

Research Article

Learning to Attend and to Ignore Is a Matter of Gains and Losses

Chiara Della Libera and Leonardo Chelazzi

Department of Neurological and Vision Sciences, Section of Human Physiology, University of Verona, and National Institute of Neuroscience, Verona, Italy

ABSTRACT—*Efficient goal-directed behavior in a crowded world is crucially mediated by visual selective attention (VSA), which regulates deployment of cognitive resources toward selected, behaviorally relevant visual objects. Acting as a filter on perceptual representations, VSA allows preferential processing of relevant objects and concurrently inhibits traces of irrelevant items, thus preventing harmful distraction. Recent evidence showed that monetary rewards for performance on VSA tasks strongly affect immediately subsequent deployment of attention; a typical aftereffect of VSA (negative priming) was found only following highly rewarded selections. Here we report a much more striking demonstration that the controlled delivery of monetary rewards also affects attentional processing several days later. Thus, the propensity to select or to ignore specific visual objects appears to be strongly biased by the more or less rewarding consequences of past attentional encounters with the same objects.*

Efficient goal-directed behavior in a crowded world is crucially mediated by visual selective attention (VSA), which regulates deployment of cognitive resources toward selected, behaviorally relevant visual objects. Acting as a filter on perceptual representations, VSA allows preferential processing of relevant objects and concurrently inhibits traces of irrelevant items, thus preventing harmful distraction (Posner, 2004; Serences & Yantis, 2006). Recently, we demonstrated that the ongoing deployment of VSA may be strongly affected by the delivery of monetary rewards (Della Libera & Chelazzi, 2006). In that study, observers performed a task in which they had to respond to prime and probe displays, presented as sequential pairs within individual trials. After each correct response to a prime stimulus, observers were given a high or low monetary reward. The

level of the reward did not depend on actual performance, but subjects were misleadingly told that high and low rewards signified optimal and suboptimal performance, respectively. Under these conditions, *negative priming* (impaired response to a probe target that had served as the distractor in the preceding prime display; Tipper, 2001) occurred only if the attentional selection of the prime target had been highly rewarded and was therefore deemed successful by the subject (Della Libera & Chelazzi, 2006). This study revealed for the first time that attentional processes are subject to an “efficiency check” system that dynamically adjusts attentional deployment toward specific items on the basis of previous outcomes. Every time a selection occurs, a memory trace is stored, and the strength of that trace, in turn, is modulated by how successful the selection turns out to be. A highly rewarded attentional selection will leave a stronger—and longer-lasting—trace than a selection that has poor consequences.

Our demonstration that rewards can adjust the immediate deployment of VSA raises the possibility that they may also be capable of shaping attentional deployment toward specific objects in the long term. For instance, one might expect that objects whose selection has been consistently followed by higher rewards would acquire a privileged status and become more salient than objects whose selection has been followed by lower rewards, so as to facilitate performance. This prediction goes far beyond what we showed previously (Della Libera & Chelazzi, 2006), and if confirmed would force a radical shift in perspective concerning a whole class of observations in attention research. Specifically, it would suggest that long-term learning effects in attentional tasks (including visual search; e.g., Chun, 2000) do not simply reflect development of expertise with given stimulus materials, but are instead critically dependent on the more or less rewarding outcomes of past attentional encounters with those materials.

To test this hypothesis, we developed a new experimental paradigm made up of a *training phase*, during which correct attentional selection of specific visual items was rewarded with differential monetary gains, and a delayed *test phase*, run several days later, in which the effects of the history of rewards on at-

Address correspondence to Leonardo Chelazzi, Department of Neurological and Vision Sciences, Section of Physiology, University of Verona, Strada Le Grazie 8, I-37134 Verona, Italy, e-mail: leonardo.chelazzi@univr.it.

tentional selection could be assessed in the absence of any ongoing reward manipulation. This approach was applied in two related experiments.

GENERAL METHOD

The two experiments were identical with respect to the training procedure, which is outlined in this section. The experiments differed in the testing paradigm used, as explained in the sections that follow.

Participants, Stimuli, and Apparatus

Participants were students at Verona University. All had normal or corrected-to-normal vision. A different group of 16 subjects participated in each experiment (Experiment 1: 11 females and 5 males with a mean age of 24 years; Experiment 2: 11 females and 5 males with a mean age of 23 years). None of them had previously participated in similar experiments, and they were all naive as to the purposes of the study. When queried after the end of the experiments, they did not express any intuition or suspicion regarding the actual schedule of reward delivery during the training phase.

The stimuli were a set of 16 outline nonsense shapes ($2^\circ \times 2^\circ$), the same ones used in our earlier research (Della Libera & Chelazzi, 2006), and were selected from the set used by Strayer and Grison (1999). Stimulus displays were presented on a 15-in. CRT monitor in a quiet and dimly lit room, at a viewing distance of 57 cm. The experiments were created and run on a personal computer with E-Prime software (Schneider, Eschman, & Zuccolotto, 2002).

Training Task and Reward Schedules

Throughout the training phase, subjects performed a task in which stimuli were to be selected or ignored. As in our previous work (Della Libera & Chelazzi, 2006), correct selections resulted in high or low monetary feedback, intended to imply optimal or suboptimal performance, respectively, to the subjects. Each trial started with a cue, a green or red square ($0.5^\circ \times 0.5^\circ$) displayed centrally for 400 ms. The color of this cue signaled which shape in the following display was the target for that trial. The following stimulus display consisted of three shapes. Two overlapping shapes, one red and one green, appeared at 3° of eccentricity on the left, along the horizontal meridian; the shape with the same color as the cue was the target, and the other shape was the distractor. A single black shape, the comparison, appeared at the same eccentricity on the right. Subjects were to make a same/different judgment on the shapes of the target and the comparison, and then indicate their decision by pressing one of two response keys with the index or middle finger of the right hand. The stimulus display was presented 600 ms after cue offset and remained visible for 3 s or until the subject responded, whichever came first. Correct responses were followed by a re-

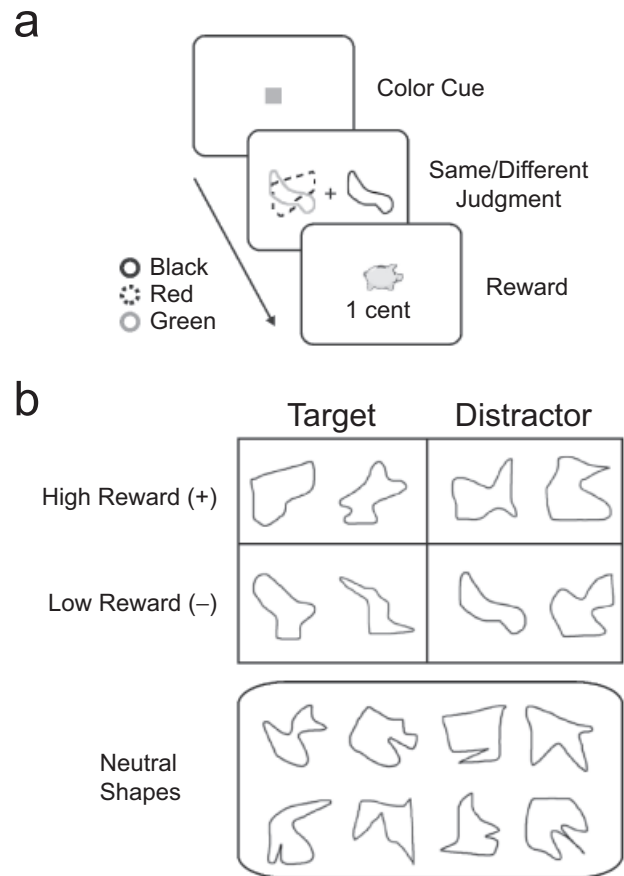


Fig. 1. Illustration of the training paradigm: examples of (a) the display sequence and (b) the assignment of stimulus shapes to reward categories. In each trial, two overlapping shapes, one red and one green, appeared on the left of fixation, while a black shape appeared on the opposite side. A previously presented green or red central square cued the shape that was relevant in the trial. Subjects performed a same/different judgment between the shape displayed in the relevant color and the black shape. Correct responses were followed by a reward, which could be high or low, and the amount gained was indicated on the monitor. The assignment of the stimuli to reward levels (high, low) and roles (target, distractor, neutral) was different for each subject, in order to avoid possible confounds. See the main text for details.

ward, which could be *high* (€0.10) or *low* (€0.01); the amount was shown on the monitor for 600 ms (see Fig. 1a). A new trial started after a 1-s intertrial interval. Errors were followed by an 800-ms auditory tone.

Subjects completed three training sessions on consecutive days, followed by a delayed test session. Each training session comprised 960 trials and lasted approximately 1 hr. Stimulus displays were designed so that each of the 16 nonsense shapes appeared equally often as the cued target, the distractor, and the comparison shape. The target and comparison shapes were identical in 50% of the trials; these trials required a “same” response. In the remaining trials, the comparison shape differed from both the target and the distractor, and the correct response was “different.”

A fundamental aspect of the procedure is that we dissociated reward levels from task performance. Although subjects were

told that rewards depended on their performance, the size of the rewards was completely decoupled from actual response parameters and was balanced across all experimental conditions so that high and low rewards occurred with the same overall probability (50%). However, in order to effectively deceive participants, we varied the total monetary compensation across individuals (€85–100).

Crucially, we manipulated the schedule of reward delivery so that the stimulus shapes could be divided into categories according to the proportion of high versus low rewards received when the shapes served as targets or distractors (see Fig. 1b). Four shapes had a biased probability of leading to a high or low reward when they were targets. Two of these items led to a high reward in 80% of cases and a low reward in 20% of cases (*T+* shapes); the other 2 led to a high reward in 20% of cases and a low reward in 80% of cases (*T-* shapes). Four other shapes had a biased probability of leading to a high or low reward when they were distractors. Two of them led to a high reward in 80% of cases and a low reward in 20% of cases (*D+* shapes); the other 2 led to a high reward in 20% of cases and a low reward in 80% of cases (*D-* shapes). When these 8 shapes were displayed in the role that was not associated with a bias in the reward schedule (e.g., when a *T+* item was presented as a distractor), they led to high and low rewards with equal probability. The remaining 8 shapes were used as neutral fillers and led to a high reward in 50% of cases and a low reward in 50% of cases, both when they were targets and when they were distractors. In order to avoid any possible confounds caused by fixed assignment of individual shapes to the experimental categories, we created 16 possible shape-category combinations, so that the 16 shapes were categorized differently for each participant.

EXPERIMENT 1

Five days after training, participants in Experiment 1 were tested using a task identical to the one performed during training, except for the absence of rewards (Fig. 2a). Performance was analyzed in terms of reaction time (RT) for correct responses and error rates.

First, we evaluated whether subjects' speed in selecting target shapes was affected by the shapes' reward contingencies during training. That is, we tested whether RTs in the test trials differed according to the categories the target shapes had belonged to in the training phase. We grouped RTs according to the type of shape presented as the target, excluding trials in which the target was a neutral shape, and submitted the data to a two-way analysis of variance (ANOVA) with factors of *reward bias* (80% high/20% low or 20% high/80% low) and *item history* (whether the item was a target or a distractor when bias was applied during training). The analysis revealed no significant effects, indicating that task performance was unaffected by targets' reward history (Fig. 3a, left panel).

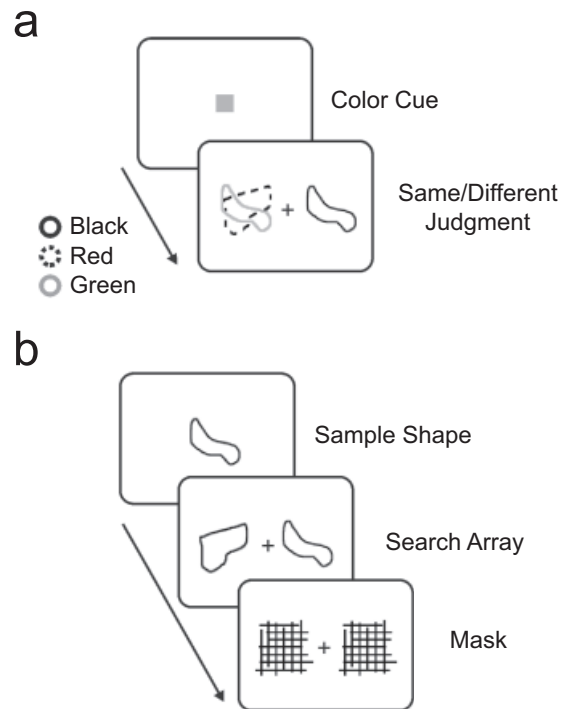


Fig. 2. Illustration of the display sequence in the test trials of (a) Experiment 1 and (b) Experiment 2. In Experiment 1 (a), as in the training task, each trial started with the presentation of a central cue that signaled the color of the target shape in the upcoming stimulus display. Two overlapping shapes, one red and one green, were then shown on the left of the screen, while a black shape, the comparison shape, appeared on the right. Subjects performed a same/different judgment between the target and the comparison shape. Unlike in training, correct responses were not followed by reward. In Experiment 2 (b), each trial started with the presentation of a target shape that participants were to search for in the upcoming search array. The search array was then shown briefly, before being replaced by two masking patterns. Subjects had to report as quickly and as accurately as possible whether the target shape was present or absent in the search array.

A second ANOVA was conducted after grouping data according to the type of shape presented as the distractor (i.e., *T+*, *T-*, *D+*, or *D-*). This ANOVA, again with reward bias and item history as the main factors, revealed no significant main effects, but a highly significant interaction, $F(1, 15) = 17.868, p < .001, \eta_p^2 = .544$ (Fig. 3a, right panel). Post hoc *t* tests revealed that responses to trials with *T+* or *D-* items as distractors were significantly slower than responses to trials with *T-* or *D+* items as distractors—*T+* vs. *T-*: $t(15) = 3.715, p < .002, p_{\text{rep}} = .979$; *T+* vs. *D+*: $t(15) = 3.738, p < .002, p_{\text{rep}} = .979$; *D-* vs. *T-*: $t(15) = 2.169, p < .05, p_{\text{rep}} = .878$; *D-* vs. *D+*: $t(15) = 2.281, p < .04, p_{\text{rep}} = .892$.

The average error rate was 5.3% across subjects and conditions. Error rates were analyzed following the same approach as for RTs, but none of the crucial effects in the ANOVAs were significant.

These results show that our manipulation was successful at establishing long-term attentional biases in relation to specific items. Specifically, the inhibitory mechanisms involved in fil-

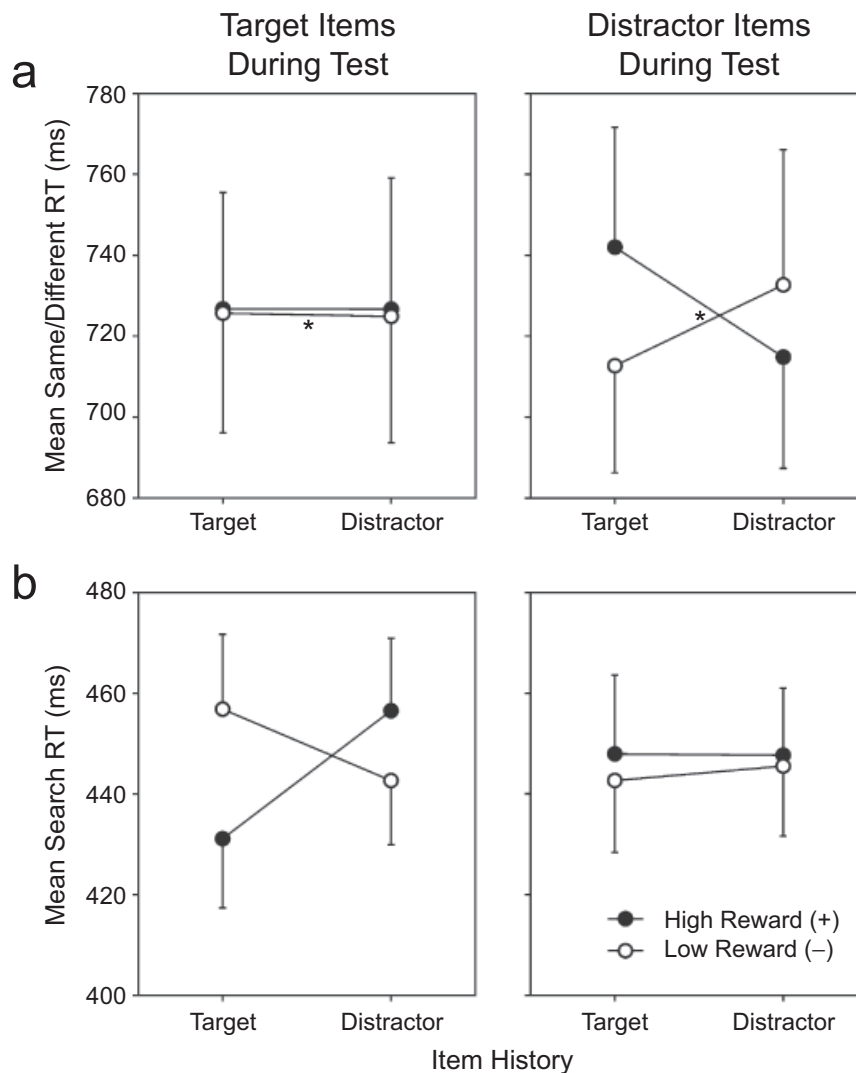


Fig. 3. Mean reaction time (RT) on correct trials as a function of reward bias (high: 80% high/20% low; low: 20% high/80% low) and item history (whether the item was a target or a distractor when bias was applied in training). Results are shown separately for target and distractor items in test trials of (a) Experiment 1 and (b) Experiment 2. Error bars represent standard errors of the means. Asterisks indicate average performance in responding to neutral items.

tering out distractors were especially sensitive to the reward history associated with each individual shape. Items that were particularly advantageous targets during training because they led to higher rewards when correctly selected (T+ shapes) became highly interfering when later displayed as distractors. Comparable, but opposite, results were observed for shapes previously associated with a biased reward schedule when presented as distractors. Items that, when filtered out during training, were associated with high rewards on most occasions (D+ shapes) became less interfering distractors, as if they had become less salient and therefore more easily disregarded. The opposite effects observed for T+ and D+ items are particularly important because they show that the observed influence of reward on performance depended critically on the combination of the amount of reward given when a particular shape was presented and whether the shape was selected or ignored.

Symmetrical results were found for shapes that during training were more often associated with low than high rewards. Items that were scarcely advantageous targets (T- shapes) became relatively weak distractors and could be readily ignored; conversely, items that had been filtered out but resulted in poor outcomes (D- shapes) retained relatively high interfering power. These striking results demonstrate that VSA mechanisms keep track of the outcomes of past attentional episodes with specific objects, so that past outcomes affect attentional performance several days later. The next experiment assessed the degree of generality of these effects.

EXPERIMENT 2

In Experiment 2, we tested whether the differences in attentional saliency of the various shape categories, created by

training, were tied to the specific selective processes involved in the training task, or whether these differences could generalize across diverse attentional paradigms. Moreover, we explored whether the influence of reward contingencies is linked exclusively to the processing of distractors, as found in Experiment 1, or whether the influence could be revealed also for the processing of targets, given the appropriate circumstances. The general procedure used in Experiment 1 was replicated in Experiment 2, except that a simple visual search task was used in the test phase. Key features of this task were that, unlike in Experiment 1, the location of the target in the display (on target-present trials) was unpredictable from trial to trial and the distractor did not overlap with the target.

Task and Procedure

The stimuli, apparatus, and training procedure were the same as in Experiment 1, except that a simple visual search task was introduced in the test phase. In each trial of this task, a single black outline was presented centrally for 400 ms, designating the target shape for the trial. After a 300-ms interval, two black shapes appeared, one on each side of fixation; these shapes were centered at 3° of eccentricity along the horizontal meridian. This array was visible for 180 ms and then immediately replaced by masks (see Fig. 2b). Subjects had to report as quickly and accurately as possible whether the target shape was present or absent, indicating their answer by pressing one of two keys on the numerical keypad. They were allowed 1,500 ms after the onset of the array to make a response. A new trial started 800 ms after the response. Errors were followed by an 800-ms auditory tone. The test session lasted approximately 1 hr and comprised 960 trials, evenly divided into target-present and target-absent trials. Each search array included one shape from the neutral set and one shape belonging to one of the four critical categories (T+, T-, D+, D-). The latter, the critical item, was equally likely to be the target or the distractor. Moreover, each stimulus was used the same number of times as the target and as the distractor.

Results and Discussion

Analyses were performed on mean RTs for correct responses and on error rates. Separate analyses were performed for target-present and target-absent trials.¹ Our primary interest was the pattern of performance in the target-present condition, which allowed us to assess the effects of our manipulation in relation to both targets and distractors. In this condition (as in the target-absent condition), one critical item was always included in the array, so the factors in the ANOVA were reward bias of the critical item (80% high/20% low or 20% high/80% low), item

history (whether the critical item was a target or a distractor when bias was applied in training), and item status in the search array (whether the critical item was the target or the distractor). Most relevant to our purposes, the three-way interaction involving all factors in the ANOVA was highly significant, $F(1, 15) = 20.5, p < .0001, \eta_p^2 = .577$.² Systematic analyses revealed that a significant interaction between reward bias and item history was found only when the critical item in the array was the target (Fig. 3b, left panel). Searches for T+ shapes were significantly faster than those for either T- or D+ shapes—T+ vs. T-: $t(15) = 2.517, p < .03, p_{rep} = .908$; T+ vs. D+: $t(15) = 3.515, p < .003, p_{rep} = .974$. Moreover, searching for a D- shape resulted in faster RTs than searching for a T- item, $t(15) = 3.188, p < .006, p_{rep} = .962$, or for a D+ item, although the latter effect did not reach statistical significance, $t(15) = 1.542, p = .144$. In contrast, analyses showed that the four stimulus categories (T+, T-, D+, and D-) yielded statistically indistinguishable performance when the critical item in the array was the distractor (Fig. 3b, right panel).

Less revealing results emerged from analyses on target-absent trials, which is not surprising given the finding that in the target-present condition, our manipulation affected performance only in relation to the target. The factors in this ANOVA were item history of the target of search, or sample item (whether the item was a target or a distractor in the display when bias was applied during training); reward bias associated with the target of search, or sample item (80% high/20% low or 20% high/80% low); item history of the critical distractor (whether the item was a target or a distractor when bias was applied in training); and reward bias associated with the critical distractor (80% high/20% low or 20% high/80% low). We excluded trials in which the sought-for target was a neutral shape. The main effect of item history of the target of search was highly significant, $F(1, 15) = 35.313, p < .0001, p_{rep} = .995, \eta_p^2 = .702$; there was a reliable advantage of 14 ms in search for items that were special targets during training (T+ or T-; 493 ms), relative to search for items that were special distractors during training (D+ or D-; 507 ms). A marginally significant interaction was found between reward bias associated with the target of search and reward bias associated with the distractor, $F(1, 15) = 5.442, p < .04, p_{rep} = 0.901, \eta_p^2 = .266$.

The average error rate across subjects and conditions was 7.1%. An ANOVA on error rates in the target-present condition ($M = 8.4\%$) did not reveal any significant effects.³

Crucial results from this experiment were obtained by analyzing performance on target-present trials. Responses were significantly faster for targets whose correct selection had led to higher monetary gains during training (T+ shapes) than for

¹An omnibus ANOVA including both target-present and target-absent trials could not be performed because different factors were available in the two cases.

²The interaction between reward bias and item history was also significant, $F(1, 15) = 8.437, p < .02, p_{rep} = .947, \eta_p^2 = .360$. Average RT was 440 ms for T+ items, 450 ms for T- items, 452 ms for D+ items, and 444 ms for D- items.

³The ANOVA on error rates revealed two marginally significant interactions in the target-absent condition, but they were of no obvious meaning.

targets whose correct selection had led to less rewarding outcomes during training (T− shapes). We speculate that the increased saliency of T+ shapes that was acquired during training rendered them more promptly accessible to attentional processing even when they were encountered in a new experimental context. Conversely, it was easier to select a target that during training had been a less rewarded distractor (D− shapes) than to select a target that during training had been a highly rewarded distractor (D+ shapes), although this was only a trend in the data. These findings suggest—in agreement with the results of Experiment 1—that memory traces for each selection episode were formed with reference to both the selected target and the ignored distractor, and that these traces were independently modulated by our reward manipulation.

GENERAL DISCUSSION

These results demonstrate for the first time that the attentional processing of specific objects is adjusted durably according to the more or less rewarding consequences of prior attentional episodes concerning the same objects. Therefore, the long-term learning to select and to ignore specific objects in the environment is shaped by a cumulative measure of gains (and losses) resulting from past encounters with those objects.

In brief, items whose correct selection as targets had typically been followed by the high reward during training (T+ items) became more difficult to reject when serving as distractors (Experiment 1) and easier to select when serving as targets (Experiment 2). In contrast, items whose correct selection as targets had typically been followed by the low reward during training (T− items) became easier to reject when serving as distractors (Experiment 1) and more difficult to select when serving as targets (Experiment 2). Similarly, if the rejection of an item serving as the distractor was often followed by relatively favorable outcomes during training (D+ items), the item became easier to ignore when later presented as a distractor (Experiment 1) and harder to select when later presented as the target (Experiment 2). Conversely, stimuli that had typically led to poor outcomes during training when they served as distractors (D− items) retained considerable interfering power when shown as distractors (Experiment 1), but became easier to select when shown as targets (Experiment 2). On the basis of these findings, we propose that every episode of attentional selection leaves behind a memory trace that incorporates information about the specific items involved, the specific attentional processes applied to them, and, crucially, the adaptive value associated with the episode.

Item-specific biases in reward delivery produced different effects in Experiments 1 and 2, affecting either the mechanisms responsible for selecting targets (Experiment 2) or those involved in rejecting distractors (Experiment 1). This divergence can be explained by the nature of the tasks used. The test task in Experiment 1 allowed more efficient control of target selection.

Because the target's color was signaled in advance and its position was constant, selection was relatively immune to our manipulation. In addition, the overlap of target and distractor items might have posed especially high demands on inhibitory attentional mechanisms (e.g., Lavie & Tsal, 1994), making suppression of nonrelevant distractors more susceptible to our manipulation.

Experiment 2 instead used a simple visual search task during the test phase, and this task presumably tapped the mechanisms of target selection more directly, placing less emphasis on distractor suppression. In this case, the sample shape had to be memorized anew on each trial, and the target's location (on target-present trials) varied randomly across trials. Moreover, given that the target and distractor occupied well-separated positions in space, the role of inhibitory processes might have been less decisive in sustaining performance. Still, under appropriate circumstances, reward-based effects for the target and the distractor might occur together.

It is well established that the ability to select and to ignore specific items in cluttered displays undergoes substantial improvement after extensive practice (learning; e.g., Chun, 2000; Kyllingsbæk, Schneider, & Bundesen, 2001; Shiffrin & Schneider, 1977; Tipper, Grison, & Kessler, 2003; Vidnyànszky & Sohn, 2005). Such learning is also reflected in how gaze is preferentially directed toward certain objects, and not others, during visual exploration of familiar scenes (e.g., Bichot & Schall, 1999; Hayhoe & Ballard, 2005; Jagadeesh, Chelazzi, Mishkin, & Desimone, 2001; Rothkopf, Ballard, & Hayhoe, 2006). The experiments reported here offer novel evidence concerning the way in which learning affects attentional processes. Specifically, attention seems to be influenced not only by past encounters with specific objects and contexts, but also by the previous consequences of selecting or discarding specific objects. This possibility fits well with an *instance-based* view of behavioral development and control (Logan, 1988, 2002). According to this theory, learning is mediated by the accumulation of instances, each of which adds new information that can affect future performance in similar contexts and situations. Extending this idea, the present work indicates that “instances” are created during attentional selections, and that such instances include both the representation of the objects involved in a selection episode and information about the outcomes of that episode. This combination of information comes into play automatically when the same objects compete for attentional processing on future occasions.

Our results can also be readily interpreted in relation to current reinforcement learning theories (Dayan & Balleine, 2002; Montague, Hyman, & Cohen, 2004; Sutton & Barto, 1998). According to these theories, rewards (and punishments) instantiate plastic changes in neural circuits responsible for decision making and behavioral control, and these changes serve the goal of maximizing future gains and minimizing future losses. Similarly, we propose that in our paradigm, the controlled delivery of

rewards adaptively adjusts attentional biases toward specific items, with the primary consequence of leading to enhanced selection of targets that have been associated with higher rewards when presented in the same role during training, and to enhanced rejection of distractors that have been associated with higher rewards when presented in the same role during training.

Functional neuroimaging studies in humans and single-cell recordings in behaving macaques have recently documented reward-related signals in numerous brain structures, including the basal ganglia, orbitofrontal and anterior cingulate cortices, amygdala, and dopaminergic midbrain (O'Doherty, 2004, 2007; Schultz, 2000, 2006). Moreover, it has been shown that reward-related modulation in additional brain areas, including several sectors of frontal and prefrontal cortex, posterior parietal cortex, the basal ganglia, and the superior colliculus, has a direct impact on sensorimotor processing and behavioral control (Hikosaka, Nakamura, & Nakahara, 2006). Many of the modulated regions (especially posterior parietal and prefrontal cortices and the superior colliculus) are key players in the brain network controlling attention and gaze (e.g., Glimcher, 2001; Ikeda & Hikosaka, 2003; Leon & Shadlen, 1999; Platt & Glimcher, 1999; Roesch & Olson, 2004; Sugrue, Corrado, & Newsome, 2004). Given these findings, it is perhaps not surprising that our reward-based manipulation was so successful at shaping long-term attentional deployment toward specific stimuli.

Acknowledgments—We acknowledge the generous support from Fondazione Cariverona, Verona, Italy.

REFERENCES

- Bichot, N.P., & Schall, J.D. (1999). Effects of similarity and history on neural mechanisms of visual selection. *Nature Neuroscience*, 2, 549–554.
- Chun, M.M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, 4, 170–177.
- Dayan, P., & Balleine, B.W. (2002). Reward, motivation, and reinforcement learning. *Neuron*, 36, 285–298.
- Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science*, 17, 222–227.
- Glimcher, P.W. (2001). Making choices: The neurophysiology of visual-saccadic decision making. *Trends in Neurosciences*, 24, 654–659.
- Hayhoe, M.M., & Ballard, D.H. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, 9, 188–194.
- Hikosaka, O., Nakamura, K., & Nakahara, H. (2006). Basal ganglia orient eyes to reward. *Journal of Neurophysiology*, 95, 567–584.
- Ikeda, T., & Hikosaka, O. (2003). Reward-dependent gain and bias of visual responses in primate superior colliculus. *Neuron*, 39, 693–700.
- Jagadeesh, B., Chelazzi, L., Mishkin, M., & Desimone, R. (2001). Learning increases stimulus salience in anterior inferior temporal cortex of the macaque. *Journal of Neurophysiology*, 86, 290–303.
- Kyllingsbæk, S., Schneider, W.X., & Bundesen, C. (2001). Automatic attraction of attention to former targets in visual displays of letters. *Perception & Psychophysics*, 63, 85–98.
- Lavie, N., & Tsai, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. *Perception & Psychophysics*, 56, 183–197.
- Leon, M.I., & Shadlen, M.N. (1999). Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron*, 24, 415–425.
- Logan, G.D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492–527.
- Logan, G.D. (2002). An instance theory of attention and memory. *Psychological Review*, 109, 376–400.
- Montague, P.R., Hyman, S.E., & Cohen, J.D. (2004). Computational roles for dopamine in behavioural control. *Nature*, 431, 760–767.
- O'Doherty, J.P. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology*, 14, 769–776.
- O'Doherty, J.P. (2007). Lights, camembert, action! The role of human orbitofrontal cortex in encoding stimuli, rewards, and choices. In G. Shoenbaum, J.A. Gottfried, E.A. Murray, & S.J. Ramus (Eds.), *Linking affect to action: Critical contribution of the orbitofrontal cortex* (Annals of the New York Academy of Sciences Vol. 1121, pp. 254–272). New York: New York Academy of Sciences.
- Platt, M.L., & Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233–238.
- Posner, M.I. (2004). *Cognitive neuroscience of attention*. New York: Guilford Press.
- Roesch, M.R., & Olson, C.R. (2004). Neuronal activity related to reward value and motivation in primate frontal cortex. *Science*, 304, 307–310.
- Rothkopf, C.A., Ballard, D.H., & Hayhoe, M.M. (2007). Task and context determine where you look. *Journal of Vision*, 7, 1–20.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2002). *E-Prime user's guide*. Pittsburgh, PA: Psychology Software Tools.
- Schultz, W. (2000). Multiple reward signals in the brain. *Nature Reviews Neuroscience*, 1, 199–207.
- Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annual Review of Psychology*, 57, 87–115.
- Serences, J.T., & Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*, 10, 38–45.
- Shiffrin, R.M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127–190.
- Strayer, D.L., & Grison, S. (1999). Negative identity priming is contingent on stimulus repetition. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 24–38.
- Sugrue, L.P., Corrado, G.S., & Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, 304, 1782–1787.
- Sutton, R., & Barto, A.G. (1998). *Reinforcement learning*. Cambridge, MA: MIT Press.
- Tipper, S.P. (2001). Does negative priming reflect inhibitory mechanisms? A review and integration of conflicting views. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, 54, 321–343.
- Tipper, S.P., Grison, S., & Kessler, K. (2003). Long-term inhibition of return of attention. *Psychological Science*, 14, 19–25.
- Vidnyànszky, Z., & Sohn, W. (2005). Learning to suppress task-irrelevant visual stimuli with attention. *Vision Research*, 45, 677–685.

(RECEIVED 7/11/08; REVISION ACCEPTED 11/14/08)