

Rewarding context accelerates implicit guidance in visual search

Yuan-Chi Tseng · Alejandro Lleras

Published online: 30 November 2012
© Psychonomic Society, Inc. 2012

Abstract It is well known that observers can implicitly learn the spatial context of complex visual searches, such that future searches through repeated contexts are completed faster than those through novel contexts, even though observers remain at chance at discriminating repeated from new contexts. This *contextual-cueing* effect arises quickly (within less than five exposures) and asymptotes within 30 exposures to repeated contexts. In spite of being a robust effect (its magnitude is over 100 ms at the asymptotic level), the effect is implicit: Participants are usually at chance at discriminating old from new contexts at the end of an experiment, in spite of having seen each repeated context more than 30 times throughout a 50-min experiment. Here, we demonstrate that the speed at which the contextual-cueing effect arises can be modulated by external rewards associated with the search contexts (not with the performance itself). Following each visual search trial (and irrespective of a participant's search speed on the trial), we provided a reward, a penalty, or no feedback to the participant. Crucially, the type of feedback obtained was associated with the specific contexts, such that some repeated contexts were always associated with reward, and others were always associated with penalties. Implicit learning occurred fastest for contexts associated with positive feedback, though penalizing contexts also showed a learning benefit. Consistent feedback also produced faster learning than did variable feedback,

though *unexpected penalties* produced the largest immediate effects on search performance.

Keywords Rewards · Contextual cuing · Implicit learning · Visual search · Guidance

Human vision is very sensitive to regularities in the environment (Purves, Wojtach, & Lotto, 2011; Simoncelli, 2003). Furthermore, we use learned regularities about the environment to guide our behavior, a well-known example being the phenomenon known as *contextual cueing* (Chun & Jiang, 1998). Contextual cueing refers to the finding that searching for a target occurs faster if the target appears within the same spatial context and at the same location as on previous search instances (old display), as compared to when the target appears within a completely novel spatial context (new display).

Contextual cueing emerges quite quickly: On average, after as few as five repetitions, search reaction times become significantly shorter on old-display than on new-display trials (Chun & Jiang, 1998; Tseng & Li, 2004). In addition, performance in contextual-cueing experiments also asymptotes after about 30 exposures to old displays. That is, further repetitions of old displays provide no additional search time advantage. Figure 1a illustrates the emergence and asymptoting characteristics of this phenomenon (Brockmole & Henderson, 2006; Chun, 2000; Chun & Jiang, 1998, 2003; Jiang & Wagner, 2004; Schankin, Hagemann, & Schubö, 2011).

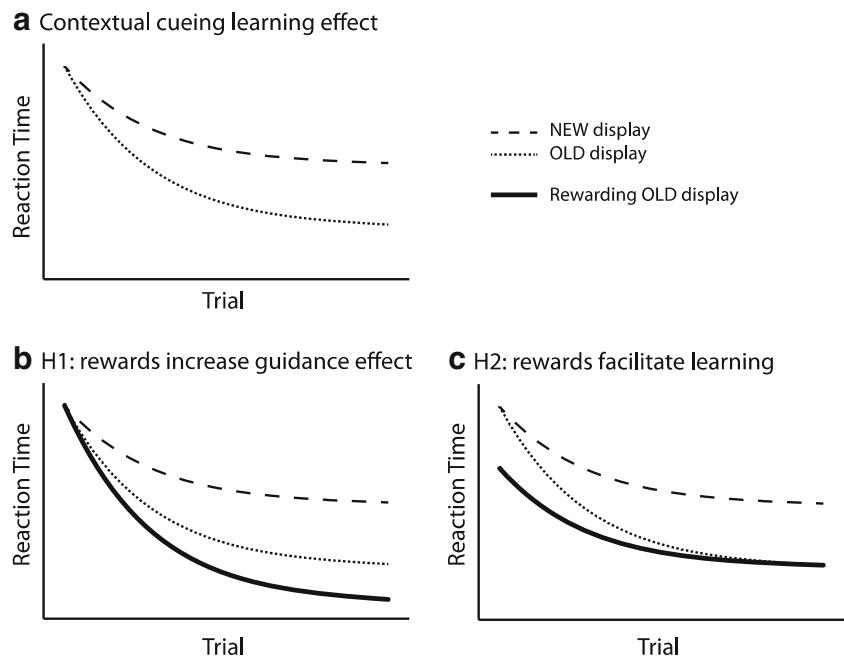
In spite of its robustness and large magnitude, the contextual-cueing effect is an entirely implicit effect (Chun & Jiang, 1998, 2003): Even though participants find targets faster in old than in new contexts, and even though they may see the same set of repeated contexts as many as 30 times in a given experimental session, participants remain at chance at discriminating old from new spatial contexts in an explicit old-versus-new discrimination task at the end of the experiment (Chun & Jiang, 1998, 2003; Tseng & Li, 2004). This finding has been widely accepted as reflecting a case in

Y.-C. Tseng · A. Lleras
Department of Psychology and Human
Factors Division, Beckman Institute,
University of Illinois, Urbana, IL 61801, USA

Y.-C. Tseng
Smith-Kettlewell Eye Research Institute,
San Francisco, CA 94115, USA

Y.-C. Tseng (✉)
Department of Industrial Design, National
Cheng Kung University, 1 University Road,
Tainan 701, Taiwan
e-mail: yctsens@mail.ncku.edu.tw

Fig. 1 **a** Emergence and asymptoting characteristics of the contextual-cueing learning effect (modified from Chun & Jiang, 2003). **b** Hypothesis that associated reward modulates the strength of the contextual guidance effect. **c** Hypothesis that associated reward accelerates the acquisition of a given context



which implicit learning affects the attentional system: The learned implicit spatial context guides attention toward the target location (Chun & Jiang, 1998, 1999; Chun & Nakayama, 2000; van Asselen & Castelo-Branco, 2009). This finding has been supported by eye movement studies (Manginelli & Pollmann, 2009; Peterson & Kramer, 2001; Tseng & Li, 2004; van Asselen, Sampaio, Pina, & Castelo-Branco, 2011) and electrophysiological studies (Olson, Chun, & Allison, 2001; Schankin & Schubö, 2009). Furthermore, the contextual-cueing effect is impenetrable to awareness, and knowledge of the repetition of displays in an experiment changes little with regard to the observed effects. For example, in Chun and Jiang (2003) participants were informed that half of the display configurations would repeat throughout the experiment and that each time a configuration repeated, the target would be at the same location. Awareness of the repetition of displays did not improve performance, and participants still remained at chance at discriminating old from new contexts at the end of the experimental session.

Although contextual cueing is implicit and cognitively impenetrable, one might be able to modulate this effect through motivational manipulations. Here we tested one such motivational manipulation that has successfully produced strong modulations of the effect: Implementing a reward schedule associated with the spatial contexts. A number of previous studies have demonstrated that motivational rewards can facilitate selective attention and visual search (e.g., Hickey, Chelazzi, & Theeuwes, 2010a, 2010b, 2011; Kiss, Driver, & Eimer, 2009; Della Libera & Chelazzi, 2006, 2009; Navalpakkam, Koch, & Perona, 2009).

For example, Kiss et al. (2009) found that reward-associated targets elicit earlier and larger N2pc components

in a singleton-search task, which has been taken as evidence that the deployment of attention is biased toward reward-associated targets. Similarly, Raymond and O'Brien (2009) found that reward-associated images are easier to detect in RSVP streams and are less vulnerable to the attentional blink (Raymond & O'Brien, 2009). Furthermore, in a series of studies looking at the effects of reward on priming during visual search, Hickey and colleagues (Hickey et al., 2010a, 2010b, 2011) demonstrated that attention is more likely to be guided toward high-reward-associated objects; that is, objects that contain visual features that have been consistently associated with high-magnitude rewards draw attention to themselves more readily than do those containing visual features associated with low-magnitude rewards. Also, Kristjánsson, Sigurjónsdóttir, and Driver (2010) found that rewards can modulate the target-color repetition effect known as *priming of pop-out* (i.e., in a color pop-out task, repeating the target color across successive trials speeds up search performance). The authors found the magnitude of rewards can modulate the strength of this intertrial priming, with larger color priming being observed in conditions with larger rewards. And furthermore, Hickey, Chelazzi, and Theeuwes (2010b) found that the magnitude of this modulation of reward on priming was predicted very well by personality characteristics: Individuals who had large effects of rewards on color priming also scored high on the Behavioral Activation System scale (Carver & White, 1994), a personality test that is thought to reflect the extent to which an individual engages in reward-driven behaviors. In other words, a seemingly automatic behavior (color priming) is actually modulated by situational motivational factors (relative rewards within the experiment), as well as by more general personality traits.

In Hickey and colleagues' experiments (Hickey et al., 2010a, 2010b) the reward-associated feature was color, but target status was determined within the shape dimension: The targets were shape singletons. These findings suggest that rewards modulate the impact that irrelevant features can have in a search task by presumably changing the visual representation (saliency) of the reward-associated feature, and therefore facilitate visual search. Here, we extended this logic to contextual cueing by modulating the reward associated with each search context. If searching through a set of old contexts reliably triggered a reward, more so than did searching through a second set of old contexts, one would expect that the guidance observed from the former contexts would be stronger than the guidance observed from contexts that did not trigger a reward, or that even induced a penalty (e.g., Fig. 1b). Alternatively, it could be that rewards do not change the strength of guidance, but rather alter the speed of acquisition of a given context. That is, rewarding contexts (i.e., reward-associated contexts) may be learned faster (because of reinforcement) than nonrewarding contexts, yet, asymptotically, all contexts may end up producing equivalent levels of guidance (once learned), irrespective of the level of reward obtained during the learning phase (e.g., Fig. 1c).

In sum, we were interested in exploring whether rewarding spatial contexts would alter the implicit-learning process that gives rise to the contextual-cueing effect: If participants are unaware that they are searching through an old context, yet they receive a reward at the end of the trial, will that reward change the manner in which participants implicitly learn the old context? To answer this question, we looked at the time of occurrence of the contextual-cueing effect as well as at the degrees of guidance across different manipulations of the rewards (and penalties) associated with repeated and novel contexts.

General method

Participants

Undergraduate students from the University of Illinois at Urbana Champaign participated for course credit. They were 18–22 years old, had normal or corrected-to-normal vision, and were naïve as to the purpose of the experiment. Informed consent was obtained from all participants before the experiments were conducted. A total of 32 participants were recruited for Experiment 1, and 20 participants were recruited for Experiment 2.

Contextual configuration and stimuli

We used the *spatial contextual-cueing* paradigm that was first introduced by Chun and Jiang (1998), in which participants search for a target T among distractor Ls and report

the T's orientation. The key manipulation is spatial context: Unbeknownst to participants, 12 of the contexts are repeated throughout the experiment (12 distinct associations between a specific layout of items in the display and a corresponding target location). These contexts are intermixed with new contexts, which are randomly generated, never repeat, and provide no predictive information regarding the target location. The contextual-cueing effect is the response time (RT) advantage for finding targets faster on repeated-context as compared to new-context trials.

Here, we assigned a level of outcome (reward [R], penalty [P], or no outcome [N]) to every context in the experiment. The outcome was in the form of points, which were important to the participants because they had been told that the experiment would end once they had reached either 1,800 (Exp. 1) or 3,600 (Exp. 2) points. The participants saw the points after completing each trial, and the points were not associated with their search speeds. Experiments 1 and 2 differed in the manners in which outcomes were assigned to different contexts. To prevent participants from arbitrarily responding throughout the experiment, errors in target identification were penalized by taking away 25 points from the participants' total number of points. This incentive worked well: Most participants had very high accuracy rates, although in each experiment some participants had to be discarded because of overly low accuracy.

The stimuli were presented in an invisible 8×6 (= 48-cell) grid (approximately $53.2^\circ \times 39.9^\circ$). They were randomly jittered to prevent collinearities with other stimuli. The background color of the search display was gray, and the search stimuli were colored blue, green, orange, or red. The stimuli were about $3.1^\circ \times 3.1^\circ$ in size (see Fig. 2a).

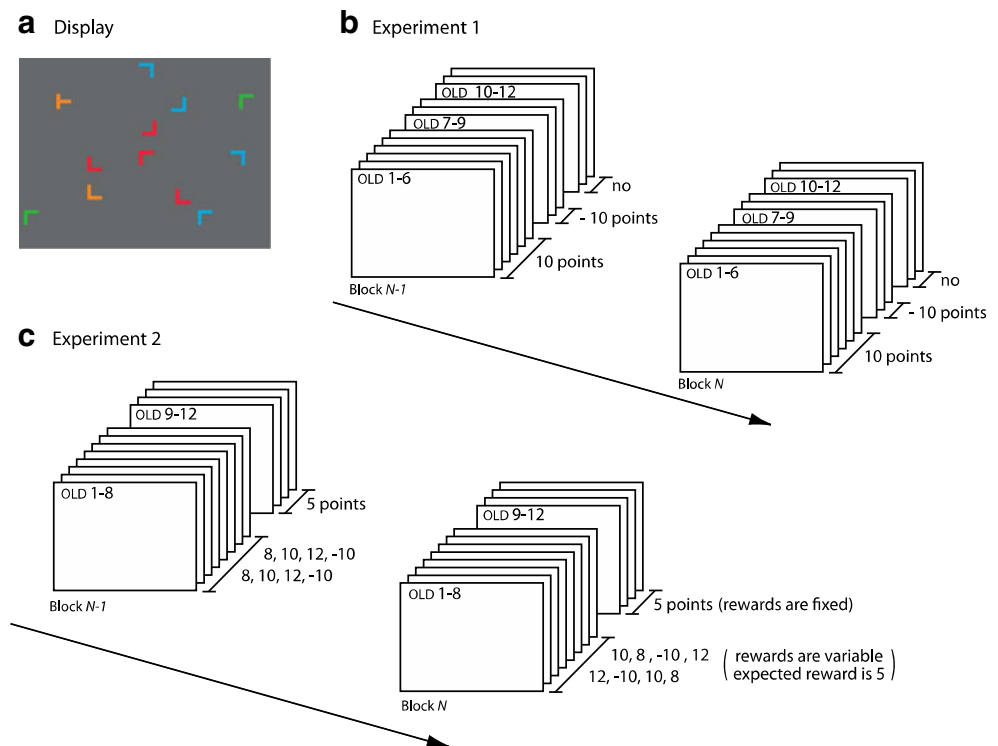
The target was a T rotated 90° to the right or left. Participants were to report the orientation of the T. The distractors were L stimuli presented randomly in one of four orientations (0° , 90° , 180° , and 270°). A target was always present among 11 distractors in a search display. The target orientation was chosen randomly on each trial so that there would be no correlation between a given context and the identity of the target in that context.

Each block consisted of 24 trials (12 repeated, old contexts and 12 new contexts). The locations of the targets in these 24 displays did not overlap and were arranged in an equiprobable fashion among the 48 cells in the display grid, between the old and new conditions.

Trial sequence and procedure

Each trial started with a fixation point that lasted 700 ms, followed by a visual search display presented until a response was made or until 4,000 ms had elapsed. A feedback display immediately followed, showing two numbers: the outcome points for the trial and the running total of all points

Fig. 2 **a** Schematic search displays. **b** Context–outcome assignments in Experiment 1: Consistently, six of the contexts were associated with 10 points, three with –10 points, and the other three with no reward outcome. **c** Context–outcome assignments in Experiment 2: Eight of the contexts were associated with variable rewards (indicated by two sets of 8, 10, 12, or –10 points) from one block to another, and four contexts were consistently associated with 5-point rewards. Note that 12 new displays were also presented in each block, with reward distributions similar to those in each experiment. The 24 search displays were mixed within each block



earned until that time. If the response was incorrect, an error message was displayed for 1,000 ms. After 500 ms, the next trial started. At the end of the experiment, the participants were shown a screen containing two search displays: a randomly chosen old display and a newly generated new display. They were asked to judge which of the two displays was more familiar. The recognition test (forced choice) is a simple way of testing for the implicit nature of memory (Frensch & Rüniger, 2003; Tseng & Li, 2004).

Data analysis

We pooled the RTs into six five-block epochs (Chun & Jiang, 1998). We computed mean RTs for the three types of repeated contexts (R-old, P-old, and N-old) in each of the six epochs. The RTs on new-context trials were collapsed into a single mean, irrespective of their level of reward, since the reward occurred after the response and new displays were never repeated. We only analyzed data from participants who averaged at least 90 % correct performance in the task, to minimize the effects of errors on performance and to make sure that all of our participants were at a similar point in the speed–accuracy trade-off function. This led us to discard seven of the 32 participants in Experiment 1 and two of the 20 participants in Experiment 2. This somewhat strict criterion ensured that we obtained similar numbers of observations in all search conditions and across participants. For example, in Experiment 1, 16 of the participants had fewer than one error per data point (i.e., per cell in the design), and the average number of errors

per data point was 0.88. Not a single participant had more than three errors in any cell of our design.

Given the number of points required to reach the goal in each experiment, all participants completed at least six epochs (30 blocks of trials). Because the participants made relatively few errors, after Epoch 6 they required only a relatively small number of additional trials to complete the experiment. Epoch 7 was never completed by any of the participants included in the analysis, and the data from these “catch-up” trials were not analyzed. It should be noted that all of the participants who performed the experiment for the full 50-min period allotted by the Institutional Review Board were stopped at the 50-min mark, given credit for participation, and allowed to leave. This was the case for two participants in Experiment 1, who both ended up being excluded from the analysis (one of the participants had 231 errors in the first 720 trials, and the other, 364 errors). The other five excluded participants completed the experiment in less than 50 min but were still discarded because they did not meet our inclusion criterion.

Experiment 1: Effects of rewarding contexts on the emergence of contextual cueing

Outcome manipulation

Among the 12 repeated contexts displays, six of the contexts gave participants 10 points (rewarding contexts, R-old),

three gave them -10 points (penalizing contexts, P-old), and the final three contexts did not give points nor was any outcome shown to the participants¹ (no-outcome contexts, N-old). The outcome of each repeated context was consistent throughout the experiment (Fig. 2b). Therefore, on a given block, participants could gain as much as 30 points from old displays. The 12 new displays had the same reward frequency distribution: On each block, six of the new displays produced 10 points, three took 10 points away, and three yielded no outcome; therefore, participants could earn 30 points from new displays as well. In total, the participants could earn 60 points per block. Importantly, the reward levels were constant for each old display, and for both new and old displays, the reward levels were entirely independent of RTs. The experiment ended once participants had gathered 1,800 points, forcing them to complete at least 30 blocks (720 trials). Participants were given a 24-trial practice block with only new displays, and they were not informed of the outcome manipulation during practice.

Results

Figure 3a illustrates correct mean RTs as a function of epoch for each search condition. A two-way repeated measures ANOVA was performed with Search Condition and Epoch as factors. Search condition had a significant main effect, $F(3, 72) = 55.133, p < .001, \eta_p^2 = .697$: Pairwise comparisons showed that the RTs in all three old displays were faster than RTs on new trials (all $ps < .05$), replicating the contextual-cueing effect. Moreover, the average RTs on rewarding context trials were shorter than those on penalizing context and no-outcome context trials (all $ps < .01$). The RTs on penalizing context trials were not significantly faster than those on no-outcome context trials ($p = .227$). These results showed that a learned reward (here, a “points” reward) can improve search performance exclusively for contexts that are implicitly associated with a positive valence. We also verified that none of these patterns were due to differences in accuracy across the conditions: No reliable effects of search display on accuracy were observed, as can be seen in Fig. 3b.

Time of occurrence of guidance Epoch had a significant main effect on RTs, $F(5, 120) = 13.894, p < .001, \eta_p^2 = .367$, reflecting learning throughout the experiment. Crucially, we found a significant interaction between search condition and epoch, $F(15, 360) = 11.223, p < .001, \eta_p^2 = .319$, implying that learning varied by search condition. Planned paired-samples *t* tests were conducted to compare the mean RTs between each pair of search conditions at

¹ We ran a different experiment in which we used an “X” symbol as the feedback on no-outcome context trials. We obtained the same pattern of results as in the present experiment.

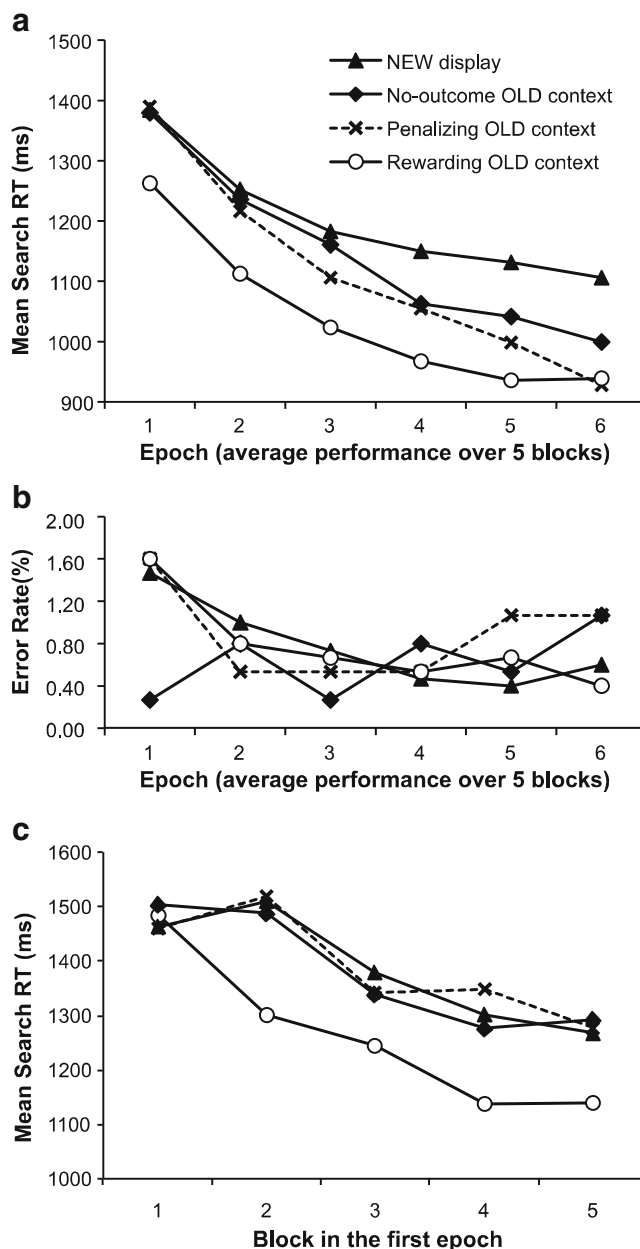


Fig. 3 Mean search performance in the new-display, no-outcome old-context, penalizing old-context, and rewarding old-context conditions in Experiment 1. **a** Mean correct search response times (RTs), plotted as a function of epoch (average of five blocks). **b** Mean error rates plotted as a function of epoch. **c** Mean correct search RTs plotted as a function of block in the first epoch. Contextual cueing is the RT difference between the new-display and each of the old-context conditions ($N = 25$)

each epoch. The contextual-cueing effects were computed separately for R-old, P-old, and N-old trials by subtracting their RTs from the RT of the new condition. Figure 3 illustrates how quickly participants learned the rewarding contexts, or more precisely, how quickly the rewarding contexts started to influence search behavior: A strong and reliable contextual-cueing effect for rewarding contexts emerged as

early as the first epoch and remained significant across the experiment (mean contextual-cueing effects: 122, 140, 159, 183, 196, and 167 ms in Epochs 1–6, respectively; all p s < .001). Contextual cueing on P-old context trials only became significant starting at the third epoch [77 ms in Epoch 3, $t(24) = 2.575$, $p = .017$; 95 ms in Epoch 4, $t(24) = 2.745$, $p = .011$; 133 ms in Epoch 5, $t(24) = 4.490$, $p < .001$; 178 ms in Epoch 6, $t(24) = 6.480$, $p < .001$], and contextual cueing on N-old context trials only emerged at the fourth epoch [87 ms in Epoch 4, $t(24) = 2.230$, $p = .035$; 90 ms in Epoch 5, $t(24) = 3.099$, $p = .005$; 107 ms in Epoch 6, $t(24) = 3.617$, $p = .001$].

Because we observed a significant contextual-cueing effect in the R-old condition as early as Epoch 1, we did a follow-up analysis, consistent with previous studies (Chun & Jiang, 1998; Jiang & Chun, 2001; Jiang & Wagner, 2004): We broke down the data and looked at them separately for each block of trials within Epoch 1 (see Fig. 3c). Significant contextual cueing was observed only for R-old contexts within this first epoch, as evidenced by faster RTs in Blocks 2–5: $t(24) = -4.347$, $p < .001$; $t(24) = -2.388$, $p = .013$; $t(24) = -2.992$, $p = .003$; and $t(24) = -2.378$, $p = .013$, respectively. Crucially, we found no difference in Block 1 RTs by reward condition, which ensures that all of the displays were initially equated in terms of search difficulty [new vs. R-old, $t(24) = -0.831$, $p = .207$; new vs. P-old, $t(24) = 0.359$, $p = .639$; new vs. N-old, $t(24) = -0.517$, $p = .305$; R-old vs. P-old, $t(24) = 0.252$, $p = .598$; R-old vs. N-old, $t(24) = -1.268$, $p = .108$; P-old vs. N-old, $t(24) = -0.623$, $p = .270$].

Asymptote of guidance The results of paired-samples t tests showed that rewarding contexts provided a greater degree of guidance than penalizing contexts in Epochs 1–5 [$t(24) = 3.273$, $p = .003$; $t(24) = 2.671$, $p = .013$; $t(24) = 3.12$, $p = .005$; $t(24) = 2.328$, $p = .029$; and $t(24) = 2.103$, $p = .046$, respectively]. However, the contextual-cueing effects among rewarding, penalizing, and no-outcome contexts were equal by the 6th epoch (all p s > .1). These results suggest that the overall degree of guidance of a spatial context is not increased by the reward manipulation after the context is fully learned, but rather that the speed at which the repeated context is learned (and begins to exert an effect on search behavior) is what changes. This difference in learning rates produces effects on performance that appear earlier in the experiment in the R-old context trials than on either P-old or N-old context trials.

Implicit guidance The results of our forced choice context recognition test (old vs. new discrimination) also showed that participants were no more familiar with the repeated displays than with new displays (10 participants selected an old display and 15 participants selected a new display, $p = .424$, binominal test). Furthermore, participants were no more familiar with reward- than with penalty-associated contexts, either (11 of

the participants selected a reward-associated context, and 14 participants selected a penalty-associated context, when they were forced to choose in a display showing a pair of a reward-associated context and penalty-associated context $p = .69$, binominal test). In verbal interviews, none of our participants reported any familiarity with the cued displays.

Discussion

The results of Experiment 1 showed that the outcome that we associate with repeated experiences with a display can strongly modulate our response to that display in the future, even when we are not aware of the learning that is taking place (i.e., the learning remains implicit). Experiment 2 was in part designed to replicate this finding. However, before proceeding, it is worth discussing one aspect of our results: the finding of a contextual-cueing condition as early as Block 2 in the rewarded-old display condition. While at first this result may be surprising, it is worth noting that a number of published reports on contextual cueing have presented similar results. For instance, in the original Chun and Jiang (1998) article, which reported the existence of contextual cueing for the first time, the authors found a significant cueing effect as early as Epoch 1 in their Experiment 4. This is a valid comparison, given that here we used the same task and stimuli that they had used in that investigation (a search task in which participants were asked to find an oriented T among 11 randomly oriented L stimuli). In their Experiment 4, the RTs to new and old displays were significantly different as early as Epoch 1 ($p = .005$ for set size 12 and $p = .01$ for set size 16). Chun and Jiang (1998) further analyzed the data as a function of block within Epoch 1 and found, as we did, that a single repetition of a search display can suffice to produce significant cueing effects: Contextual cueing was significant at both the 12- and 16-item set sizes as early as Block 2, $p = .008$ and .01, respectively. Lleras and Von Mühlhelen (2004) also found that the contextual-cueing effect can occur in the first epoch (Exp. 3). Other studies have also shown that significant contextual-cueing effects can be found with fewer than five repetitions (i.e., as early as the first epoch; see Jiang & Chun, 2001; Jiang & Wagner, 2004). Although Peterson and Kramer (2001) did not calculate statistics to directly investigate contextual cueing within Epoch 1, their data still showed a clear contextual-cueing effect as early as Epoch 1. In sum, our findings are not out of line with previous studies on contextual cueing, but simply replicate the finding that, under the right conditions, this effect can emerge as early as within the first epoch of an experiment.

Furthermore, it is worth noting that we only found this “quick” emergence of contextual cueing in one of our three display repetition conditions: Contextual cueing took longer to emerge with penalizing and no-outcome old displays.

Experiment 2: Expected versus obtained rewards

We propose that the arousal elicited by the feedback display alters the consolidation process of implicit memory of the spatial context. In Experiment 2, we investigated this possibility by manipulating reward expectations: If participants have come to expect that a positive reward will be associated with a given context (call it *context V*), even if that expectation is implicit, we predicted that unexpectedly presenting a penalty after a search through context *V* would create a state of high arousal that would further accelerate the learning of this context and increase its impact on performance on its next occurrence in the experiment. To do so, we created two types of contexts: those with consistent reward outcomes (same outcome on every trial), distinguished from other contexts with varying levels of reward outcomes. The outcome of the latter group was varied on a probabilistic basis, while maintaining the same overall expected value as the contexts with consistent reward outcomes. This manipulation allowed us to separately measure the effects due to the overall expected outcome when a repeated context was reencountered, as distinct from the more immediate consequences of having recently obtained a given level of reward associated with a repeated context. Most interestingly, it also allowed us to investigate the scenario in which a clear violation of expectations would occur: Participants would (implicitly) expect a positive reward from a task environment, and yet obtain a penalty. We hypothesized that contexts that produced unexpected penalties should have a higher impact on performance (in terms of larger contextual-cueing effects) than would contexts that produced outcomes consistent with expectations, because of the surprising and therefore arousing nature of such events.

Outcome manipulation

We manipulated two reward assignments: a consistent and a variable one. Four repeated contexts had consistent rewards (i.e., participants always received 5 points at the end of a trial with one of those contexts; C-old), and eight repeated contexts had variable rewards (i.e., participants received one of four possible outcomes at the end of a trial with one of those contexts; V-old). For these eight repeated contexts, in an arbitrary fashion on each block, two of the old displays were assigned 8 points, two were assigned 10 points, two were assigned 12 points, and two were assigned –10 points (Fig. 2c). Therefore, the expected value throughout the

experiment of a variable-reward context would be 5 points $[(8 + 8 + 10 + 10 + 12 + 12 - 10 - 10)/8 = 5]$. Importantly, the variable-reward manipulation ensured that the expected values of both consistent-reward and variable-reward conditions were the same. Thus, in the long run, participants earned as many points from consistent-reward contexts than from variable-reward contexts. The assignment of a given reward level to a given context in this condition was done anew on each block of trials through a pseudorandom permutation of the possible reward values (8, 8, 10, 10, 12, 12, –10, –10), with the only constraint being that the overall point assignments for the current set of variable-reward contexts had to be different from the point assignments on the previous block of trials. Finally, we set the 12 new displays on each block to have a similar reward distribution, so that participants could not identify new from old displays on the basis of the outcome of a given trial. Participants could earn $(12 \text{ old} + 12 \text{ new}) \times 5 = 120$ points per block. The experiment was completed when participants reached a total of 3,600 points, corresponding to 30 blocks of trials. Finally, to allow us a comparison to Experiment 1, it should be noted that across both experiments, participants experienced similar frequencies of penalizing trials (six out of 24 trials, including both the old and new conditions, were penalizing in Exp. 1, whereas four out of 24 produced penalties in Exp. 2). The value of the penalties was the same across the two experiments (loss of 10 points), although, admittedly, the impact of such a loss was less in Experiment 2 than in Experiment 1, insofar as it comprised a smaller percentage of the overall point goal (Exp. 2 ended after participants had garnered 3,600 points, but only 1,800 points were required in Exp. 1).

Results and discussion

Consistent versus variable rewards A two-way repeated measures ANOVA was performed with search condition (new, C-old, and V-old) and epoch (1–6) as within-subjects variables. Both search condition and epoch had significant main effects on RTs, $F(2, 34) = 15.322, p < .001, \eta_p^2 = .474$, and $F(5, 85) = 19.63, p < .001, \eta_p^2 = .536$, respectively, reflecting that RTs were shorter to old than to new displays and that RTs decreased as a function of epoch. The interaction between search condition and epoch was also significant, $F(10, 170) = 3.775, p < .001, \eta_p^2 = .182$, demonstrating that the emergence rates of contextual cueing differed across the different contexts.

Figure 4 illustrates that variable-reward contexts were learned more slowly than consistent-reward contexts: Whereas contextual-cueing effects were observed as early as Epoch 1 for consistent-reward contexts (and indeed, contextual cueing was present for all epochs; all $ps < .01$), this effect was initially absent in the variable-reward

condition: RTs in that condition were identical to RTs to the new displays in Epoch 1, $t(17) = -0.53$, $p = .603$. That said, contextual-cueing effects did not differ between the two types of rewarding contexts beginning in Epoch 2 (all p s > .05 in Epochs 2–6), reflecting that the magnitude of the overall contextual-cueing effect was likely driven by the expected reward of a context. In addition, when we looked separately at contextual cueing for each block within Epoch 1 (see Fig. 4c), significant contextual cueing was observed only for the consistently rewarding old contexts, as evidenced by faster RTs in Blocks 2–5, $t(17) = -1.816$, $p = .043$; $t(17) = -2.370$, $p = .015$; $t(17) = -3.358$, $p = .002$; and $t(17) = -2.041$, $p = .029$, respectively. Crucially, we found no difference in the Block 1 RTs by reward condition, which ensured that all displays were initially equated in terms of search difficulty [new vs. C-old, $t(17) = -0.226$, $p = .412$; new vs. V-old, $t(17) = -0.572$, $p = .287$; C-old vs. V-old, $t(17) = -1.504$, $p = .073$]. The fast emergence of contextual cueing in the consistent-reward condition provides a nice replication of the fast emergence (one trial) of contextual cueing that we observed in the reward condition of Experiment 1.

These results suggest that experiencing inconsistent rewards that include penalties—particularly, early in the learning process—can defer the learning of the contexts as compared to experiencing consistent levels of reward. Importantly, once learned, the overall magnitudes of the contextual-cueing effects were identical for the variable-reward and consistent-reward contexts, suggesting that the magnitude of the overall effect tracks the expected value of repeated contexts.

Expected versus obtained outcomes The *obtained* outcome after a trial with a given context might have a direct impact on the next search through that same context, above and beyond the effect of the *expected* outcome for that context on the current trial. That is, although participants may implicitly expect to obtain on average of 5 points on any given trial, the effect of this expected outcome might differ if, on the last search through the current context, an unexpected (and penalizing) outcome had been obtained. For ten of the 12 old contexts in each block, searching through the context normally led to positive reward points (e.g., 5, 8, 10, or 12 points), but occasionally (i.e., two of 12 trials) it would result in a penalty. Given this outcome distribution, we could ask the question, Does an unexpected negative outcome alter learning of the spatial context that it was paired with? To answer this question, search performance was analyzed exclusively on two types of variable-reward contexts: most recently rewarding contexts (+10 points, gray bars) and most recently penalizing contexts (−10 points, black bars). In this analysis, we ignored rewarding contexts for which the outcome was other than 10 points, in order to

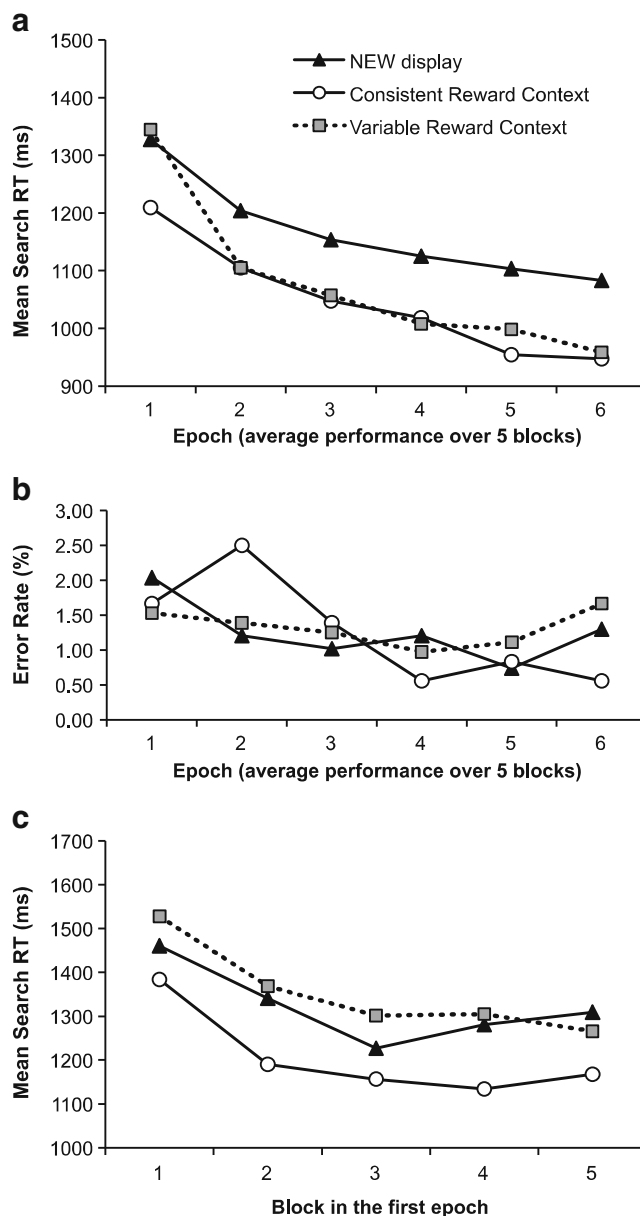


Fig. 4 Mean search performance in the new-display, consistent-reward old-context, and variable-reward old-context conditions in Experiment 2. **a** Mean correct search RTs plotted as a function of epoch. **b** Mean error rates plotted as a function of epoch. **c** Mean correct search RTs plotted as a function of block in the first epoch. Note that the consistent-reward contexts and the variable-reward contexts had the same expected reward, 5 points ($N = 18$)

equate the magnitudes of the outcomes obtained across these two conditions. On any given block of trials, there were only two occurrences of each of these types of trials (see Fig. 2c). We also compared performance on these trials to performance on a randomly chosen two of the consistent-reward contexts (out of the possible four). To obtain more stable data points, we also combined RTs from Epochs 1 and 2, 3 and 4, and 5 and 6.

Figure 5 shows the contextual-cueing effect as a function of the reward that was most recently associated with the current context. Our analyses showed that, consistent with the overall results observed in Fig. 4, the consistent-reward context produced consistent contextual-cueing effects at all epochs (Epochs 1–2, 135 ms, $p < .01$; Epochs 3–4, 83 ms, $p < .05$; Epochs 5–6, 119 ms, $p < .01$), and they also provide a replication of Experiment 1. More interesting are the results observed for the two variable-context conditions (see Fig. 5). Looking first at the variable-reward contexts (these are contexts that were last rewarded with +10 points but had an expected outcome of +5 points), contextual cueing took longer to emerge and only became significant at Epochs 5–6 (Epoch 5–6, 87 ms, $p < .05$). In contrast, contexts that were expected to deliver a positive reward but that, on their most recent occurrence, had surprisingly been associated with a point penalty were learned faster and showed consistent contextual-cueing effects throughout the experiment (Epochs 1–2, 110 ms, $p = .057$; Epochs 3–4, 157 ms, $p < .001$; Epochs 5–6, 115 ms, $p < .05$). Note that the contextual cueing effect in Epoch 1–2 is only marginally significant. Directly comparing these two conditions, we found that the contextual-cueing effect induced by the most recently penalty-associated contexts was greater than the one induced by recently rewarding contexts in Epochs 1–2 ($p < .05$) and 3–4 ($p < .05$), though by the end of the experiment, all three types of contexts produced identical cueing effects. In sum, our results show that experiencing an

unexpected point penalty (when a reward is expected) can immediately accelerate the learning of the associated context.

General discussion

The spatial contextual-cueing effect, which is driven by implicit memory of the consistent spatial relationship between the target and distractors, often develops very quickly and gradually reaches asymptotic levels after a few repetitions (Chun, 2000; Chun & Jiang, 1998, 2003; Jiang & Wagner, 2004; Schankin et al., 2011; Tseng & Li, 2004). Previous work has suggested that this implicit guidance of visual search is not influenced by explicit instructions (Chun & Jiang, 2003). Here, in two spatial contextual-cueing experiments, we manipulated the outcomes of the spatial contexts and found for the first time that the outcomes associated with spatial contexts can accelerate implicit learning of the repeating contexts. The different time courses of the contextual-cueing effect across outcome conditions that we found in Experiment 1 show that the learning process (how quickly people learn the repeated contexts and how quickly the repeated contexts come to influence search performance) is clearly mediated by valence, with rewarding contexts accelerating implicit learning of the repeated contexts (Fig. 3a and c).

Interestingly, it appeared that although penalizing contexts were learned more slowly than rewarding contexts, they were learned slightly faster than no-outcome contexts (Exp. 1). This finding is not necessarily surprising, as participants were probably aroused by the negative outcome on those trials. As a result, it is likely that the outcome-related arousal on those trials made the learning of those trials slightly faster than when no outcome was received at all. In Experiment 2, we further investigated the effect of rewards on contextual cueing by manipulating the expectation of reward associated with each context. When controlling for the overall expected outcome of a context (i.e., how many points on average a participant would receive from a given context), our results showed that contexts having consistent reward levels produced contextual-cueing effects earlier in the experiment than did those contexts with inconsistent reward levels (Fig. 4). Note that in consistent-reward display trials, the expectation of a specific reward (5 points) was never violated, and these displays were learned the quickest. Regarding variable-reward contexts, our results also suggested that *unexpected* penalties had a large impact on the learning rates of repeated displays: Unexpectedly penalized displays were responded to faster on their next occurrence than were displays that were *expectedly* rewarded.

On the other hand, our results also show that once the guiding effect of a context has reached an asymptotic level, the overall magnitude of the contextual-cueing effect can

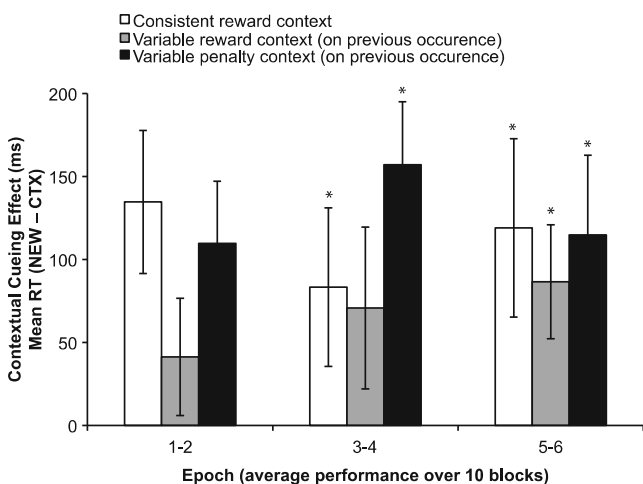


Fig. 5 Contextual-cueing effect (mean search RT of new displays – mean search RT of old displays) for consistent-reward contexts and for two types of variable-outcome contexts: (1) old contexts that were last associated with a reward of +10 points and (2) old contexts that were last associated with a penalty of –10 points. That is, the graph shows contextual-cueing effects measured on block n as a function of the outcome associated with a given repeated context on block $n - 1$. Due to the lower number of observations per condition (only two data points per block), the data were averaged across Epochs 1 and 2, Epochs 3 and 4, and Epochs 5 and 6. Asterisks indicate statistically significant contextual-cueing effects, and error bars represent SEMs

not be modulated further by its outcome. Experiment 1 showed that the asymptotic magnitude of the contextual-cueing effect was not affected by the reward or penalty associated with that context. Experiment 2 further suggested that the asymptotic magnitude of contextual guidance was not altered by whether participants had received an expected or unexpected outcome the last time that they had searched through that context.

Our results also suggests that the reward/penalty manipulation did not impact the level of awareness of the learning taking place in the task: The learning remained implicit. This is consistent with previous studies, which have repeatedly shown that contextual cueing is cognitively impenetrable (Chun & Jiang, 1998, 2003; Tseng & Li, 2004). However, contrary to suggestions in the literature (Chun & Jiang, 2003), we suggest that the learning of a context may be accelerated by motivational manipulations like rewards, and that it can also be slowed down when the outcome value is consistently negative (Exp. 1) or inconsistently positive (Exp. 2). This means that participants not only learn an association between a spatial context and the position of a target within that context, but also learn the relation between a spatial context and the subsequent level of reward that becomes associated with that context.

Previous studies have suggested that rewarding an act of selection can modulate the visual representation of the reward-associated feature, such that, for example, visual search is facilitated (e.g., Hickey et al., 2010a, 2010b, 2011; Kiss et al., 2009; Della Libera & Chelazzi, 2006, 2009; Navalpakkam et al., 2009). There were two critical differences between those studies and ours. First, we did not find that the reward manipulation (at the asymptotic level) fundamentally altered (neither increased nor decreased) the level of guidance provided by a repeated context. This is an important difference. A second difference between our study and prior research on the relation between attention and rewards was “what” was rewarded: Whereas in prior studies the reward was most often associated with a simple act of selection (e.g., select one of two shapes), here the “rewarded action” was much more complex: Our displays were not “pop-out” displays (i.e., participants could not immediately tell where the target was). Instead, they had to actively search each display to find the target, and what was rewarded, in a way, was the complete search behavior. Thus, perhaps a crucial difference between our results and previous studies was the complexity of the rewarded action. Finally, it should be noted that in our methodology, the reward was the same whether participants took 500 or 2,500 ms to find the target in a given context, and in this way, our rewards were clearly dissociated from specific behaviors.

Our results further suggest that the effects of reward on search in this paradigm were implicit, given that participants

could not discriminate old from new contexts. This is not entirely surprising, given the extensive literature showing that contextual cueing is itself an implicit effect. Yet the fact that rewards impacted the learning rate of this implicit effect seems to us interesting in several ways: The outcome of a trial was known only after the spatial context that was being learned was no longer visible, yet it appears that this post-facto information could accelerate the rate of learning of that context, probably by strengthening the consolidation of that spatial context information into memory. Furthermore, it is worth noting that, although participants might not be able to consciously distinguish an old from a new context, upon seeing a given context on the current trial, their visual systems had an implicit memory trace of the most recent outcome associated with that specific context. Access to this information occurred sufficiently quickly to affect behavior on the current trial (Exp. 2), particularly during the learning stages of the experiment. Finally, comparing the results of penalty trials across Experiments 1 and 2, we could isolate the effect of surprise on learning. Learning seemed most sensitive to *unexpected* penalties: Receiving an unexpected penalty after completing a search through a given context on block n produced faster search performance through the same context on block $n + 1$, when a *positive* outcome was expected (Exp. 2); when a penalty was expected, contextual cueing actually took longer to emerge for penalizing contexts (Exp. 1). Note that penalties occurred at similar rates in both experiments. We propose that this “surprise” effect on learning might be mediated by arousal: Receiving an expected penalty is likely to be less arousing than receiving a penalty when a *reward* is expected. Indeed, it has long been known that high states of arousal at encoding lead to greater long-term memory of the encoded events (for a classic review of the relationship between arousal and memory, see Eysenck, 1976). In other words, even though participants may have been penalized after searching through a given context, their consolidation of the context into memory was likely enhanced by the arousing unexpected outcome obtained shortly after the end of the trial. Because information that is encoded at high arousal levels is also more easily retrieved, it is then easy to see why unexpectedly penalizing contexts could go on to have large contextual-cueing effects the next time that the participant is faced with searching through such a context. Further research will be needed to determine whether the inverse result is also true: Will a context associated with an expected penalty be learned faster if it is unexpectedly associated with a reward? We believe that it will be. Finally, it is worth underscoring the implicit nature of these effects: Our participants were unable to discriminate old from new contexts and showed no overt sensitivity to either the valence of the reward associated with a given context (rewarding old contexts were not discriminated from either penalizing old or

new contexts). Yet, performance was clearly affected by the Context \times Outcome binding.

As for the larger literature on learning, although much work has focused on the effects of rewards and penalties on associative and instrumental learning, little is known about how positive and negative rewards alter implicit learning (see, e.g., Wächter, Lungu, Liu, Willingham, & Ashe, 2009). That said, some things are known. A growing literature is showing that positive and negative rewards (in general) are processed by different brain circuits (see the meta-analysis by Liu, Hairston, Schrier, & Fan, 2011), so in fact our results showing different effects for these two types of outcomes fit well with this neurological literature. Regarding implicit learning more specifically, Wächter et al. showed that implicit procedural learning was accelerated by rewards (though not by penalties). This result is well in line with the results of our Experiment 1, in which we only found a small difference in learning speed between penalizing and no-outcome old displays, yet a great improvement for rewarding old displays. Recently, Lam, Wächter, Globas, Karnath, and Luft (2012) studied implicit learning using a classification task. Their results (a) showed better learning of positively than of negatively rewarded information and, using fMRI, (b) showed increased activity in several brain regions (including the nucleus accumbens, the sensorimotor and premotor cortices, and the cingulum) during the feedback part of a trial for positive versus negative feedback. In particular, they showed that people remembered a larger number of implicitly learned associations when those associations were rewarded in a positive manner (a smiley face presented after the trial) rather than a negative manner (a frowny face presented after the trial). Overall, our results are in line with this emerging literature: Our participants showed faster implicit learning for positive (rewarding contexts) than for negative (penalizing contexts) rewards. The greater brain activation when processing positive feedback suggests that our results were probably due to better (or more in-depth) encoding of the rewarding than of the penalizing contexts.

One final point is worth discussing about our results: In both experiments, we consistently found evidence of one-trial learning that was specific to the reward condition. This type of one-trial learning effect is not unique to this study. For example, Cook and Fagot (2009) showed that animals could learn hundreds of picture–response pairs after a single exposure and, more critical to our finding, that this form of one-trial learning was more robust when the to-be remembered pair was followed by a reward rather than by a mild penalty (Cook & Fagot, 2009). In contrast, our findings do differ from previous studies on associate learning that have found one-trial learning to be boosted by negative outcomes (Armstrong, DeVito, & Cleland, 2006; Bardo & Bevins, 2000). One critical difference, though, between the latter

studies and ours (as well as Cook & Fagot, 2009) was the number of learned associations. Whereas only one association was learned in the latter studies, in ours, 12 complex associations were learned between a target location and the surrounding distractor context. In Cook and Fagot's study, even more associations were learned. Together, this evidence suggests that positive rewards can have a very fast accelerating effect on the speed of learning complex associations (even implicit ones)—to the point that a single trial may suffice for learning—and that this accelerating effect seems to be more pronounced than that of negative outcomes.

Conclusion

Our findings suggest that contextual learning is greatly sensitive to both the arousal and valence of context–outcome associations. Implicit learning, though cognitively impenetrable, is heavily modulated by external rewards associated with the search interaction. It is worth emphasizing that the learning process itself is what is modulated: Learning is accelerated, to the point that it might sometimes occur as quickly as after a single exposure. These results show that the outcomes of our interactions with the environment (both positive and negative) are remembered, and that these outcomes can exert strong implicit influences on our future behavior, not only by affecting the memory trace of the object or goal of our actions (which has been shown before), but also by affecting the memory representations of the contexts in which our actions take place.

Author note This research was partially supported by a Rachel C. Atkinson Postdoctoral Fellowship from Smith-Kettlewell Eye Research Institute and by Grant No. NSC 101-2218-E-006-021 from the National Science Council to Y.-C.T.

References

- Armstrong, C. M., DeVito, L. M., & Cleland, T. A. (2006). One-trial associative odor learning in neonatal mice. *Chemical Senses*, *31*, 343–349.
- Bardo, M., & Bevins, R. A. (2000). Conditioned place preference: What does it add to our preclinical understanding of drug reward? *Psychopharmacology*, *153*, 31–43.
- Brockmole, J. R., & Henderson, J. M. (2006). Using real-world scenes as contextual cues for search. *Visual Cognition*, *13*, 99–108. doi:10.1080/13506280500165188
- Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology*, *67*, 319–333. doi:10.1037/0022-3514.67.2.319
- Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, *4*, 170–178. doi:10.1016/S1364-6613(00)01476-5
- Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*, 28–71. doi:10.1006/cogp.1998.0681

- Chun, M. M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, *10*, 360–365. doi:10.1111/1467-9280.00168
- Chun, M. M., & Jiang, Y. (2003). Implicit, long-term spatial contextual memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 224–234. doi:10.1037/0278-7393.29.2.224
- Chun, M. M., & Nakayama, K. (2000). On the functional role of implicit visual memory for the adaptive deployment of attention across scenes. *Visual Cognition*, *7*, 65–81. doi:10.1080/135062800394685
- Cook, R., & Fagot, J. (2009). First trial rewards promote 1-trial learning and prolonged memory in pigeon and baboon. *Proceedings of the National Academy of Sciences*, *106*, 9530–9533. doi:10.1073/pnas.0903378106
- Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science*, *17*, 222–227. doi:10.1111/j.1467-9280.2006.01689.x
- Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*, 778–784. doi:10.1111/j.1467-9280.2009.02360.x
- Eysenck, M. W. (1976). Arousal, learning, and memory. *Psychological Bulletin*, *83*, 389–404. doi:10.1037/0033-2909.83.3.389
- Frensch, P. A., & Rüniger, D. (2003). Implicit learning. *Current Directions in Psychological Science*, *12*, 13–18. doi:10.1111/1467-8721.01213
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010a). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*, *30*, 11096–11103.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010b). Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *PLoS One*, *5*, e14087. doi:10.1371/journal.pone.0014087
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2011). Reward has a residual impact on target selection in visual search, but not on the suppression of distractors. *Visual Cognition*, *19*, 117–128. doi:10.1080/13506285.2010.503946
- Jiang, Y., & Chun, M. M. (2001). Selective attention modulates implicit learning. *Quarterly Journal of Experimental Psychology*, *54A*, 1105–1124. doi:10.1080/02724980042000516
- Jiang, Y., & Wagner, L. C. (2004). What is learned in spatial contextual cuing—configuration or individual locations? *Perception & Psychophysics*, *66*, 454–463. doi:10.1080/02724980042000516
- Kiss, M., Driver, J., & Eimer, M. (2009). Reward priority of visual target singletons modulates event-related potential signatures of attentional selection. *Psychological Science*, *20*, 245–251. doi:10.1111/j.1467-9280.2009.02281.x
- Kristjánsson, Á., Sigurjónsdóttir, Ó., & Driver, J. (2010). Fortune and reversals of fortune in visual search: Reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Attention, Perception, & Psychophysics*, *72*, 1229–1236. doi:10.3858/APP.72.5.1229
- Lam, J. M., Wächter, T., Globas, C., Karnath, H. O., & Luft, A. R. (2012). Predictive value and reward in implicit classification learning. *Human Brain Mapping*. doi:10.1002/hbm.21431
- Liu, X., Hairston, J., Schrier, M., & Fan, J. (2011). Common and distinct networks underlying reward valence and processing stages: A meta-analysis of functional neuroimaging studies. *Neuroscience and Biobehavioral Reviews*, *35*, 1219–1236. doi:10.1016/j.neubiorev.2010.12.012
- Lleras, A., & Von Mühlelen, A. (2004). Spatial context and top-down strategies in visual search. *Spatial Vision*, *17*, 465–482. doi:10.1163/1568568041920113
- Manginelli, A. A., & Pollmann, S. (2009). Misleading contextual cues: How do they affect visual search? *Psychological Research*, *73*, 212–221. doi:10.1007/s00426-008-0211-1
- Navalpakkam, V., Koch, C., & Perona, P. (2009). Homo economicus in visual search. *Journal of Vision*, *9*(1):31, 1–16. doi:10.1167/9.1.31
- Olson, I. R., Chun, M. M., & Allison, T. (2001). Contextual guidance of attention: Human intracranial event-related potential evidence for feedback modulation in anatomically early temporally late stages of visual processing. *Brain*, *124*, 1417–1425.
- Peterson, M. S., & Kramer, A. F. (2001). Attentional guidance of the eyes by contextual information and abrupt onsets. *Perception & Psychophysics*, *63*, 1239–1249. doi:10.3758/BF03194537
- Purves, D., Wojtach, W. T., & Lotto, R. B. (2011). Understanding vision in wholly empirical terms. *Proceedings of the National Academy of Sciences*, *108*, 15588–15595.
- Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science*, *20*, 981–988. doi:10.1111/j.1467-9280.2009.02391.x
- Schankin, A., Hagemann, D., & Schubö, A. (2011). Is contextual cueing more than the guidance of visual-spatial attention? *Biological Psychology*, *87*, 58–65. doi:10.1016/j.biopsycho.2011.02.003
- Schankin, A., & Schubö, A. (2009). Cognitive processes facilitated by contextual cueing: Evidence from event related brain potentials. *Psychophysiology*, *46*, 668–679. doi:10.1111/j.1469-8986.2009.00807.x
- Simoncelli, E. P. (2003). Vision and the statistics of the visual environment. *Current Opinion in Neurobiology*, *13*, 144–149. doi:10.1016/S0959-4388(03)00047-3
- Tseng, Y.-C., & Li, C.-S. R. (2004). Oculomotor correlates of context-guided learning in visual search. *Perception & Psychophysics*, *66*, 1363–1378. doi:10.3758/BF03195004
- van Asselen, M., & Castelo-Branco, M. (2009). The role of peripheral vision in implicit contextual cuing. *Attention, Perception, & Psychophysics*, *71*, 76–81. doi:10.3758/APP.71.1.76
- van Asselen, M., Sampaio, J., Pina, A., & Castelo-Branco, M. (2011). Object based implicit contextual learning: A study of eye movements. *Attention, Perception, & Psychophysics*, *73*, 297–302. doi:10.3758/s13414-010-0047-9
- Wächter, T., Lungu, O. V., Liu, T., Willingham, D. T., & Ashe, J. (2009). Differential effect of reward and punishment on procedural learning. *Journal of Neuroscience*, *29*, 436–443. doi:10.1523/JNEUROSCI.4132-08.2009