Object-based attention is accentuated by object reward association

Damiano GrignolioIDhttps://orcid.org/0000-0001-7699-2390David J. AcunzoIDhttps://orcid.org/0000-0003-3282-1551Clayton HickeyIDhttps://orcid.org/0000-0002-4162-5643

Centre for Human Brain Health and School of Psychology, University of Birmingham, Birmingham, UK

Word count: ~11000

Number of figures: 5

Author note:

We have no known conflict of interest to disclose. All material and data are available at [UBIRA eData repository - University of Birmingham institutional repository, link to be inserted at publication].

Correspondence concerning the article should be addressed to Damiano Grignolio or Clayton Hickey, Centre for Human Brain Health, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK. Email: dxg063@student.bham.ac.uk, <u>c.m.hickey@bham.ac.uk</u>

This work was supported by the European Research Council (ERC) under the European Union Horizon 2020 Research and Innovation Program Grant Agreement 804360 to C.H.

Abstract

Humans use selective attention to prioritize visual features, like color or shape, as well as discrete spatial locations, and these effects are sensitive to the experience of reward. Reward-associated features and locations are accordingly prioritized from early in the visual hierarchy. Attention is also sensitive to the establishment of visual objects: selection of one constituent object part often leads to prioritization of other locations on that object. But very little is known about the influence of reward on this object-based control of attention. Here we show in 4 experiments that reward prioritization and object prioritization interact in visual cognition to guide selection. Experiment 1 establishes groundwork for this investigation, showing that reward feedback does not negate object prioritization. In Experiment 2, we corroborate the hypothesis that reward prioritization and object prioritization emerge concurrently. In Experiment 3, we find that reward prioritization and object prioritization sustain and interact in extinction, when reward feedback is discontinued. We verify this interaction in Experiment 4, linking it to task experience rather than the strategic utility of the reward association. Results suggest that information gathered from locations on reward-associated objects gains preferential access to cognition.

Keywords: attention, object-based attention, associative learning, reward

Significance statements:

- Selective attention is biased to reward-associated stimuli and to locations on attended objects

- Here, we show that these effects combine, such that locations on reward-associated objects are selected preferentially.

- Information from these locations will gain preferential access to downstream cognitive processes like decision-making and motor control.

Introduction

Visuo-spatial attention has traditionally been characterized as under the control of two concurrent influences, such that stimuli with high physical salience and stimuli that match strategic attentional templates are both prioritized (eg. Egeth & Yantis, 1997). However, in recent years there has been increasing awareness of additional influences on attentional control that do not conveniently fit this dichotomy. For example, attention is directly sensitive to associative learning (Le Pelley et al., 2016, for review). In particular, stimuli that have been associated with reward will draw selective attention in the future, even when this provides no strategic benefit (eg. Anderson et al., 2011; Hickey et al., 2010a; Hickey & van Zoest, 2012). At the same time, attention has been drawn to part of an object, locations containing other parts of that object also become prioritized. There is ongoing debate regarding the degree to which reward prioritization (Anderson, 2019; Awh et al., 2012) and object prioritization (Peters & Kriegeskorte, 2021; Shomstein, 2012) can be characterized as goal-driven or strategic and relatively little is known about how these influences on attention combine.

Here, we investigate the relationship between reward prioritization and object prioritization, with particular interest in the possibility that these influences combine and interact to prioritize information gathering from reward-associated visual objects. To this end, we had participants complete experiments based on a well-established object-based attention paradigm known as the 2-rectangle task (Egly et al., 1994). In the paradigm, a trial begins with the presentation of two rectangles designed so that their ends define 4 screen positions that are equidistant from fixation. An exogenous cue - such as a flash of light subsequently draws attention to one of these locations. A target then appears, but, critically, the target does not always appear at the cued location. When the target appears at an uncued location, this can either be on the same rectangle as the cue or on the other rectangle.

The key observation in the 2-rectangle task is that when the target does not appear at the cued location, but does appear on the same object as the cue, responses are quicker and more accurate than when the target appears at any other uncued location. That is, attention appears to engage with the cued object as a whole, benefitting processing of stimuli that happen to appear at uncued locations on that object. Early interpretation suggested this reflected a low-level and automatic role of object prioritization in visual segmentation during perception of visual environments (Driver et al., 2001; Egly et al., 1994; Wannig et al., 2011). In line with this, object prioritization emerges when task confines give sufficient opportunity and motivation for objects to be derived from visual input (Chen & Cave, 2006, 2008). However, subsequent results have shown that object prioritization disappears when cue validity is high (Shomstein & Yantis, 2004; Yeari & Goldsmith, 2010), when the target location is endogenously identified (Macquistan, 1997), and when the cue location provides indirect information about the target location (Drummond & Shomstein, 2010). This has motivated the alternative proposal that object prioritization might be a default strategy adopted during search that can be discarded when the target location is unambiguous (Shomstein, 2012).

If object prioritization is strategic, this raises the possibility that it may be sensitive to the concurrent establishment of other strategies, and this idea motivated Shomstein and Johnson (2013) to conduct a series of experiments combining manipulation of reward outcome and object continuation in the 2-rectangle task. The experimental logic was that a strategy of attending to objects might be down-weighted, or even discarded, when task confines provided other, better ways to optimize performance. In Experiment 1, participants were consistently rewarded for accurately reporting targets that appeared on uncued objects. This reversed the pattern normally observed in the 2-rectangle task: responses became quicker when the target appeared on the uncued object rather than the cued object. By itself, this could simply mean that participants had learned to strategically deploy attention away from the cued object in order to optimize earnings, and that this strategic effect obscured any effect of object prioritization. However, the authors showed in a

subsequent experiment that object prioritization also disappeared when reward was not linked to any specific object. In Experiment 2, reward feedback was randomly of either highor low-magnitude whenever the target appeared at an uncued location. This removed the motivation to strategically attend to the uncued object, but the results again showed no evidence of object prioritization (see also, Lee & Shomstein, 2013).

This latter finding is puzzling. In line with the motivating hypothesis for the study, Shomstein and Johnson (2013) interpreted it as evidence of the strategic nature of object-based attention. That is, the authors suggested that participants discarded the strategy of attending to objects when another, better strategy was available. But it is unclear exactly what strategy participants might have adopted, or why they were motivated to make this strategic shift. Reward feedback in this experiment was random, so there was no opportunity to use outcome to optimize reward harvesting. As a result, there was no way to verify that participants actually changed strategy. Concerningly, this means that the interpretation offered by the authors - that participants discarded one strategy and adopted another - rests on the null observation of no object-based effect.

There are alternative accounts for the data pattern. One possibility is that the effect of object-based attention did not emerge in analysis simply because that effect is small and noisy. This possibility is consistent with subsequent results from Zhao et al. (2020). These authors employed a variant of the two-rectangle paradigm where the 'rectangles' were images of high-denomination and low-denomination monetary notes and data was collected from a larger sample than was employed in Shomstein and Johnson (2013; 30 vs 10 individuals). Results show a robust object-based effect when the cued note was of high-rectangle note, with responses particularly slow when the high-denomination note was cued but the target ultimately appeared on the low-denomination note.

However, there are limitations to the design adopted by Zhao et al. (2020). In particular, the value of the high- and low-denomination notes employed in that study reflected real-world experience over a long time frame and the visual characteristics of the notes could not be controlled or counter-balanced. It happened that the higher-value note was bright red, while the lower-value note was dull green: if the salience of the red note was greater than the green note - and red stimuli are known to be of particular salience (eg. Pomerleau et al., 2014) - the putative impact of value on object-based attention observed in this study could actually reflect the influence of note color. Zhao et al. (2020) partially addressed this issue in a control experiment, showing that the interaction of value and object status did not reliably emerge when the rectangles had note color but no other defining features. But the behavioral difference between experiments was small, and not statistically assessed, leaving ambiguity on the issue.

Another interpretation of the null result reported in Shomstein and Johnson (2013; Experiment 2) is that the object-based effect did not emerge in analysis because of conflicting sequential contingencies in the experimental design. In that experiment, reward was randomly determined to be of high- or low-magnitude in each invalidly cued trial, regardless of whether the target appeared on the same rectangle as the cue or on the other rectangle. This kind of random reward schedule is known to create inter-trial effects on behavior and brain activity that can be identified in sequential analysis (Hickey et al., 2010a, 2010b, 2015). The general observation is that when target selection in one trial results in high-magnitude reward, selective processing is biased toward similar stimuli in immediately subsequent task performance. When participants in Shomstein and Johnson (2013) successfully identified a target that appeared on the same object as the cue, and were rewarded for their performance, this may have created a bias toward cued objects in the next trial. This would create performance benefits when the relationship repeated between trials, but performance costs when this did not occur. Under these circumstances, mean results - collapsed over sequence - might show no evidence of object prioritization.

In light of these issues, the current study was designed to further investigate the relationship between reward prioritization and object prioritization in attentional control. We report results from 4 experiments using the 2-rectangle task. We begin in Experiment 1 by attempting to reproduce Shomstein and Johnson (2013; Experiment 2) in order to test the possibility that reward and object prioritization interact across trial contingencies, as

hypothesized above. To foreshadow, we find no evidence of this kind of sequential effect, but, in contrast to the original study, we do find evidence of object prioritization. In Experiment 2, we follow on from Zhao et al. (2020) to investigate how the association of reward to an object impacts object prioritization, importantly using a design that allows for counterbalancing of the association of reward to specific object features. In Experiment 3, we extend this design to measure the impact of reward on object prioritization during extinction, when reward contingencies are stopped. Finally, in Experiment 4, we directly compare results observed when reward associations are actively reinforced versus in extinction. Across experiments, the results show that reward prioritization and object prioritization guide attention concurrently. These effects initially appear as two independent influences on attention, but come to interact as participants gain task experience, regardless of the ongoing strategic utility of the reward association.

Transparency and openness

Sample sizes for each experiment are motivated from formal power analysis where possible and informal consideration of potential effect size otherwise and this is described in the 'Participants' section for each experiment. Power analyses were performed with G*Power software v3.1.9.6 with alpha of .05 (Faul et al., 2007, Erdfelder et al., 1996). All data exclusion parameters and data exclusions are described, as are all manipulations and measures. Age, sex, and nationality demographics are provided for each experiment; no other demographic information was considered. Results can be expected to generalize to the English-speaking, primarily high-income international population from which the sample was taken.

All data analysis was conducted using R 4.2.3 (R Core Team, 2020). Figures rely on output from the ggplot2 package for R 3.4.2 (Wickham, 2016) with subsequent adaptation in vector graphics software. None of the experiments were pre-registered. Data collection began in September, 2020, and completed in May, 2022. Data and research materials are

available at [UBIRA eData repository - University of Birmingham institutional repository, link to be inserted at publication].

Experiment 1

Participants

Sixty-nine participants were recruited online through Prolific (<u>www.prolific.co</u>). All participants provided informed consent, reported normal or corrected-to-normal visual activity, fluency in English and were naïve to the purpose of the experiment. Participants were excluded from analysis based on interquartile range (IQR) of cross-conditional mean accuracy or cross-conditional median reaction time (excluded when median reaction time (RT) > 3Q +1.5*IQR, mean accuracy < 1Q-1.5*IQR). This led to the exclusion of 3 individuals for a final sample of 66 (30 male, 35 female, 1 other; mean age 34 years ± 11 years SD; nationality: AU 1 CA 3 GR 1 IE 1 NO 1 NZ 1 SA 2 UK 51 US 5).



Figure 1. Task description. (**A**) In Experiment 1 the task starts with a 1 second preview screen containing only the two outline rectangles and a fixation dot, followed by the appearance of a red cue at the end of one of the rectangles. After a delay of 200ms, the target (an "L" or a "T") appears at one of three locations at the end of a rectangle. At the same time the 3 other corners are filled with distractors. The target screen is followed by another 100ms delay screen and then all the 4 positions are masked with a grid until response. A feedback screen subsequently indicates the amount of points earned or lost in the current trial and the total points accumulated so far by the participant. (**B**) Experiments 2 through 4 differ from Experiment 1 in 5 ways: (1) A 500ms screen containing only a fixation point was inserted at the start of every trial; (2) The two rectangles had different colors, which randomly changed for each trial; (3) The cue was a box, rather than outline; (4) The colored rectangle sustained during the feedback screen with 'Incorrect' was showed, no point was deducted. Additionally, in Experiment 3 and 4, the mask screen was sustained until response or 1700ms. (**C**) Feedback conditions for correctly performed trials differed across was initially the case, but reward was discontinued in a later experimental phase. In Experiment 4, the LL group completed a task very similar to that used in Experiment 2 and the LE group completed a task very similar to Experiment 3.

The sample size was guided by power analysis of the object-based effect identified in Shomstein et al. (2013, Experiment 1a, ANOVA $\eta_p^2 = .315$). This identified the need for 22 participants to reproduce, assuming power of .8. We approached data collection with the expectation that any sequential effect of reward on object prioritization - the *a priori* target of this experiment - would have a substantially smaller effect than this. Accordingly, we increased our sample to triple this estimate in the hope this would provide sufficient power to detect a sequential effect of unknown size.

Design and procedure

The task was built in Python using Opensesame software (version 3.3.6), converted to Javascript using OSWeb (Mathôt et al., 2012), and further adapted in Javascript where necessary. The web application Jatos (Lange et al., 2015) was used to host the experiment on a computer server and participants completed the experiment in a web browser on a personal computer in their own environment. In order to standardize stimulus size across settings, participants were asked to maintain a distance of 60 cm from their computer screen and asked to adapt the web browser magnification such that an example rectangle on the screen fit the size of a standard bank card. The experiment began with detailed instructions that emphasized both speed and accuracy.

The experiment design is similar to Shomstein and Johnson (2013, Experiment 2). As illustrated in Figure 1a, in each experimental trial participants were asked to discriminate the identity of the target, which could either be a 'L' or 'T' (1.2° x 0.9° visual angle) and was presented at one of the four corners of the computer screen (4.4° visual angle from center). The other three locations contained non-targets, which were shapes created by superimposing the target letters and randomly orienting this image 90°, 180°, or 270° off vertical.

In order to link targets and non-targets to the same or different visual objects, each trial began with the presentation of two rectangles that each encompassed 2 of the 4 target and nontarget stimuli locations ($10^{\circ} \times 2.5^{\circ}$ visual angle). The rectangles were either oriented

vertically, such that the two left stimuli locations and two right stimuli locations appeared on separate objects, or horizontally, such that the two top stimuli locations and the two bottom stimuli locations appeared on separate objects, and this was counterbalanced across participants. The rectangles were presented for 1 s before one of the target locations was cued. The cue took the form of a brightening and widening of the lines defining one end of the rectangle (200 ms; 2.4° visual angle). The cue appeared with equal likelihood at the location of the upcoming target, at the other location on the same rectangle, or at the equidistant location on the other rectangle. A pause of 200 ms followed, after which the target and nontargets appeared for 60 ms before a delay of 100 ms and subsequent onset of mask stimuli (crossed squares; 2.4° x 2.4° visual angle). Mask stimuli sustained until response was registered, at which point a new trial began. Participants were instructed to report the target identity via key-press with the index finger on a standard keyboard. For all participants, left-hand response indicated that the target was a 'T' and right-hand response indicated that the target was an 'L'.

Each trial concluded with reward feedback. When the cue correctly identified the target location, correct performance earned 1 point, but when the cue was invalid, correct performance could randomly lead to either 1 or 6 points. Errors led to the loss of equivalent points (-1 point in valid trials and randomly -1 or -6 points in invalid trials). Participants were informed before taking part in the experiment that each 35 points led to payment of £0.10. Accuracy feedback was provided at the end of every experimental block, alongside information on points earned in that block and points earned overall, and participants were paid at the end of the experiment based on the total points accumulated. Participants completed 24 practice trials followed by 9 blocks of 72 trials, the experiment took approximately 45 minutes to finish, and average pay was £4.10 (\pm £0.41 SD).

Results

Trials with RT greater than 1700 ms were excluded from the analysis (1.81% of trials, \pm 3.31% SD). This cutoff was applied to Experiments 1 and 2 in order to equate analytic

parameters with Experiments 3 and 4, where there was no opportunity to respond beyond 1700 ms after stimulus onset (see the treatment of these experiments below for more detail).

The experimental design generated 3 equally-likely configurations between target and cue: the target either appeared on the same position of the cue (valid), on the other location on the same object (invalid same-object; SO), or at the nearest location on the other object (invalid different-object; DO). Figure 2A illustrates reaction times for these three conditions. Unexpectedly, given existing results from Shomstein and Johnson (2013; Experiment 2), we identified an object-based attentional effect, with faster RT in the SO condition than in the DO condition (683 ms vs 701 ms; t(65) = 4.13, p < .001, *d* = 0.162). A significant cueing effect was also evident in the difference between valid and invalid trials (collapsed across SO and DO conditions; 664 ms vs 692 ms; t(65) = 2.91, p = .005, *d* = 0.285). These effects were mirrored in accuracy: participants were significantly better in validly cued trials than invalidly cued trials (88.9% vs 85.7%; t(65) = 2.70, p=.009, *d* = 0.304) and were better in the SO condition than the DO condition (86.2% vs.



Figure 2. RT results from Experiment 1. Here and in subsequent figures, black circles and squares represent mean conditional reaction times. In gray, we illustrate the distribution of participant mean performance per condition and the conditional effect for each participant. The median, first, and third quartile are indicated for each distribution. Panel **A** illustrates results per cue-target relationship. An object prioritization effect is evident in the contrast of DO and SO conditions. Panel **B** illustrates results as a function of trial sequence (inconsistent: the target appeared on the cued object in the current trial, but appeared on the uncued object in the preceding trial, or, vice versa, the target appeared on the uncued object in the current trial but the cued object in the preceding trial; consistent: the target appeared on the cued object in two sequential trials, or the target appeared on the uncued object in two sequential trials) and the magnitude of reward received in the preceding trial (high or low).

85.1%; t(65) = 2.10, p = .040, d = 0.104). Statistical analysis of accuracy here and in subsequent experiments is based on data transformed to approximate normality (Box & Cox, 1964).

As described above, we approached Experiment 1 with the idea that intertrial sequences might influence performance in this task, creating implicit expectations regarding the relationship between the cue and target locations. To assess this possibility, we divided invalidly-cued trials based on repeated cue-target relationship. Consistent trials were those where either the target appeared on the same object as the cue for consecutive invalidly

cued trials, or the target appeared on the uncued object for consecutive invalidly cued trials. Inconsistent trials were invalidly cued trials where the target appeared on the same object as the cue in one trial but on the uncued object in the next trial, or vice versa. Our expectation was that consistency would interact with the magnitude of reward received for the preceding trial: when participants received high-magnitude reward in trial n-1, and the cue-target relationship from trial n-1 was repeated into trial n, participants should be faster to respond to the target. In contrast, when participants received high-magnitude reward in trial n-1, but the cue-target relationship was not repeated, participants would be slower. However, as illustrated in Figure 2B, no effect of trial sequence emerged in our results. A two-way repeated measures ANOVA (RANOVA) with factors for consistency and reward magnitude in the immediately preceding trial showed no significant effect of prior reward, consistency or interaction of these factors (all Fs<1).

Discussion

Experiment 1 was motivated by the idea that object-based attention and reward-driven attention might co-exist in visual cognition, but that these influences were hidden in results from Shomstein and Johnson (2013; Experiment 2) by inter-trial effects on performance. We expected to reproduce the null result observed by Shomstein and Johnson (2013; Experiment 2) in core analyses and add new perspective in analysis of trial sequence. Instead, we found a robust effect of object prioritization but no significant effect of sequence on object prioritization. As a null result, the non-significant effect of inter-trial sequence is, of course, ambiguous. However, the results suggest that if sequence has the impact on performance we expected, this effect is small and will require a very large sample to detect.

We have no clear account for the disparity between the current results and those from Shomstein and Johnson (2013; Experiment 2), other than that the null object prioritization effect observed in the earlier work was a type II error.

Results from Experiment 1 open the possibility that object prioritization and reward prioritization might co-exist. If this is the case, what influence do they have on one another?

One possibility is that they interact, as would be the case if attention were to engage with objects with speed or strength that differs as a function of the object reward association. As noted in the Introduction, results from Zhao et al. (2020) suggest such an effect, but are arguably ambiguous due to low-level confounds in the design. To further test this idea while controlling for low-level visual properties of the stimuli we conducted a second experiment where the two rectangles each carried a unique color. For each participant, one color was associated with high-magnitude reward: correct response to a target appearing on the object characterized by this color garnered more points. If reward interacts with object-based attention, we expected the object-based effect to be accentuated when it was the reward-associated object that was cued.

Experiment 2

Participants

Experiment 2 was based on a new sample of 33 participants and used the same recruitment and exclusion procedures as for Experiment 1. Five participants were rejected from analysis due to outlier performance, leading to a final sample of 28 individuals (9 male, 19 female, 0 other; mean age 32 years ± 10 years SD; nationality CA 2 SA 2 UK 22 US 2). The final sample size was determined by power analysis of the object-based effect observed in Experiment 1. Calculation of power was based on the effect size for the paired t-test contrasting SO and DO conditions ($d_z = 0.508$) with assumed power of .8.

Design and procedure

All data collection took place online and the procedure closely matched that of Experiment 1, except that each of the two rectangles had a filled color and the reward schedule was linked to these colors. Rectangle colors were pseudo-randomly selected for each trial so that the two rectangles never had the same color for a given trial and the colors were drawn from a

set of five possibilities: red (RGB: 160,0,0), blue (34, 5, 255), purple (141, 21, 124), green (18, 90, 23), and brown (127, 58, 6). For each participant, one of the five colors was associated with high-magnitude reward: when the target appeared on the rectangle with this color, a correct response earned 100 points. If the target appeared on an object of any other color, correct performance earned 5 points. The association of reward outcome to specific colors was described to participants before they began the experiment and was counterbalanced across participants. Participants earned £0.05 for every 95 points and no points were deducted for incorrect responses. The design otherwise differed from Experiment 1 in that rectangle onset was preceded by a shorter fixation duration (500 ms), the cue was defined by a white square ($2.4^{\circ} \times 2.4^{\circ}$ visual angle) that appeared at the end of one of the rectangles (see Figure 1B) and participant were presented with 10 blocks of 54 trials. The experiment took approximately 45 minutes to complete and average pay was £5.15 (± £0.38 SD).

Results

As in Experiment 1, trials with RT greater than 1700 ms were excluded from analysis (2.51% of trials, ± 3.01% SD).

To assess a possible influence of reward on the object-based effect, we divided trials into conditions based on cue validity and reward association. High-reward trials (1/6 of trials) occurred when one of the two rectangles was characterized by the high-reward color and the target appeared on the high-reward object, whereas low-reward trials (1/6 of trials) occurred when one of the two rectangles was characterized by the high-reward color but the target appeared on the low-reward object. Finally, baseline trials (2/3 of trials) occurred when neither rectangle was characterized by the high-reward color.

Validly Cued Condition

Figure 3A illustrates the results for validly cued trials. A one-way RANOVA with a factor for reward (baseline, high and low reward conditions) identified a significant main effect (Greenhouse-Geisser correction; F(1.26,34.12) = 82.65, p < .001, η_p^2 = .754). Pairwise comparisons showed that RT was quicker in high-reward trials than in baseline (t(27) = 2.59, p = .015, *d* = 0.251) and slower in low-reward trials than in baseline (t(27) = 9.81, p < .001, *d* = 1.223). Analysis of accuracy also showed a significant main effect of reward (Greenhouse-Geisser correction; F(1.30,35.03) = 30.94, p < .001, η_p^2 = .534). Accuracy was improved for high-reward compared to low-reward trials (94.0% vs 79.6%, t(27) = 5.72, p < .001, *d* = 1.328), for baseline compared to low-reward trials (91.5% vs 79.6%, t(27) = 6.30, p < .001, *d* = 0.850), and approached corrected significance for baseline compared to high-reward trials (94.0% vs 91.5%, t(27) =2.24, p = .033, *d* = 0.457).

Invalidly Cued Condition

As illustrated in Figure 3B, the pattern observed in validly cued conditions also appears in analysis of invalidly cued trials, with faster responses in high-reward trials than in baseline and slower responses in low-reward trials than in baseline. An additional effect of object prioritization emerged in these results, with faster responses when the target appeared on



Figure 3. Results from Experiment 2. Panel **A** illustrates reaction times from the valid cue conditions. Panel **B** illustrates reaction times from the invalid cue conditions. Notably, the object prioritization effect emerges both when the target appears on the high-reward associated object and the low-reward associated object, with no appreciable distinction between these conditions.

the cued object versus when it appeared on the uncued object. We began analysis by identifying a significant object prioritization effect in the baseline condition, when both rectangles were characterized by low-reward color (t(27) = 5.58, p < .001, *d* = 0.351). The same effect of object prioritization emerged in analysis of accuracy in the baseline condition (91.9% vs 90.0%, t(27) = 2.70, p = .011, *d* = 0.399). We subsequently conducted a RANOVA of invalidly-cued data with factors for reward (target on high-reward colored object, target on low-reward colored object) and cue-target relationship (cue identified rectangle where target appeared, cue identified other rectangle). This identified a main effect of reward (F(1,27) = 93.57, p < .001, η_p^2 = .776) and a main effect of cue-target relationship (F(1,27) = 13.73, p < .001, η_p^2 = .337) but no significant interaction (F(1,27) = 1.44, p = .240, η_p^2 = .051). Results from the invalidly-cued trials thus show an effect of reward prioritization and an effect of

object prioritization, but no significant relationship between these effects. Accuracy results show the same main effect of reward (93.6% high vs 78.3% low reward, F(1,27) = 45.56, p < .001, η_p^2 = .628), but no significant effect of cue-target relationship (86.9% same vs 84.9% different object, F(1,27) = 2.46, p = .128, η_p^2 = .084) or interaction (F<1).

Discussion

Results from Experiment 2 suggest that while reward prioritization has a robust impact on performance, it does not substantively change the impact of object prioritization. There is the possibility that the absence of significant interaction reflects lack of statistical power or limitations to the sensitivity of ANOVA to interactions in the presence of main effects. Experiment 3 both addresses this ambiguity and tests new hypotheses.

In Experiment 2, participants were explicitly informed of the association between reward and color, and we expect this motivated them to establish a strategy to prioritize the reward-associated visual feature. This kind of strategic attentional control can have a strong impact on the prioritization and perception of features and objects (eg. Folk, Remington, & Johnston, 1992). However, as noted in the Introduction, reward can also have a more direct and automatic impact on perception and prioritization that sustains when it is no longer useful (Hickey, Chelazzi, & Theeuwes, 2010a, b; Anderson et al., 2011). It is unclear if the reward effect identified in Experiment 2 reflects the strategic effect of reward, associative learning, or both. Moreover, there is the possibility that a low-level, automatic effect of reward - driven by associative learning - might interact with object prioritization in ways that do not robustly emerge when reward is used to strategically guide selection.

Experiment 3 was designed to address these issues. In an initial phase of Experiment 3 - the *learning phase* - participants completed a task identical to that of Experiment 2, with color consistently predicting the reward outcome of correct performance. However, in a second stage of the experiment - the *extinction phase* - participants were informed that reward was no longer available, but they would have to complete the remainder of the task

to receive the reward they had earned in the earlier stage. We approached this experiment with 3 key questions. First, will reward prioritization sustain in the extinction phase, when its strategic utility is removed? Second, does object prioritization also sustain in this phase of the experiment? And, finally, if both effects occur in extinction, do they show the additive relationship identified in Experiment 2?

Experiment 3

Participants

Experiment 3 was based on a new sample of 61 participants and used the same recruitment and exclusion procedures identified for Experiment 1. Three participants were rejected, leading to a final sample of 58 individuals (34 male, 24 female, 1 other; mean age 37 years \pm 12 years SD, one participant withheld report of age; nationality CA 2 HU 1 IN 1 IR 4 PL 2 SA 1 UK 40 US 4 ZW 1, one participant withheld report of nationality). The sample size was determined by consideration of the object-based effects observed in Experiments 1 (d_z



Figure 4. Results from Experiment 3. Reaction times from the learning phase are illustrated in panel **A.** The object prioritization effect emerges when the target appears on the high-reward associated object, and also when it appears on the low-reward associated object, with no appreciable difference in magnitude. Reaction times from the extinction phase are illustrated in panel **B.** The object prioritization effect appears larger when the target appears on the low-reward associated object.

= 0.508) and 2 (η_p^2 =.337). Our expectation was that the raw magnitude of this effect might reduce with the increased length of Experiment 3, as overall reaction times became faster, and thus that the effect size after training would be smaller than observed in Experiments 1 and 2. We therefore targeted a sample size twice that employed in Experiment 2. Fewer participants exhibited outlier performance than was predicted by results from Experiment 2, leading to slight over-recruitment (of 58 individuals rather than 56).

Design and procedure

In Experiments 1 and 2, we observed that participants occasionally failed to respond promptly to a trial, suggesting that they took impromptu breaks within an experimental block, and this introduced minor complications in analysis of results. To ensure that participants took breaks only within the block structure, we introduced a response time limit. Participants were required to respond within 1700 ms of stimulus onset and a trial ended either when response was made or this interval had passed.

The learning phase of Experiment 3 was similar to the design of Experiment 2, with the rectangle colors predicting reward outcome (2 blocks of 72 trials). Reward was discontinued in the subsequent extinction phase (8 blocks of 72 trials; see Figure 1C). Participants were informed that they needed to maintain accuracy of 85% or greater in order for earnings from the learning phase to be paid out at the end of the experiment (or they would receive a lesser base rate payment of £5.85). During the learning phase, participants earned £0.26 for every 100 points accumulated, with no points lost for incorrect answers. Participants completed the experiment in approximately 70 minutes and average pay was £7.03 (± £0.3 SD).

Results

In 0.52% (\pm 0.92% SD) of trials participants did not respond within 1700 ms of stimulus onset. These trials were excluded from calculation of accuracy.

Validly Cued Conditions

For the sake of completeness, we analyzed the validly cued experimental conditions, though these results test no experimental hypotheses. A two-way RANOVA with factors for reward (baseline, high-reward and low-reward) and phase (learning and extinction) showed statistical significance of both main effects and the interaction (Greenhouse-Geisser correction; reward: 647 vs 646 vs 706 ms, F(1.32,75.30) = 26.88, p < .001., η_p^2 = .320;

phase: 728 vs 604 ms, F(1.00,57.00) = 191.74, p < .001, η_p^2 = .771; phase x reward: (693 vs 691 vs 800) vs (601 vs 600 vs 611) ms, F(1.37,78.37) = 20.80, p < .001, η_p^2 = .267). Analysis of accuracy showed similar results (Greenhouse-Geisser correction; reward: 94.3% vs 95.2% vs 88.7%, F(1.47,83.60) = 24.08, p < .001., η_p^2 = .297; phase: 89.5% vs 96.0%, F(1.00,57.00) = 74.73, p < .001, η_p^2 = .567; phase x reward: (92.1% vs 94.2% vs 82.1%) vs (96.5% vs 96.2% vs 95.3%), F(1.72,98.29) = 21.18, p < .001, η_p^2 = .271). The interaction effects highlighted by the analysis could be due to floor and ceiling effects that affected the baseline and high reward condition more than the low reward condition, as the latter had more room for improvement over the course of the experiment.

We followed up on these results with separate analysis of validly cued conditions in each of the learning and extinction phases separately. Analysis of the learning phase constituted a reproduction of Experiment 2, and results were accordingly similar. A one-way RANOVA of RT from the learning phase identified a significant main effect of reward (Greenhouse-Geisser correction; F(1.32,75.24) = 25.82, p < .001, η_p^2 = .312). Pairwise comparisons identified quicker RT in both baseline and high-reward trials compared to low-reward trials (693 vs 800 ms, t(57) = 5.82, p < .001, *d* = 0.714; 691 vs 800 ms, t(57) = 5.04, p < .001, *d* = 0.731). Unlike Experiment 2, no significance difference was found between baseline and high-reward trials (693 vs 691 ms, t(57) = 0.23, p = .816, *d* = 0.020). Analysis of accuracy mirrored the RT results. A one-way RANOVA showed a significant main effect of reward (Greenhouse-Geisser correction; F(1.58,90.31) = 26.19, p < .001, η_p^2 = .315) with no significant difference between high-reward condition and baseline (94.2% vs 92.1%, t(27) = 1.92, p = .060, *d* = .298), but a significant difference between high-reward and low-reward conditions (94.2% vs 82.1%, t(27) = 5.79, p < .001, *d* = 1.013) and between baseline and low-reward conditions (92.1% vs 82.1%, t(27) = 5.25, p < .001, *d* = 0.795).

In validly cued conditions of the extinction phase, reward did not have a significant effect on RT (601 vs 600 vs 611 ms, Greenhouse-Geisser correction; F(1.56,88.87) = 2.84, p = .076, η_p^2 = .047) or accuracy (96.5 vs 96.2 vs 95.3 ms, Greenhouse-Geisser correction; F(1.66,94.72) = 1.74, p = .186, η_p^2 = .030). Descriptive statistics for RT and accuracy are in line in the direction of a facilitation for baseline and high-reward over low-reward trials. For this reason we calculated and analyzed combined accuracy and RT scores (inverse efficiency; Townsend & Ashby, 1978, 1983). This identified a main effect of reward (624 vs 626 vs 645, Greenhouse-Geisser correction; F(1.55,88.4) = 4.48, p = 0.022, η_p^2 = .073). Pairwise comparisons showed a significance difference between baseline and low-reward conditions (t(57) = 2.72, p = .009, *d* = 0.187), but no other significant differences (high-reward vs low-reward: t(57) = 2.01, p = .049, *d* = 0.170; baseline vs high reward: t(57) = 0.40, p = .69, *d* = 0.023).

Invalidly Cued Conditions

Figure 4 presents RT results from invalidly cued conditions in Experiment 3. Analysis of these results tests core experimental hypotheses. As illustrated in Figure 4A, results from the learning phase replicated Experiment 2. A RANOVA with factors for reward (target at high-reward colored object, target at low-reward colored object) and cue-target relationship (same rectangle, different rectangle) identified main effects of reward (F(1,57) = 75.90, p < .001, η_p^2 = .57) and cue-target relationship (F(1,27) = 9.03, p = .004, η_p^2 = .137) with no significant interaction between these factors (F<1). Similar analysis of accuracy revealed a main effect of reward with improved performance with the target appearing on the rectangle with high-reward color (92% vs. 83.1%; F(1,57) = 25.60, p < .001, η_p^2 = .310). No other effects emerged (88.3% vs 86.8%, cue-target relationship: F(1,57) = 1.29, p = .261, η_p^2 = .022; interaction: F<1).

Analysis of results from the extinction phase also identified main effects of reward $(F(1,57) = 13.01, p < .001, \eta_p^2 = .186)$ and cue-target relationship $(F(1,57) = 5.03, p = .029, \eta_p^2 = .081)$. However, a significant interaction between these factors also emerged $(F(1,57) = 5.32, p = .025, \eta_p^2 = .085)$. A corresponding effect of reward emerged in analysis of accuracy (94.2% vs 92.6%; $F(1,57) = 4.02, p = .050, \eta_p^2 = .066$) but there was no evidence of the

main effect of cue-target relationship (93.5% vs 93.2%, F<1) or interaction (F(1,57) = 2.40, p = .127, η_p^2 = .040).

We conducted an omnibus RANOVA to determine if the interaction identified in the extinction phase was reliably different from the interaction identified in the learning phase. This was based on results from all invalidly-cued trials and had factors for experiment phase (learning, extinction) as well as reward and cue-target relationship. The 3-way interaction was not significant (F<1). In line with group-wise results described above, this analysis otherwise identified 3 significant main effects (reward: F(1,57) = 78.12, p < .001, η_p^2 = .578; cue-target relationship: F(1,57)=13.78, p < .001, η_p^2 =0.195; phase: F(1,57) = 186.22, p < .001, η_p^2 =0.766) and an interaction of reward and phase (F(1,57) = 58.93, p < .001, η_p^2 = .508), but no other effects (reward X cue-target relationship: F<1; cue-target relationship X phase: F(1,57)= 3.23, p = .078, η_p^2 = .054). Analysis of accuracy identified a complementary pattern, with effects of reward (93.1% vs 87.8%; F(1,57) = 23.66, p < .001, η_p^2 = .293), experimental phase (87.5% vs 93.4%; F(1,57) = 57.01, p < .001, η_p^2 = .500) and an interaction of reward with phase (F(1,57) = 57.01, p < .001, η_p^2 = .239), but no other effects (90.9% vs 90.0%, cue-target relationship: F(1,57) = 1.23, p = .272, η_p^2 = .021; three-way interaction: F(1,57) = 1.15, p = .288, η_p^2 = .020; all other Fs<1).

Discussion

Results from Experiment 3 indicate that reward prioritization sustains in the extinction phase, consistent with the idea of mechanism that is relatively insensitive to shifts in strategy (cf. Hickey, Chelazzi, & Theeuwes, 2010a, b; Anderson et al., 2011; Anderson & Yantis, 2013). Importantly, the effect of cue-target relationship - reflecting object prioritization - emerged in both the learning phase and the extinction phase.

Results from the learning phase reproduce results from Experiment 2. However, in contrast to Experiment 2 and the learning phase of Experiment 3, a statistical interaction of

reward prioritization and object prioritization emerged in the extinction phase of Experiment 3. The object prioritization effect had greater strength when the target ultimately appeared on the low reward rectangle. This interaction is similar to that observed in Zhao et al. (2020; Experiment 2), where a nominal interaction of reward prioritization and object prioritization also emerged. As described above, visual objects in Zhao et al. (2020; Experiment 2) were images of monetary notes, which presumably had associations of value, but the presence of these objects did not signal actual receipt of reward. Similarly, in the extinction phase of the current experiment the visual objects had been associated with reward, but did not signal actual receipt of reward. In both studies, participants appear to have a particularly hard time orienting attention away from locations on the high-reward object when that object has been cued, but the target has appeared elsewhere.

Results from Experiments 2 and 3 tentatively suggest that when an object reward association is strategically useful - when it validly predicts monetary outcome - reward prioritization and object prioritization emerge as independent, additive influence on attentional control. However, when a reward association is discontinued - in extinction reward prioritization and object prioritization come to interact. But there are cogent reasons to delay this conclusion. First, the interaction of reward and cue-target relationship identified in the extinction period of Experiment 3 (which was significant) was not reliably different from the interaction identified in the learning phase (which was not significant). This highlights limitations in the experimental design. First, comparison of results from learning and extinction phases confounds the experimental manipulation of reward feedback with task sequence: the extinction phase necessarily follows the learning phase. This raises the possibility that the interaction emerges as a function of task familiarity and practice, rather than the manipulation of reward feedback. This is complicated by the fact that Experiment 3 was substantively longer than Experiment 2. Second, the extinction phase necessarily follows the learning phase, and performance is therefore more stable in the extinction phase. There is therefore more variance in behavior in the learning phase than in the extinction

phase, which can render statistical estimates unreliable. This difference in variance is exacerbated by the difference in duration of the experimental phases: the short first phase of the experiment has fewer trials than the longer second phase, meaning that estimates of performance are noisier.

These shortcomings of Experiment 3 motivated the need for an additional experiment to explicitly determine if the interaction of object prioritization and reward prioritization observed in the extinction phase of Experiment 3 was caused by the manipulation of reward utility. In Experiment 4, two separate groups of participants each completed an independent version of our task. For one group - the *learning-extinction* (LE) group - the task was nearly identical to that of Experiment 3. That is, participants initially completed a learning phase, when points could be earned, followed by an extinction phase, where there was no reward feedback. For the other group - the *learning-learning* (LL) group - the task was more similar to that of Experiment 2, with the learning phase sustaining until the end of the experiment. This design meant that in the second phase of the experiment LE participants had the same amount of practice as LL participants and that performance estimates in each group were based on an equivalent number of trials. If reward prioritization and object prioritization come to interact as a function of the manipulation of reward feedback, results in Experiment 4 should show an interaction of reward and cue-target relationship in the extinction phase of the experiment in the LE group, but not the LL group.

Experiment 4

Participants

Experiment 4 was based on a new sample of 53 participants and used the same recruitment and exclusion procedures identified for Experiment 1. Five participants were rejected, leading to a final sample of 48 individuals that were randomly assigned into two equal groups (LL group: 15 male, 9 female, 0 other; mean age 27 years ± 5 years SD; nationality BD 1 BE 1 CA 1 GR 1 HU 1 IT 1 LV 1 MX 1 NG 1 PO 7 PT 1 SA 4 UK 3; LE group: 17 male, 7 female, 0 other; mean age 25 years ± 7 years SD; nationality CZ 1 DK 1 GR 1 HU 1 IT 2 MX 1 PO 4 PT 5 SA 5 UK 2 US 1). The sample size was determined using the effect size of the interaction between factors for reward and cue-target in Experiment 3 (η_p^2 = .085). Power analysis suggested the need for 24 participants to detect an effect of reward on difference scores derived from the effect of cue-target relationship on RT, based on assumed power of .8.

Design and procedure

Experiment 4 had two phases: a learning phase (2 blocks of 72 trials) that was closely modeled on the learning phase of Experiment 3, and a second phase (4 blocks of 72 trials) that differed between groups. For the LL group, the second phase was identical to the first. For the LE group, the second phase was an extinction phase, with no reward feedback. In order to equate the two conditions in terms of total reward incentive, participants in the LE group were informed that participation in the entire experiment would lead to earnings of 4 times the amount accumulated in the learning phase. The base pay rate adopted in Experiment 4 was £4. As more total points were accumulated by participants in Experiment 4 than in Experiment 3, points were associated with less cash value (£0.10 per 100 points). All other other design characteristics were as in Experiment 3. The experiment took approximately 70 minutes and average pay was £8.38 (\pm £0.46 SD) for the LL group and £8.14 (\pm £0.61 SD) for the LE group.

Results

In 0.74% (\pm 1.49% SD) of trials participants did not respond within 1700 ms of stimulus onset. These trials were excluded from calculation of accuracy. Performance in the common learning phase was similar in both groups (mean RT 687 ms, \pm 76.8 SD; 682 ms, \pm 119 SD). As there was no *a priori* motivation to expect a difference between these groups, who

completed the same task under the same task instructions, we collapse across groups in analysis of the learning phase of the experiment.

Validly Cued Conditions

First experimental phase

As in experiment 3, we report statistical analysis of validly cued conditions, though no critical hypotheses are tested. A one-way RANOVA on RT observed in the learning phase, collapsed across the LL and LE groups, identified a significant effect of reward (baseline with both rectangles characterized by low reward color, target on high reward rectangle, target on low reward rectangle; F(1.38,65.08) = 24.15, p < .001, η_p^2 = .339). A pairwise comparison identified significant differences between baseline and low reward conditions (641 vs 758, t(47) = 5.50, p < .001, *d* = 0.783) and high reward and low reward conditions (651 vs 758, t(47) = 4.90, p < .001, *d* = 0.704) but not between baseline and high reward conditions (641 vs 651, t(47) = 0.93, p = 0.358, *d* = 0.090). Accuracy results also showed a main effect of reward (F(1.59,74.83) = 4.06, p = .029, η_p^2 = .080). Pairwise comparison showed higher accuracy for high-reward trials compared to low-reward trials (93.0% vs 87.2%, t(47) = 2.34, p = .024, *d* = 0.468). No significant difference was identified for the remaining contrasts (baseline vs low reward: 91.1% vs 87.2%, t(47) = 1.78, p = .082, *d* = 0.293; high reward vs baseline: 93.0% vs 91.1%, t(47) = 1.42, p = .161, *d* = 0.227).

Second Experimental Phase

Analysis of the second experimental phase began with an omnibus RANOVA with between-participant factor group (LL, LE), and within-participant factor reward (a baseline with only low reward colored rectangles on the screen, target appears on rectangle with high-reward color, target appears on rectangle with low-reward color). A significant main effect of reward emerged (Greenhouse-Geisser correction; F(1.29,59.45) = 27.55, p < .001, η_p^2 = .375), as well as the interaction (F(1.29,59.45) = 15.55, p < .001, η_p^2 = .253), but the

effect of group was not significant (F(1.00,46.00) = 1.49, p = 0.229, η_p^2 = .031). Further analysis showed that the interaction is driven by a prominent difference in RTs in the low reward condition for LL and LE groups (680 vs 586, Welch-Satterthwaite correction; t(40.6) = 3.10, p = .004, *d* = 0.894), in contrast to no significant difference in the baseline (566 vs 569, t(45.6) = 0.11, p = .915, *d* = 0.031) or high reward conditions (563 vs 570, t(41.6) = 0.30, p = .768, *d* = 0.086).

Analysis of accuracy also identified a significant main effect of reward (Greenhouse-Geisser correction; F(1.57,72.06) = 5.16, p = .013, η_p^2 = .0.101) and interaction (F(1.57,72.06) = 7.88, p = .002, η_p^2 = .0.146), but no significant effect of group (F<1). Further analysis of accuracy mirrored the RT analysis with a significant difference between LL and LE group when the target appeared on the rectangle with the low-reward color (88.8% vs 95.1%, Welch-Satterthwaite correction; t(31.3) = 2.33, p = .027, d = 0.672), but no significant difference between groups for baseline (94.3% vs 96.0%, t(39.9) = 1.02, p = .314, d = 0.294) or high reward conditions (96.6% vs 93.8%, t(42.1) = 1.59, p = .0.12, d = 0.459).

Invalidly Cued Conditions

First Experimental Phase

Figure 5 illustrates results from invalidly cued trials. Analysis of these results tests core experimental hypotheses. Figure 5A presents RT results for the learning phase collapsed across the LL and LE groups. We conducted a RANOVA analysis of RT observed in invalidly cued trials when a high-reward object was present on the screen. This had factors for reward (target appears on high reward object, target appears on low reward object) and cue-target relationship (target appears on cued rectangle, target appears on other rectangle) and replicated results of Experiments 2 and 3, showing significant main effects of reward (F(1,47) = 29.56, p < .001, $\eta_p^2 = .386$) and cue-target relationship (F(1,47) = 4.19, p = .046, $\eta_p^2 = .082$) but no significant interaction (F(1,47) = 1.50, p = .227, $\eta_p^2 = .031$).

Corresponding analysis of accuracy identified a main effect of reward (F(1,47) = 27.75, p < .001, η_p^2 = .371) but no effect of cue-target relationship (F(1,47) = 2.10, p = .154, η_p^2 = .094). In contrast to RT, analysis of accuracy identified an interaction (F(1,47) = 4.90, p = .032, η_p^2 = .094). Accuracy was better when the target appeared on a cued low-reward object than when it appears on an uncued low-reward object (81.8% vs 75.5%). This did not emerge when the target appeared on a high-reward object (90.4% vs 90.9%). We expect that this is driven at least in part by a ceiling effect on results from the high-reward condition. Analysis of inverse efficiency was in line with RT results, with a significant effect of reward (751 vs 1072, F(1,47) = 37.40, p < .001, η_p^2 = .443) but no evidence of cue-target relationship (881 vs 942, F(1,47) = 2.29, p = .137, η_p^2 = .046) or interaction (F<1).

Second Experimental Phase

Figures 5B and 5C present RT results from the second experimental phase for each of the LL and LE groups. Experiment 4 was conducted to test if the interaction of reward prioritization and object prioritization would emerge in each of the LL and LE groups, and analysis accordingly began with independent ANOVAs examining RT results in each group.

In the LE group, analysis identified a main effect of reward (614 vs 636, F(1,23) = 8.78, p = .007, η_p^2 = .276), no significant effect of cue-target relationship (622 vs 627, F<1), but, critically, an interaction between these factors (F(1,23) = 5.11, p = .034, η_p^2 = .182). Analysis of accuracy garnered similar results (reward: 93.2 vs 90.5, F(1,23) = 3.74, p = .066, η_p^2 = .140 ; cue-target relationship: 92.1 vs 91.7, F<1; interaction: F<1). In the LL group, analysis identified a main effect of reward (591 vs 724, F(1,23) = 28.12, p < .001, η_p^2 = .550), a main effect of cue-target relationship (643 vs 672, F(1,23) = 6.24, p = 0.02, η_p^2 = .213), and an interaction that approached significance (F(1,23) = 3.06, p = .093, η_p^2 = .117). The broad similarity in statistical results for each of the groups was also evident in analysis of inverse efficiency scores (LE 2-way interaction: F(1,23) = 4.21, p = 0.052, η_p^2 = .155; LL 2-way interaction: F(1,23) = 4.02, p = 0.057, η_p^2 = .149).

The interaction effects identified in each of the two groups have the same direction and are of similar standardized effect size. In raw data, the RT interaction is in fact larger in the LL group. The 2-way interaction of object prioritization and reward prioritization therefore does not appear constrained to the LE group. In line with this, in a larger analysis of RT- with factors for group (LL, LE), reward, and cue-target relationship - the 2-way interaction was significant (F(1,46) = 7.01, p = .011, η_p^2 = .132) but the 3-way interaction was not (F<1). This analysis otherwise paralleled the separate analyses of each group, identifying additional main effects of reward (F(1,46) = 35.12, p < .001, η_p^2 = .433) and cue-target relationship (F(1,46) = 5.23, p = .027, η_p^2 = .102), and interactions between group and reward (F(1,46) = 17.99, p < .001, η_p^2 = .281) - reflecting an accentuated reward effect in the LL group - and between cue-target relationship and reward (F(1,46) = 7.01, p=0.011, η_p^2 = .132). The only remaining effect in this analysis -the main effect of group - was not significant (F(1,46) = 2.25, p = .140, η_p^2 = .047).

Three-factor analysis of accuracy provided no substantive insight beyond that provided by analysis of the groups separately. A significant main effect of reward emerged, with participants more accurate when the cue appeared on the low reward rectangle (93.6% vs 87.5%, F(1,46) = 21.66, p < .001, η_p^2 = .320). This appears to reflect a speed accuracy tradeoff; participants were generally quicker when the cue appeared on the high reward rectangle, but slightly less accurate, regardless of the ultimate target location. A significant interaction of group and reward also emerged (F(1,46) = 6.48, p = .014, η_p^2 = .123): when the cue identified the high reward rectangle, this decreased accuracy more substantively in the LL group (93.9% vs 84.5%) than in the LE group (93.2% vs 90.5%). This appears to reflect an impact of the continuing strategic importance of reward association for the LL group. No other effects were significant (group: 89.2% vs 91.2%, F(1,46) = 1.12, p = .296, η_p^2 = .024; cue-target relationship: 91.2% vs 89.8%, F(1,46) = 1.48, p = .230, η_p^2 = .031; group X cue-target relationship: F(1,46) = 1.13, p = .294, η_p^2 = .024; all other effects Fs < 1).



Figure 5. Results from Experiment 4. Reaction times from the first phase of the experiment are collapsed across LL and LE groups and illustrated in panel **A**. The object prioritization effect emerges both when the target appears on the high-reward associated object and the low-reward associated object; analysis identifies no reliable difference in these effects (see body of paper). Reaction times from the second phase of the experiment for the LL group are illustrated in panel **B**. The object prioritization effect appears to emerge more strongly when the target appears on the low-reward associated object. Reaction times from the second phase of the experiment for the LE group are illustrated in panel **C**. Again, the object prioritization effect appears to emerge more strongly when the target appears on the low-reward associated object.

Discussion

In Experiment 3, reward prioritization and object prioritization interacted in the extinction phase of the experiment, raising the possibility that this interaction was caused by manipulation of reward feedback. This idea predicts that the interaction should emerge in the LE group of Experiment 4, but not in the LL group. In contrast, results show an interactive relationship in both groups. In RT, this effect is significant in the LE group and approaches significance in the LL group. These effects do not statistically differ between the groups, and the effect on raw RT is in fact nominally larger in the LL group than it is in the LE group. This pattern suggests that the interaction of reward prioritization and object prioritization emerges not as a product of the manipulation of reward feedback, but as a function of task practice.

As in earlier experiments, the interaction of reward prioritization and object prioritization in Experiment 4 is driven in particular by a slowing of response when the cue has identified a high-reward object, but the target ultimately appears on the low-reward object.

General Discussion

Selective attention is sensitive to the prior experience of reward, causing reward-associated stimuli to draw attention even under circumstances where this has no immediate benefit. Similarly, attention is sensitive to the definition of visual objects, prioritizing locations on a cued object over locations elsewhere, often when this is not useful. Here we investigate the relationship between these effects. We focussed on the possibility of an interaction between reward prioritization and object prioritization: is attentional engagement of visual objects stronger when the object has been linked to reward in prior experience?

Our experiments relied on the 2-rectangle paradigm (Egly et al., 1994), in which participants search for a target that appears at one of 4 locations. These locations fall on two task-irrelevant rectangles, such that each rectangle contains two possible target locations. When an exogenous cue identifies one location, participants are faster to detect the target at that location than when it appears elsewhere. More importantly, when a location is cued, but the target does not appear at this position, participants are faster to detect a target appearing on the same rectangle as the cue.

Our investigation began with a simple test of the impact of reward feedback on object prioritization in the 2-rectangle task. Existing experimental work from Shomstein and Johnson (2013) has suggested that the introduction of reward feedback to the 2-rectangle task causes object prioritization to disappear. We had the idea that this might reflect the influence of sequential contingencies on attentional deployment (eg. Hickey, Chelazzi, & Theewues, 2010a, 2010b), and Experiment 1 was designed as a replication of Shomstein and Johnson (2013, Experiment 2) in order that we might isolate intertrial effects in the results. Though we ultimately found no evidence of such intertrial contingencies, we did observe a robust effect of object prioritization. Experiment 1 thus re-opened a possibility that had been explicitly closed by Shomstein and Johnson (2013), namely that reward prioritization and object prioritization might coexist in visual cognition. This coexistence was corroborated and investigated in Experiments 2 through 4.

In critical conditions of Experiments 2 through 4, one rectangle in the display was characterized by a color that either validly predicted reward outcome (Experiment 2; learning stages of Experiments 3 and 4), or had predicted reward outcome in recent experience (extinction stage of Experiment 3; extinction stage of Experiment 4 for the LE group). Results from Experiment 2 suggested that when the reward association had strategic utility and predicted feedback magnitude, reward prioritization and object prioritization appeared as independent, additive effects on attentional selection. In Experiment 3, we tested if these effects would sustain when reward feedback was discontinued and the reward association entered into extinction. As in Experiment 2, reward prioritization and object prioritization emerged in the results of Experiment 3, but these effects now interacted: object prioritization emerged with greater strength when the target appeared on an object associated with low-magnitude reward, with RT particularly slowing when a high-reward objects was cued but the target appeared on a low-reward object.

A possible interpretation of results from Experiments 2 and 3 is that reward prioritization and object prioritization independently impacted performance when reward prioritization had strategic utility, but interacted when reward was discontinued, and this could have interesting implications for our understanding of these effects. That is, the emergence of additive effects on RT tentatively suggests an influence on different cognitive mechanisms, whereas interaction can be interpreted as reflecting an impact on the same processing stage (Sternberg, 1969). There was the possibility that a low-level, non-strategic impact of reward influenced the same cognitive stage as object prioritization, but a high-level, strategic impact of reward expressed elsewhere.

However, the design of Experiment 3 confounded the discontinuation of reward feedback with task practice. That is, the extinction phase necessarily came at the end of the

experiment, raising the possibility that it was not the manipulation of strategic utility that caused the interaction of effects, but task practice. This was exacerbated by the fact that Experiment 3 was substantially longer than Experiment 2. With this in mind, Experiment 4 compared two groups of participants. The LE group completed a task very similar to that employed in Experiment 3, with discontinuation of reward feedback in an extinction period. The LL group completed a task very similar to that employed in Experiment 3, with discontinuation of reward feedback in an extinction period. The LL group completed a task very similar to that employed in Experiment 2, with reward feedback sustaining throughout the experiment. Each group had the same amount of task experience, and this design provided the opportunity to identify whether it was the strategic utility of the reward association that impacted the relationship between object prioritization and reward prioritization, or task experience. Results suggest that it was task experience that was important, with the interaction of object prioritization and reward prioritization emerging in both LL and LE groups with roughly equal strength.

We approached the current study with the broad idea that there may be evolutionary utility to the gathering of information from locations on reward-associated visual objects, and therefore that reward prioritization and object prioritization might interact in visual cognition to benefit processing of stimuli appearing at these locations (Toates, 1990; Hickey, Chelazzi, & Theeuwes, 2010). Our results are consistent with this idea, but they do not actually demonstrate a benefit to the processing of target stimuli on the high-reward object. The interaction of reward prioritization and object prioritization identified in Experiments 3 and 4 appears rather to be driven by a slowing of RT when attention was drawn to a location on an object associated with high-magnitude reward, but participants had to subsequently reorient to a location on an object associated with low-magnitude reward (see also Zhao et al, 2020). That is, when the target appeared on a high-reward object, it did not much matter if that object had been cued, suggesting that locations on the high-reward object were prioritized regardless of cue location. However, when cue location and reward status combined, it became difficult for participants to orient attention elsewhere. This consistent prioritization of reward-associated objects in the environment appears to develop as a product of task practice. This is in line with the idea that it reflects incremental reward learning. The

experience of reward feedback in individual trials appears to build up, ultimately leading to prioritization of the reward-associated object of such degree that it can not be additionally influenced by the effect of object prioritization.

Participants in Experiments 3 and 4 show a lingering influence of reward association during extinction, when an ideal observer would not. At first blush, the insensitivity of the interaction of object prioritization and reward prioritization to the utility of the reward association is surprising. However, reward prioritization is well known in the broader literature to be relatively impervious to countermanding strategy (eg. Hickey, Chelazzi, & Theeuwes, 2010a, b; Hickey & van Zoest, 2012) and to extinguish slowly over the course of many trials (eg. Anderson et al., 2011; Stankevitch & Geng, 2015). In the extreme case, Anderson and Yantis (2013) found that a reward-associated object drew attention 9 months after training. This is in contrast to the broader motivational impact of reward prospect, which disappears quickly when the strategic utility of the reward association is removed. The motivational impact of reward is evident in the current results in the substantive main effect of reward in analysis of Experiments 3 and 4; when the utility of the reward signal is discontinued, the presence of the reward-associated object at the beginning of each trial no longer triggers the same broad arousal and investment of resources. However, the reward-associated object continues to be attentionally prioritized.

The current results add to a developing literature investigating the impact of reinforcement learning on the deployment of visuospatial attention. This literature has been influenced by empirical and theoretical work on animal associative learning, and specifically by the idea that mesencephalic dopamine may play a core role in cognition through an impact on selective control (eg. Ikemoto & Panksepp, 1999; Redgrave, Prescott, & Gurney, 1999; Jeong et al., 2022). For example, the *incentive salience hypothesis* of Berridge and Robinson (1998) suggests that reward-related dopamine acts to prime the perceptual and attentive representation of reward-associated stimuli, making it more likely that animals will notice similar stimuli in the future and ensuring that information from reward-associated stimuli gains preferential access to downstream cognitive operations like decision making

and motor control. Pathological misattribution of incentive salience has been linked to a variety of clinical disorders, including eating disorders, depression, paranoia, obsessive compulsive behaviour, and - prominently - addiction (Robinson & Berridge, 1993; Olney, Warlow, Naffziger, & Berridge, 2018). In addiction, direct drug stimulation of dopaminergic circuitry is thought to drive the attribution of incentive salience to drug-related stimuli. As a result, drug-related paraphernalia become salient and attention-drawing, and, once noticed, trigger craving and drug-seeking behaviour (ie. the 'drug trigger' phenomenon; Dackis & O'Brien, 2005). As discussed in the Introduction, the study of human incentive salience to date has largely focussed on the association of reward to low-level visual features, like color (Hickey, Chelazzi, & Theeuwes, 2010; Anderson, Laurent, & Yantis, 2011), or to categories of real-world objects (Hickey & Peelen, 2015; Donohue, Hopf, Bartsch, Schoenfeld, Heinze, & Woldorff, 2016; Hickey, Acunzo, & Dell, 2023). The current results provide the basis for further investigation of the interaction of raw visual salience, reward, and object status in the core phenomenon of incentive salience, as well as in disorders reflecting the misattribution of incentive salience.

As noted in the introduction, recent theoretical interpretation of object prioritization has suggested the effect may be strategic in nature and contingent on goal-driven attentional control settings (Shomstein & Johnson, 2013; Lee & Shomstein, 2013; Taylor et al., 2016). For example, Taylor and colleagues (2016) have shown that object prioritization manifests when the object matches the top-down perceptual filters strategically adopted by participants, but does not emerge otherwise. Results from the current study do not directly speak to this issue, but do support the idea hat object prioritization can emerge when the environment offers other strategic opportunities for the optimization of performance (cf. Shomstein & Johnson, 2013, Lee & Shomstein, 2013). This could reflect the adoption of multiple concurrent strategic approaches to visual search (Wolfe, 2021).

In conclusion, we find that reward prioritization and object prioritization have concurrent influence on the deployment of selective attention. These effects appear to combine and interact such that attention is particularly engaged with locations on visual objects that have been associated with reward.

References

- Anderson, B. A. (2019). Neurobiology of value-driven attention. Current opinion in psychology, 29, 27-33. https://doi.org/10.1016/j.copsyc.2018.11.004
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. Proceedings of the National Academy of Sciences, 108(25), 10367–10371. https://doi.org/10.1073/pnas.1104047108
- Anderson, B. A., & Yantis, S. (2013). Persistence of value-driven attentional capture. Journal of Experimental Psychology: Human Perception and Performance, 39(1), 6.
- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443. https://doi.org/10.1016/j.tics.2012.06.010
- Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience?. Brain research reviews, 28(3), 309-369. https://doi.org/10.1016/S0165-0173(98)00019-8
- Chen, Z., & Cave, K. R. (2006). Reinstating object-based attention under positional certainty: The importance of subjective parsing. *Perception & Psychophysics*, 68(6), 992–1003. https://doi.org/10.3758/BF03193360
- Chen, Z., & Cave, K. R. (2008). Object-based attention with endogenous cuing and positional certainty. *Perception & Psychophysics*, 70(8), 1435–1443. https://doi.org/10.3758/PP.70.8.1435
- Chen, Z. (2012). Object-based attention: A tutorial review. Attention, *Perception, & Psychophysics*, 74, 784-802.
- Dackis, C., & O'Brien, C. (2005). Neurobiology of addiction: treatment and public policy ramifications. Nature neuroscience, 8(11), 1431-1436. https://doi.org/10.1038/nn1105-1431

- Donohue, S. E., Hopf, J. M., Bartsch, M. V., Schoenfeld, M. A., Heinze, H. J., & Woldorff, M. G. (2016). The rapid capture of attention by rewarded objects. Journal of cognitive neuroscience, 28(4), 529-541. https://doi.org/10.1162/jocn_a_00917
- Driver, J., Davis, G., Russell, C., Turatto, M., & Freeman, E. (2001). Segmentation, attention and phenomenal visual objects. *Cognition*, *80*(1), 61–95. https://doi.org/10.1016/S0010-0277(00)00151-7
- Drummond, L., & Shomstein, S. (2010). Object-based attention: Shifting or uncertainty?. Attention, Perception, & Psychophysics, 72, 1743-1755.
- Erdfelder, E., Faul, F., & Buchner, A. (1996). GPOWER: A general power analysis program. Behavior research methods, instruments, & computers, 28, 1-11.
- Egeth, H. E., & Yantis, S. (1997). VISUAL ATTENTION: Control, Representation, and Time Course. *Annual Review of Psychology*, *48*(1), 269–297. https://doi.org/10.1146/annurev.psych.48.1.269
- Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, *123*(2), 161–177. https://doi.org/10.1037/0096-3445.123.2.161
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods*, 39(2), 175-191.
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human perception and performance*, 18(4), 1030.
- Hickey, C., Acunzo, D., & Dell, J. (2023). Suppressive control of incentive salience in real-world human vision. *Journal of Neuroscience*, *43*(37), 6415-6429.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010a). Reward Changes Salience in Human Vision via the Anterior Cingulate. *Journal of Neuroscience*, *30*(33), 11096–11103. https://doi.org/10.1523/JNEUROSCI.1026-10.2010

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*(11), e14087. https://doi.org/10.1371/journal.pone.0014087

- Hickey, C., Kaiser, D., & Peelen, M. V. (2015). Reward guides attention to object categories in real-world scenes. *Journal of Experimental Psychology: General*, 144(2), 264–273. https://doi.org/10.1037/a0038627
- Hickey, C., & Peelen, M. V. (2015). Neural mechanisms of incentive salience in naturalistic human vision. Neuron, 85(3), 512-518.

http://dx.doi.org/10.1016/j.neuron.2014.12.049

- Hickey, C., & van Zoest, W. (2012). Reward creates oculomotor salience. *Current Biology*, 22(7), R219–R220. https://doi.org/10.1016/j.cub.2012.02.007
- Ikemoto, S., & Panksepp, J. (1999). The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking. Brain Research Reviews, 31(1), 6-41. https://doi.org/10.1016/S0165-0173(99)00023-5
- Jeong, H., Taylor, A., Floeder, J. R., Lohmann, M., Mihalas, S., Wu, B., ... & Namboodiri, V.
 M. K. (2022). Mesolimbic dopamine release conveys causal associations. Science, 378(6626), eabq6740. https://www.science.org/doi/10.1126/science.abq6740
- Lange, K., Kühn, S., & Filevich, E. (2015). "Just Another Tool for Online Studies" (JATOS):
 An Easy Solution for Setup and Management of Web Servers Supporting Online
 Studies. *PLOS ONE*, *10*(6), e0130834. https://doi.org/10.1371/journal.pone.0130834
- Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: An integrative review. *Psychological bulletin*, 142(10), 1111.
- Lee, J., & Shomstein, S. (2013). The Differential Effects of Reward on Space- and Object-Based Attentional Allocation. *Journal of Neuroscience*, *33*(26), 10625–10633. https://doi.org/10.1523/JNEUROSCI.5575-12.2013
- Lou, H., Lorist, M., & Pilz, K. S. (2021). *Effects of cue validity on the temporal dynamics of attentional selection*. PsyArXiv. https://doi.org/10.31234/osf.io/sp38b

Macquistan, A. D. (1997). Object-based allocation of visual attention in response to

exogenous, but not endogenous, spatial precues. *Psychonomic Bulletin & Review*, 4(4), 512–515. https://doi.org/10.3758/BF03214341

- Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, 44(2), 314–324. https://doi.org/10.3758/s13428-011-0168-7
- O'Brien, F., & Cousineau, D. (2014). Representing Error bars in within-subject designs in typical software packages. *The Quantitative Methods for Psychology*, *10*(1), 56–67. https://doi.org/10.20982/tqmp.10.1.p056
- Olney, J. J., Warlow, S. M., Naffziger, E. E., & Berridge, K. C. (2018). Current perspectives on incentive salience and applications to clinical disorders. Current opinion in behavioral sciences, 22, 59-69. https://doi.org/10.1016/j.cobeha.2018.01.007
- Peters, B., & Kriegeskorte, N. (2021). Capturing the objects of vision with neural networks. Nature human behaviour, 5(9), 127-1144. https://doi.org/10.1038/s41562-021-01194-6
- Pomerleau, V. J., Fortier-Gauthier, U., Corriveau, I., Dell'Acqua, R., & Jolicœur, P. (2014).
 Colour-specific differences in attentional deployment for equiluminant pop-out colours: Evidence from lateralised potentials. *International Journal of Psychophysiology*, *91*(3), 194–205. https://doi.org/10.1016/j.ijpsycho.2013.10.016
- *R: The R Project for Statistical Computing*. (n.d.). Retrieved 29 June 2022, from https://www.r-project.org/
- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem?. Neuroscience, 89(4), 1009-1023. http://dx.doi.org/10.1016/S0306-4522(98)00319-4
- Robinson, T. E., & Berridge, K. C. (1993). The neural basis of drug craving: an incentive-sensitization theory of addiction. Brain research reviews, 18(3), 247-291. https://doi.org/10.1016/0165-0173(93)90013-P
- Rousselet, G. A., & Wilcox, R. R. (2019). *Reaction times and other skewed distributions: Problems with the mean and the median* (p. 383935). bioRxiv.

https://doi.org/10.1101/383935

- Shomstein, S. (2012). Object-based attention: Strategy versus automaticity. *WIREs Cognitive Science*, *3*(2), 163–169. https://doi.org/10.1002/wcs.1162
- Shomstein, S., & Johnson, J. (2013). Shaping Attention With Reward: Effects of Reward on Space- and Object-Based Selection. *Psychological Science*, *24*(12), 2369–2378. https://doi.org/10.1177/0956797613490743
- Shomstein, S., & Yantis, S. (2004). Configural and contextual prioritization in object-based attention. *Psychonomic Bulletin & Review*, *11*(2), 247–253. https://doi.org/10.3758/BF03196566
- Stankevich, B. A., & Geng, J. J. (2015). The modulation of reward priority by top-down knowledge. *Visual Cognition*, 23(1-2), 206-228.
- Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. Acta psychologica, 30, 276-315. https://doi.org/10.1016/0001-6918(69)90055-9
- Taylor, J. E. T., Rajsic, J., & Pratt, J. (2016). Object-based selection is contingent on attentional control settings. *Attention, Perception, & Psychophysics*, 78, 988-995.

Toates, F. (1990). *Motivational Systems*. New York, USA: Cambridge University Press

- Townsend, J.T., & Ashby, F.G. (1978). Methods of modeling capacity in simple processing systems. In J. Castellan & F. Restle (Eds.), Cognitive theory. Vol. 3. (pp. 200-239). Hillsdale, N.J.: Erlbaum.
- Townsend, J.T., & Ashby, F.G. (1983). Stochastic modeling of elementary psychological processes. *Cambridge: Cambridge University Press*
- Wannig, A., Stanisor, L., & Roelfsema, P. R. (2011). Automatic spread of attentional response modulation along Gestalt criteria in primary visual cortex. *Nature Neuroscience*, *14*(10), 1243–1244. https://doi.org/10.1038/nn.2910
- Wickham, H. (2016). ggplot2: Elegant Graphics for Data Analysis. Springer-Verlag New York. https://doi.org/10.1002/wics.147
- Wolfe, J. M. (2021). Guided Search 6.0: An updated model of visual search. Psychonomic Bulletin & Review, 28(4), 1060-1092.

- Yeari, M., & Goldsmith, M. (2010). Is object-based attention mandatory? Strategic control over mode of attention. *Journal of Experimental Psychology: Human Perception and Performance*, 36(3), 565–579. https://doi.org/10.1037/a0016897
- Zhao, J., Song, F., Zhou, S., Hu, S., Liu, D., Wang, Y., & Kong, F. (2020). The impact of monetary stimuli on object-based attention. *British Journal of Psychology*, *111*(3), 460–472. https://doi.org/10.1111/bjop.12418