targets were presented at combinations of spatial locations for which a lesser imbalance in competitive strength might result following the delivery of rewards in the training phase, namely, 50Hh-20Lh and 80Hh-50Lh pairs, once again leading to a single report. As before, we computed the baseline-test differences (Δ probability) in reporting the target at the location that was expected to have gained a stronger competitive advantage after training. Although no significant change in the probability of single reports was found between the baseline and test phase when pairs of targets were displayed at the above spatial combinations, the numeric trend was in the expected direction (Δ probability = 0.016 in single reports for targets at the 50Hh location in 50Hh-20Lh pairs, $t_{(23)} = 0.575$, $p = 0.571$; Δ probability = 0.007 in single reports for targets at the 80Hh location in 80Hh-50Lh pairs, $t_{(23)} = 0.254$, $p = 0.802$). Figure 3B illustrates the correlation between the theoretical reward imbalance (Δ reward value; e.g., for the 80Hh-20Lh, the theoretical reward imbalance corresponds to 60) for targets displayed at each location within each of the four different combinations of spatial locations mentioned above (80Hh-20Lh, 50Hh-20Lh, 80Hh-50Lh, and 50Hh-50Lh) and the observed acquired competitive advantage (Δ probability; i.e., the baseline-test difference in the probability of reporting the target at the prioritized location in the pair, e.g., 80Hh in a 80Hh-20Lh pair; filled squares). For the sake of completeness, we also reported the theoretical reward imbalance and the acquired competitive “disadvantage” (Δ probability), calculated for the target displayed at the “weaker” location in each pair (e.g., 20Lh in a 80Hh-20Lh pair; empty squares). We underscore that values plotted in Figure 3B were obtained by using different subsets of the data to calculate acquired competitive advantage and disadvantage for each given pair, to avoid plotting the same values in the upper right and lower left quadrant, except for a change in sign. For example, we subdivided into two sets the data collected in double target trials in which the two targets appeared in a 80Hh and 20Lh spatial location, respectively (80Hh-20Lh pair), and leading to a single report (i.e., to the correct identification of one target only; operationally, trials were alternatively assigned to the first and second dataset until completion). We then used the first subset of data to calculate the change in the probability that the single reported target was the one displayed at the 80Lh location in 80Hh-20Lh pairs, and the second subset of data to calculate the probability that the single reported target was the one displayed at the 20Lh location. This procedure was repeated for each reward pair represented in the graph to provide a complete picture of reward-dependent changes in the attentional priority of specific locations. By applying a simple linear regression analysis, we

the difference in the probability of single reports of a given target in the pair during the test and the same probability during the baseline phase (Δ probability, mean $\pm$ SEM). Data points represented by filled squares correspond to Δ probabilities of report calculated for the theoretically “prioritized” spatial location in the pair, whereas data points represented by empty squares correspond to Δ probabilities of report calculated for the target location in the pair for which a disadvantage is expected. The red solid line indicates the fitted linear regression. **C**, Δ probability of a correct response in the single target condition. The difference in the probability of correctly reporting the target during the test versus baseline session (mean $\pm$ SEM) is shown as a function of reward level.

could establish a strong correlation ($R^2 = 0.921$; $F_{(1,6)} = 69.9494$; $p = 0.0016$; Fig. 3B, red line) between the theoretical reward imbalance (Δ reward value) and the induced change in the attentional priority of a given location in a pair (Δ probability). Finally, we tested whether the acquired change in attentional priority (competitive advantage or disadvantage) varied significantly across the different reward-associated spatial locations by performing a one-way ANOVA on Δ probability values for each reward combination, as reported in Figure 3B. The results confirmed a significant modulation of Δ probability of report across reward combinations ($F_{(7,161)} = 2.065$, $p = 0.05$, $\eta_p^2 = 0.082$).

In sum, we demonstrated that our reward-based training produced robust changes in the attentional priority of specific spatial locations, which manifested as an imbalanced probability to detect and correctly identify a target in one or the other of two critical spatial locations in a pair. Figure 4 represents a spatial priority map of the visual display used in the experiment where, for any given reward level, an average priority gain was reported (see figure legend for details), corresponding to the average competitive advantage (or disadvantage) acquired by the specific location associated with that reward level during the training phase. Because reward assignments to specific spatial locations during training varied from one subject to the other (see Materials and Methods), they were conventionally realigned here as corresponding to the example reward arrangement shown in Figure 1C. The spatial priority map nicely illustrates the competitive advantage (or disadvantage) acquired by spatial locations associated with specific reward levels as a result of the training procedure.

Finally, we analyzed performance in the single target condition to test whether changes in the attentional priority of specific locations following reward-based learning would also impact on accuracy of report in the presence of a single relevant stimulus among distractors. As assessed by a two-way ANOVA, including the main factors session (baseline vs test) and reward level (80Hh, 50Hh, 50Lh, and 20Lh), accuracy of report improved significantly in the post-training session compared with the baseline session ($F_{(1,23)} = 8.937$, $p = 0.007$, $\eta_p^2 = 0.280$; Fig. 3C), reflecting a strong practice effect. However, this improvement was uninfluenced by the specific reward manipulation assigned to spatial locations during the training phase ($F_{(3,69)} = 0.214$, $p = 0.886$; Fig. 3C). Moreover, the overall effect of reward level was also nonsignificant ($F_{(3,69)} = 0.525$, $p = 0.667$; Fig. 3C). A two-way ANOVA, including the main factors session and spatial location (1–8), was also performed. Here, both main effects were highly significant (session: $F_{(1,23)} = 8.932$, $p = 0.007$, $\eta_p^2 = 0.280$; spatial location: $F_{(7,161)} = 22.023$, $p < 0.001$, $\eta_p^2 = 0.489$). Conversely, performance improvement in the single target condition from

Figure 4. Plasticity of the spatial priority map. The figure illustrates the average priority gain computed for each reward-associated spatial location, both in a 2D plane (middle) and in a 3D representation (top). Bottom, The array of eight spatial locations used in the experiment (Fig. 1A,B), where associated rewards were again spatially arranged on the basis of the example of reward contingencies in Figure 1C. For a given reward level, the average priority gain was calculated by averaging the Δ probability (baseline test difference) of single report for targets at the spatial location associated with that reward level, as separately calculated for all imbalanced-reward pairs containing that reward level. For example, for the two 80Hh spatial locations, the priority gain, reported in the map with color coding (middle and top) and as z-axis value (top only), corresponds to the average of the Δ probabilities separately calculated in 80Hh-20Lh and 80Hh-50 pairs. In this case, we pooled together 80Hh-50Hh and 80Hh-50Lh pairs corresponding to an identical reward imbalance.

the baseline to the test session did not vary for different spatial locations ($F_{(7,161)} = 0.992$, $p = 0.439$). Importantly, lack of a reliable influence of our reward-based training protocol on accuracy of report in the single target condition provides a likely explanation for why we could not detect any reliable influence of reward imbalance during training, as in both cases performance was measured in relation to a singly presented target. Therefore, it appears that the influence of reward-based training on spatial priority maps can be detected much more easily under conditions of strong cross-target competition.

Discussion

The present study demonstrates that a controlled reward-based learning protocol can durably alter priority maps, deemed responsible for directing spatial attention and arbitrating selection under conditions where multiple potential targets compete for central resources (Itti and Koch, 2001; Thompson and Bichot, 2005; Fecteau and Munoz, 2006; Gottlieb, 2007; Bisley and Goldberg, 2010; Ptak, 2012). Specifically, during a training phase, the correct discrimination of a single target among distractors was associated with a predetermined probability of obtaining high versus low reward, which varied across locations within the display. At test, spatial loca-

tions associated with an overall more positive outcome during training (80Hh) were prioritized with respect to spatial locations associated with an overall less positive outcome (20Lh). The described change in the attentional priority of specific locations was evident only in the presence of cross-target competition (i.e., when two targets engaged in strong competition for attentional resources). Conversely, in the absence of cross-target competition (single target condition), there was no appreciable change in the performance across locations, likely because, in this condition, a sufficiently strong bottom-up signal conveyed by the single target in the array is enough to resolve competition against the irrelevant distractors.

Critically, the described effects are purely spatial in nature because they were measured with an experimental paradigm using stimulus material (and a task) that differed entirely from that used during training while sharing the same spatial configuration. This generalization allows us to conclude that the effects of our reward manipulation were specifically associated with locations in space and in no way reflected an association of rewards with the critical stimuli (or elemental features) to be identified, unlike what demonstrated in previous studies (e.g., Della Libera and Chelazzi, 2006, 2009; Raymond and O'Brien, 2009; Hickey et al., 2010a, b, 2011; Kristjánsson et al., 2010; Rutherford et al., 2010; Anderson et al., 2011a, b).

Reward-dependent alterations in target identification performance occurred with a high degree of spatial resolution, affecting spatial locations associated with a critical reward imbalance (80Hh and 20Lh), without spreading to nearby locations (50Hh or 50Lh) for which rewards were delivered in a balanced fashion. Hence, our results do not reflect a general facilitation in attentional orienting toward the visual hemifield more frequently associated with positive outcome during training, but rather an alteration in the attentional priority of discrete spatial locations based on the associated high (or low) probability of positive outcome. Importantly, the observed changes in performance for certain spatial locations were measured after a 4 d delay from the training phase, and in an extinction regimen, when rewards were no longer available, being therefore nonstrategic in nature (Chelazzi et al., 2013). Based on previous evidence collected in the nonspatial domain, it can be hypothesized that the reward effects described here are established through a reinforcement learning mechanism such that, when facing a complex array after training, attention is preferentially and more readily deployed to some locations relative to other locations (Chelazzi et al., 2013).

In sum, the present research provides clear-cut evidence that a reward-based training regimen is able to produce long-lasting alterations in the ability of observers to locate and identify task-relevant information at specific locations in space. The observed changes generalized across stimuli and task contexts, were long-lasting, affected performance with a high degree of spatial resolution, and were only evident in the context of cross-target competition. Together, these findings support the notion that our protocol was able to engender durable plastic changes in spatial priority maps.

Spatial priority maps are topographically organized maps of the external visual world, in which the behavioral priority of locations (or objects) is proportionally represented by differential neuronal activity (Thompson and Bichot, 2005; Fecteau and Munoz, 2006; Bisley and Goldberg, 2010). This concept originates from the notion of saliency map, a theoretical and computational construct representing space (and objects) in terms of their bottom-up salience (Koch and Ullman, 1985; Itti and Koch,

2000, 2001; Walther and Koch, 2006; Soltani and Koch, 2010), and extends to accommodate for task- and context-related top-down signals (e.g., Thompson et al., 2005; Serences and Yantis, 2007; Ipata et al., 2009). Spatial priority maps guide attention (and behavior) on the basis of a “winner-take all” principle, wherein attention is deployed to the location corresponding to peak activity in the map (Itti and Koch, 2001; Thompson and Bichot, 2005; Fecteau and Munoz, 2006; Bisley and Goldberg, 2010). Although spatial priority maps have been described as a real-time representation of behavioral salience resulting from convergent bottom-up and top-down signals, our study suggests that priority maps can be subject to durable plastic changes.

Based on electrophysiological evidence, cortical regions in the posterior parietal (Gottlieb et al., 1998; Kusunoki et al., 2000; Ipata et al., 2009; Bisley and Goldberg, 2010; e.g., LIP) and frontal cortex (Thompson and Bichot, 2005; Thompson et al., 2005; e.g., FEF) have been identified as candidate spatial priority maps (Serences and Yantis, 2007; Jerde et al., 2012; Ptak, 2012; Jerde and Curtis, 2013; Sprague and Serences, 2013; as also confirmed by neuroimaging studies in humans), and can be regarded as potential substrates of the learning effects described here, although the contribution of subcortical structures, including the superior colliculus (Noudoost et al., 2010; Bisley, 2011; Knudsen, 2011; Wurtz et al., 2011; Krauzlis et al., 2013), and the caudate nucleus in the basal ganglia (Yamamoto et al., 2012; Kim and Hikosaka, 2013) should not be excluded.

Interestingly, some of the same brain structures have been demonstrated to be highly sensitive to reward signals, although within the context of behavioral paradigms that are clearly different from the one used here, namely, paradigms investigating value-driven choice behavior (with reward acting as incentive) and not involving any learning component. Specifically, it has been documented that the prospect of differential reward in relation to alternative locations can affect neural activity in different nodes of the brain network controlling attentional deployment in the spatial domain and corresponding to putative sites of spatial priority maps, including regions of the frontal and prefrontal cortex (e.g., Leon and Shadlen, 1999; Glimcher, 2003; Roesch and Olson, 2007; Kim et al., 2012) and the posterior parietal cortex (e.g., Platt and Glimcher, 1999; Sugrue et al., 2004; Mohanty et al., 2008), as well as subcortical structures (Hikosaka, 2007; Yamamoto et al., 2013), notably the caudate nucleus of the basal ganglia. Neural activity in these regions has been shown to differentially encode the expected reward value associated with a given object or location, demonstrating the sensitivity of parietal and frontal regions, and of the caudate nucleus, to the motivational salience and/or valence of visual stimuli and locations (Leon and Shadlen, 1999; Platt and Glimcher, 1999; Glimcher, 2003; Sugrue et al., 2004; Hikosaka, 2007; Roesch and Olson, 2007; Mohanty et al., 2008; Serences, 2008; Kim et al., 2012; Leathers and Olson, 2012; Yamamoto et al., 2013).

As an alternative, one could hypothesize that reward biases were encoded in other regions of the nervous system deemed responsible for contextual spatial memory (including the encoding of reward-related information), such as, for example, the hippocampus and related structures (e.g., Luo et al., 2011; Lansink et al., 2012). In this view, the presence of a task-relevant stimulus (target) at a given location would trigger the activation of the reward-related memories associated with that specific location, in turn affecting the online encoding of the behavioral relevance of that specific location and conferring a competitive advantage to locations associated with more pos-

itive outcomes. Interestingly, a recent study demonstrated that reward associations can magnify the influence exerted by acquired spatial memories on current visual search processes within natural visual scenes, resulting in a reward-associated memory-based orienting of attention in space (Doallo et al., 2013). If our data were to be explained in this framework, however, it would be less obvious why reward-related representations did not influence performance in the single target condition.

To our knowledge, the present study is the first to describe long-term reward-based attentional learning effects for specific locations in space. In a few previous studies exploring the impact of motivational factors on space-based attention, rewards were used as an incentive to drive behavior and attention, thus addressing a completely different phenomenon from that described here. For example, attentional resources can be systematically oriented toward the spatial location(s) associated with the maximum expected reward in the current trial (Serences, 2008; Navalpakkam et al., 2009). In the same vein, an elegant study (Lucas et al., 2013) has recently demonstrated that visual exploration and attentional selection of visual targets can be biased in the short-term by an asymmetrical distribution of available rewards across space. Specifically, if available rewards were asymmetrically distributed from the left to the right of the target array in a “gambling” search task, both healthy participants and neglect patients developed a short-lasting leftward bias in target choice (Lucas et al., 2013), likely reflecting the enactment of specific cognitive strategies aimed at maximizing the earning of reward during the experiment. Importantly, the described effects developed during the course of a single experimental session in which rewards, acting as an incentive driving behavior and attention, were continuously available. Instead, in our study, reward was used as a feedback on performance and was demonstrated to act as a teaching signal to shape attention and behavior for future episodes of selection involving the same locations (Chelazzi et al., 2013).

More directly related to the present research, a recent study revealed a significant intertrial influence of valence information associated with specific spatial locations, dynamically affecting the online deployment of spatial attention (Camara et al., 2013). Although elements of the experimental design suggest that the pattern of results in this study might be interpreted as reflecting a form of priming of the motor response, it can also be taken to indicate that valence information affects attentional orienting in space, as claimed by the authors (Camara et al., 2013). At any rate, the described effects only marginally relate to our findings as they entail a short-term influence of reward, transitorily affecting performance without systematically altering the attentional priority of specific locations. Compatible results were obtained in other studies showing that value associated stimuli are capable of strong attentional capture (Rutherford et al., 2010; Anderson et al., 2011b). Participants learned to associate specific stimuli (or features) with more or less positive outcomes in a value learning training phase; afterward, when the same stimuli were used as uninformative cues in a typical probe task (Rutherford et al., 2010) or as irrelevant distractors in a visual search array (Anderson et al., 2011b), responses were slower for probes (or targets) appearing in the spatial location previously occupied by highly rewarded items, reflecting the typical cost in performance associated with inhibition of return (e.g., Klein, 2000). However, in these studies, modulations of spatial attention reflected an asso-

ciation of rewards with critical stimuli (or elemental features) and not with spatial locations.

What we address here is a different and more dramatic effect of reward on attentional processing in the spatial domain. Specifically, we provide the first evidence that reward is able to trigger a learning process wherein a specific bias in priority is acquired by spatial locations that were associated with more or less positive outcomes when selected in previous episodes of attentional deployment. Conceptually, the possibility that attentional representations of space undergo durable plastic changes is far from obvious because space is a sort of primordial medium for perception, attentional deployment, and action planning/execution, and as such, it could be somewhat resistant to learning effects based on the application of relatively short training regimens. Interestingly, a recent study demonstrated that space-based attentional guidance following an informative exogenous sensory cue is relatively impermeable to reward-based influences (Shomstein and Johnson, 2013), perhaps implying that basic, hard-wired principles governing attentional deployment in space are rather strong and inflexible. On the other hand, the pervasive functional significance of spatial encoding and attentional orienting in space for the enactment of goal-efficient behavior makes long-term learning of spatial priority of paramount importance for the attentional system to provide behavioral planning processes with the most efficient and informed-by-experience representation of the outer world (Gottlieb, 2012; Chelazzi et al., 2013). It is then highly plausible that plastic changes do occur at different levels in the brain corresponding to the representation of space in different spatial reference frames, or coordinate systems (e.g., Silver and Kastner, 2009; Ptak, 2012; Humphreys et al., 2013; Szczepanski and Saalmann, 2013), such that activation of contextual memories may help select the most relevant spatial priority map for the current goals of the individual.

References

- Alvarez GA, Cavanagh P (2005) Independent resources for attentional tracking in the left and right visual hemifields. *Psychol Sci* 16:637–643. [CrossRef Medline](#)
- Anderson BA (2013) A value-driven mechanism of attentional selection. *J Vis* 13:1–16. [CrossRef Medline](#)
- Anderson BA, Laurent PA, Yantis S (2011a) Value-driven attentional capture. *Proc Natl Acad Sci U S A* 108:10367–10371. [CrossRef Medline](#)
- Anderson BA, Laurent PA, Yantis S (2011b) Learned value magnifies salience-based attentional capture. *PLoS One* 6:e27926. [CrossRef Medline](#)
- Awh E, Belopolsky AV, Theeuwes J (2012) Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn Sci* 16:437–443. [CrossRef Medline](#)
- Beck DM, Kastner S (2009) Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vision Res* 49:1154–1165. [CrossRef Medline](#)
- Bisley JW (2011) The neural basis of visual attention. *J Physiol* 589:49–57. [CrossRef Medline](#)
- Bisley JW, Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1–21. [CrossRef Medline](#)
- Camara E, Manohar S, Husain M (2013) Past rewards capture spatial attention and action choices. *Exp Brain Res* 230:291–300. [CrossRef Medline](#)
- Carrasco M, Talgar CP, Cameron EL (2001) Characterizing visual performance fields: effects of transient covert attention, spatial frequency, eccentricity, task and set size. *Spat Vis* 15:61–75. [CrossRef Medline](#)
- Chakravarthi R, Cavanagh P (2009) Bilateral field advantage in visual crowding. *Vision Res* 49:1638–1646. [CrossRef Medline](#)
- Chelazzi L, Della Libera C, Sani I, Santandrea E (2011) Neural basis of visual selective attention. *Wiley Interdiscip Rev Cogn Sci* 2:392–407. [CrossRef Medline](#)
- Chelazzi L, Perlato A, Santandrea E, Della Libera C (2013) Rewards teach visual selective attention. *Vision Res* 85:58–72. [CrossRef Medline](#)

- Corbetta M, Shulman GL (2011) Spatial neglect and attention networks. *Annu Rev Neurosci* 34:569–599. [CrossRef Medline](#)
- Della Libera C, Chelazzi L (2006) Visual selective attention and the effects of monetary rewards. *Psychol Sci* 17:222–227. [CrossRef Medline](#)
- Della Libera C, Chelazzi L (2009) Learning to attend and to ignore is a matter of gains and losses. *Psychol Sci* 20:778–784. [CrossRef Medline](#)
- Doallo S, Patai EZ, Nobre AC (2013) Reward associations magnify memory-based biases on perception. *J Cogn Neurosci* 25:245–257. [CrossRef Medline](#)
- Fecteau JH, Munoz DP (2006) Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn Sci* 10:382–390. [CrossRef Medline](#)
- Glimcher PW (2003) The neurobiology of visual-saccadic decision making. *Annu Rev Neurosci* 26:133–179. [CrossRef Medline](#)
- Gottlieb J (2007) From thought to action: the parietal cortex as a bridge between perception, action, and cognition. *Neuron* 53:9–16. [CrossRef Medline](#)
- Gottlieb J (2012) Attention, learning, and the value of information. *Neuron* 76:281–295. [CrossRef Medline](#)
- Gottlieb JP, Kusunoki M, Goldberg ME (1998) The representation of visual saliency in monkey parietal cortex. *Nature* 391:481–484. [CrossRef Medline](#)
- Hickey C, Chelazzi L, Theeuwes J (2010a) Reward changes saliency in human vision via the anterior cingulate. *J Neurosci* 30:11096–11103. [CrossRef Medline](#)
- Hickey C, Chelazzi L, Theeuwes J (2010b) Reward guides vision when it's your thing: trait reward-seeking in reward-mediated visual priming. *PLoS One* 5:e14087. [CrossRef Medline](#)
- Hickey C, Chelazzi L, Theeuwes J (2011) Reward has a residual impact on target selection in visual search but not on the suppression of distractors. *Vis Cogn* 19:117–128. [CrossRef](#)
- Hikosaka O (2007) Basal ganglia mechanisms of reward-oriented eye movement. *Ann N Y Acad Sci* 1104:229–249. [CrossRef](#)
- Humphreys GW, Gillebert CR, Chechlacz M, Riddoch MJ (2013) Reference frames in visual selection. *Ann N Y Acad Sci* 1296:75–87. [CrossRef](#)
- Ipata AE, Gee AL, Bisley JW, Goldberg ME (2009) Neurons in the lateral intraparietal area create a priority map by the combination of disparate signals. *Exp Brain Res* 192:479–488. [CrossRef Medline](#)
- Itti L, Koch C (2000) A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Res* 40:1489–1506. [CrossRef Medline](#)
- Itti L, Koch C (2001) Computational modelling of visual attention. *Nat Rev Neurosci* 2:194–203. [CrossRef Medline](#)
- Jerde TA, Curtis CE (2013) Maps of space in human frontoparietal cortex. *J Physiol Paris* 107:510–516. [CrossRef Medline](#)
- Jerde TA, Merriam EP, Riggall AC, Hedges JH, Curtis CE (2012) Prioritized maps of space in human frontoparietal cortex. *J Neurosci* 32:17382–17390. [CrossRef Medline](#)
- Kim HF, Hikosaka O (2013) Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron* 79:1001–1010. [CrossRef Medline](#)
- Kim S, Cai X, Hwang J, Lee D (2012) Prefrontal and striatal activity related to values of objects and locations. *Front Neurosci* 6:108. [CrossRef Medline](#)
- Klein RM (2000) Inhibition of return. *Trends Cogn Sci* 4:138–147. [CrossRef Medline](#)
- Knudsen EI (2011) Control from below: the role of a midbrain network in spatial attention. *Eur J Neurosci* 33:1961–1972. [CrossRef Medline](#)
- Koch C, Ullman S (1985) Shifts in selective visual attention: towards the underlying neural circuitry. *Hum Neurobiol* 4:219–227. [Medline](#)
- Kraft A, Müller NG, Hagendorf H, Schira MM, Dick S, Fendrich RM, Brandt SA (2005) Interactions between task difficulty and hemispheric distribution of attended locations: implications for the splitting attention debate. *Brain Res Cogn Brain Res* 24:19–32. [CrossRef Medline](#)
- Krauzlis RJ, Lovejoy LP, Zénon A (2013) Superior colliculus and visual spatial attention. *Annu Rev Neurosci* 36:165–182. [CrossRef Medline](#)
- 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. *Atten Percept Psychophys* 72:1229–1236. [CrossRef Medline](#)
- Kusunoki M, Gottlieb J, Goldberg ME (2000) The lateral intraparietal area as a saliency map: the representation of abrupt onset, stimulus motion, and task relevance. *Vision Res* 40:1459–1468. [CrossRef Medline](#)
- Lansink CS, Jackson JC, Lankelma JV, Ito R, Robbins TW, Everitt BJ, Pennartz CM (2012) Reward cues in space: commonalities and differences in neural coding by hippocampal and ventral striatal ensembles. *J Neurosci* 32:12444–12459. [CrossRef Medline](#)
- Leathers ML, Olson CR (2012) In monkeys making value-based decisions, LIP neurons encode cue saliency and not action value. *Science* 338:132–135. [CrossRef Medline](#)
- Leon MI, Shadlen MN (1999) Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron* 24:415–425. [CrossRef Medline](#)
- Lucas N, Schwartz S, Leroy R, Pavin S, Diserens K, Vuilleumier P (2013) Gambling against neglect: unconscious spatial biases induced by reward reinforcement in healthy people and brain-damaged patients. *Cortex* 49:2616–2627. [CrossRef Medline](#)
- Luck SJ, Hillyard SA, Mangun GR, Gazzaniga MS (1989) Independent hemispheric attentional systems mediate visual search in split-brain patients. *Nature* 342:543–545. [CrossRef Medline](#)
- Luo AH, Tahsili-Fahadan P, Wise RA, Lupica CR, Aston-Jones G (2011) Linking context with reward: a functional circuit from hippocampal CA3 to ventral tegmental area. *Science* 333:353–357. [CrossRef Medline](#)
- Macaluso E (2010) Orienting of spatial attention and the interplay between the senses. *Cortex* 46:282–297. [CrossRef Medline](#)
- Macaluso E, Doricchi F (2013) Attention and predictions: control of spatial attention beyond the endogenous-exogenous dichotomy. *Front Hum Neurosci* 6:51–12. [CrossRef Medline](#)
- Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* 8:261–265. [CrossRef Medline](#)
- Mohanty A, Gitelman DR, Small DM, Mesulam MM (2008) The spatial attention network interacts with limbic and monoaminergic systems to modulate motivation-induced attention shifts. *Cereb Cortex* 18:2604–2613. [CrossRef Medline](#)
- Navalpakkam V, Koch C, Perona P (2009) Homo economicus in visual search. *J Vis* 31:1–16. [CrossRef Medline](#)
- Nobre AC (2001) The attentive homunculus: now you see it, now you don't. *Neurosci Biobehav Rev* 25:477–496. [CrossRef Medline](#)
- Noudoost B, Chang MH, Steinmetz NA, Moore T (2010) Top-down control of visual attention. *Curr Opin Neurobiol* 20:183–190. [CrossRef Medline](#)
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238. [CrossRef Medline](#)
- Ptak R (2012) The frontoparietal attention network of the human brain: action, saliency, and a priority map of the environment. *Neuroscientist* 18:502–515. [CrossRef Medline](#)
- Ptak R, Fellrath J (2013) Spatial neglect and the neural coding of attentional priority. *Neurosci Biobehav Rev* 37:705–722. [CrossRef Medline](#)
- Raymond JE, O'Brien JL (2009) Selective visual attention and motivation: the consequences of value learning in an attentional blink task. *Psychol Sci* 20:981–988. [CrossRef Medline](#)
- Roesch MR, Olson CR (2007) Neuronal activity related to anticipated reward in frontal cortex: does it represent value or reflect motivation? *Ann N Y Acad Sci* 1121:431–446. [CrossRef](#)
- Rutherford HJ, O'Brien JL, Raymond JE (2010) Value associations of irrelevant stimuli modify rapid visual orienting. *Psychon Bull Rev* 17:536–542. [CrossRef Medline](#)
- Serences JT (2008) Value-based modulations in human visual cortex. *Neuron* 60:1169–1181. [CrossRef Medline](#)
- Serences JT, Yantis S (2007) Spatially selective representations of voluntary and stimulus-driven attentional priority in human occipital, parietal, and frontal cortex. *Cereb Cortex* 17:284–293. [CrossRef Medline](#)
- Sereno AB, Kosslyn SM (1991) Discrimination within and between hemifields: a new constraint on theories of attention. *Neuropsychologia* 29:659–675. [CrossRef Medline](#)
- Shomstein S, Johnson J (2013) Shaping attention with reward: effects of reward on space- and object-based selection. *Psychol Sci* 24:2369–2378. [CrossRef Medline](#)
- Silver MA, Kastner S (2009) Topographic maps in human frontal and parietal cortex. *Trends Cogn Sci* 13:488–495. [CrossRef Medline](#)
- Soltani A, Koch C (2010) Visual saliency computations: mechanisms, constraints, and the effect of feedback. *J Neurosci* 30:12831–12843. [CrossRef Medline](#)
- Sprague TC, Serences JT (2013) Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nat Neurosci* 16:1879–1887. [CrossRef Medline](#)

- Squire RF, Noudoost B, Schafer RJ, Moore T (2013) Prefrontal contributions to visual selective attention. *Annu Rev Neurosci* 36:451–466. [CrossRef Medline](#)
- Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304:1782–1787. [CrossRef Medline](#)
- Szczepanski SM, Saalman YB (2013) Human fronto-parietal and parieto-hippocampal pathways represent behavioral priorities in multiple spatial reference frames. *Bioarchitecture* 5:147–152. [CrossRef Medline](#)
- Thompson KG, Bichot NP (2005) A visual salience map in the primate frontal eye field. *Prog Brain Res* 147:251–262. [CrossRef Medline](#)
- Thompson KG, Bichot NP, Sato TR (2005) Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. *J Neurophysiol* 93:337–351. [CrossRef Medline](#)
- Walther D, Koch C (2006) Modeling attention to salient proto-objects. *Neural Netw* 19:1395–1407. [CrossRef Medline](#)
- Wurtz RH, McAlonan K, Cavanaugh J, Berman RA (2011) Thalamic pathways for active vision. *Trends Cogn Sci* 15:177–184. [CrossRef Medline](#)
- Yamamoto S, Monosov IE, Yasuda M, Hikosaka O (2012) What and where information in the caudate tail guides saccades to visual objects. *J Neurosci* 32:11005–11016. [CrossRef Medline](#)
- Yamamoto S, Kim HF, Hikosaka O (2013) Reward value-contingent changes of visual responses in the primate caudate tail associated with a visuomotor skill. *J Neurosci* 33:11227–11238. [CrossRef Medline](#)
- Yantis S, Serences JT (2003) Cortical mechanisms of space-based and object-based attentional control. *Curr Opin Neurobiol* 13:187–193. [CrossRef Medline](#)