

This is a repository copy of *Reward guides attention to object categories in real-world scenes*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/153411/>

Version: Accepted Version

Article:

Hickey, Clayton, Kaiser, Daniel orcid.org/0000-0002-9007-3160 and Peelen, Marius V (2015) Reward guides attention to object categories in real-world scenes. *Journal of experimental psychology. General*. pp. 264-273. ISSN: 1939-2222

<https://doi.org/10.1037/a0038627>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Running head: Reward in real scenes

Reward guides attention to object categories in real-world scenes

Clayton Hickey^{1,2}, Daniel Kaiser¹, & Marius V. Peelen¹

¹ Center for Mind / Brain Sciences (CIMEC), University of Trento, Italy

² Cognitive Psychology, VU University Amsterdam, The Netherlands

Word count text: ~5700

Word count abstract: 201

Number of figures: 5

Number of tables: 1

Keywords: attention, reward, visual search, natural scenes

Address correspondence to: Clayton Hickey, Center for Mind / Brain Sciences (CIMEC), University of Trento, Corso Bettini 31, Rovereto, Italy, email: clayton.hickey@unitn.it, tel: +39 046 480 8632

Abstract

Reward is thought to motivate animal approach behaviour in part by automatically facilitating the perceptual processing of reward-associated visual stimuli. Studies have demonstrated this effect for low-level visual features such as colour and orientation. However, outside the laboratory it is rare that low-level features uniquely characterize objects relevant for behaviour. Here we test whether reward can prime representations at the level of object category. Participants detected category exemplars (cars, trees, people) in briefly presented photographs of real-world scenes. On a subset of trials successful target detection was rewarded and the effect of this reward was measured on the subsequent trial. Results show that rewarded selection of a category exemplar caused other members of this category to become visually salient, disrupting search when subsequently presented as distractors. Importantly, this occurred even when there was little opportunity for the repetition of visual features between examples, with the rewarded selection of a human body increasing the salience of a subsequently presented face. Selection of a category example thus appears to activate representations of prototypical category characteristics even when these are not present in the stimulus. In this way reward can guide attention to categories of stimuli even when individual examples share no visual characteristics.

Introduction

Theories of reward learning propose that animal approach behaviour is determined in part by perceptual biases created by reward feedback (eg. Toates, 1986; Ikemoto & Panksepp, 1999; Schultz, 2002; Franken, Booij, & van den Brink, 2005). For example, the incentive salience hypothesis of Berridge and Robinson (1998) suggests that reward signals encoded in the phasic release of mesencephalic dopamine ultimately prime the perceptual response to reward-associated stimuli. These objects become salient and attention-drawing, in this way increasing the likelihood that they will be noticed and approached when encountered in the future. This mechanism is thought to have the evolutionary function of biasing vision towards objects likely to provide basic rewards like food.

To date, experimental evidence of this effect in human vision has been provided by studies pairing primary or secondary rewards with the successful detection or discrimination of simple objects defined by low-level visual features. Reward causes processing of these objects to be facilitated, changing how they are subsequently perceived and increasing the likelihood that they will be attended (eg. Anderson, Laurent, & Yantis, 2011, 2012; Anderson & Yantis, 2013; Baldassi & Simoncini, 2011; Della Libera & Chelazzi, 2006, 2009; Kiss, Driver, & Eimer, 2009; Raymond & O'Brien, 2009; Seitz & Watanabe, 2003). Critically, this occurs under circumstances where the priming is of no strategic benefit (Hickey, Chelazzi, & Theeuwes, 2010a, 2010b, 2011; Seitz, Kim, & Watanabe, 2009), and even when it is counter-productive (Hickey et al. 2010a; Hickey & van Zoest, 2012, 2013), indicating that reward may have an impact on vision that is independent of strategy and top-down set.

Existing work examining reward's impact on human vision has thus focused on low-level

visual features like orientation, direction of motion, and colour. However, outside the laboratory the low-level features that define a target vary dramatically with changes in environmental variables like perspective, lighting, and distance. Moreover, search in real life is often for members of a category of stimuli where individual exemplars can be characterized by a wide array of low-level features with little overlap between instances. Under these circumstances the evolutionary utility of a feature-priming mechanisms seems questionable.

Studies of visual attention employing real-world stimuli have demonstrated that humans do in fact have a striking ability to detect the presence of members of familiar real-world object categories, such as cars, people, or animals, even when these are presented in cluttered natural scenes (eg. Potter, 1976; Thorpe, Fize, & Marlot, 1996; see for review Peelen & Kastner, 2014). This occurs in spite of tremendous variation in the low-level visual features that characterize the individual exemplars of the target category and may reflect sensitivity for a set of highly overlearned and closely associated intermediate-level shape features that are diagnostic of a semantic category (Evans & Treisman, 2005; Ullman, Vidal-Naquet, & Sali, 2002; Treisman, 2006; Reeder & Peelen, 2013). Recent evidence has additionally demonstrated that when distractor objects are taken from the same conceptual category as a target they will rapidly and involuntarily draw selective resources, even when the specific target and distractor images share very little in the way of visual characteristics (Wyble, Folk, & Potter, 2013).

Here we test the possibility that reward primes the salience of categories of stimuli in naturalistic images. We had participants search through static landscapes for members of real-world categories - cars, trees, people - where individual examples in each category could be

characterized by very different low-level visual features (see Figure 1). Each trial began with a semantic cue informing the participant of the target category for the trial and each correctly-performed trial could result in reward feedback. The likelihood that reward was received in any trial was random and participants knew this to be the case. This is a critical design feature: reward was not tied to any particular feature of the visual stimuli. There was therefore no strategic incentive to use this feedback to prepare for the next trial. To determine if reward had an impact on object saliency we examined search efficiency as a function of a.) whether a distractor in the display was a member of the target category from the immediately preceding trial, and b.) whether that preceding trial garnered high- or low-magnitude reward outcome (see Hickey et al., 2010a). If reward acts to prime the visual representation of a category of real-world visual objects, search for a target should be disrupted when an exemplar of a recently reward-associated category is present as a distractor in the scene.

*** INSERT FIG 1 ABOUT HERE ***

*** INSERT TABLE 1 ABOUT HERE ***

Experiment 1

Method

Participants

Twenty-four neurologically typical students of the VU University Amsterdam took part in Experiment 1. All gave informed consent prior to participation, reported normal or corrected-to-

normal vision, and received financial compensation. Four participants had false alarm rates of 50% or greater and were discarded from analysis. Seven of the remaining 20 participants were male (mean age: 21.9 +/- 2.8 years SD).

Experimental stimuli and procedure

The experiment took place in a sound-attenuated room. Stimuli were presented to participants via a CRT monitor located 1 meter from the eyes. Participants were presented with black-and-white photos of natural scenes (13.4° x 10.7° visual angle) and asked to report the presence or absence of exemplars of a target category (see Figure 1a). Each trial began with a cue informing the participant of the target category for the current trial - cars, trees, or people - and an exemplar of this category was present in 66% of trials. Scenes were presented for 40 ms before being replaced by a mask, which was created by randomizing x and y pixel coordinates of the image itself. The mask remained onscreen until keyboard response: the 'z' button (left index finger) indicated target presence and the 'm' key target (right index finger) indicated its absence.

Immediately following response the mask disappeared and the scene was presented alone for 1500 ms. Reward feedback was subsequently overlaid at the center of the screen for 2000 ms. If response was correct, feedback was '+10' or '0' in green font, indicating the receipt of 10 points or no points. If response was incorrect feedback was '-10' or '0' in red font, indicating the loss of 10 points or no loss. Feedback magnitude was randomly determined for each trial: when response was correct there was an equal chance that 10 or 0 points would be awarded, and when response was incorrect there was an equal chance of that 10 or 0 points would be lost. Each point had an approximate cash value of €0.006 and participants were paid

based on the number of points accumulated throughout the experiment. Each participant completed 18 blocks of 25 trials and mean pay was €10.96 +/- €1.40 SD. Participants were informed of total accumulated score and cash equivalent at the end of every block, were instructed to try and maximize their earnings through accurate performance, and were explicitly told that the reward magnitude for each correctly-performed trial was random.

Stimuli were taken from a database of 200 black-and-white photographs depicting landscapes and cityscapes. These were manually selected from an online repository of labelled photos (Russell, 2008) such that each image contained at least two of the three possible target categories: 50 contained examples of trees and cars, 50 trees and people, 50 cars and people, and 50 cars, trees, and people. In each target-present trial an image was randomly selected without replacement from one of the three scene types that contained an example of the target, and in each target-absent trial an image was randomly selected without replacement from the scene type that did not contain an example of the target. Images were used in the experiment such that all photos of each scene type were presented before repeating any image of that type. Participants therefore saw all images twice and a few three times. Importantly, category examples embedded within each scene varied dramatically in appearance, spatial location, size, and perspective. For example, a car could be located at any position in the scene, could be partially occluded by other objects, could be viewed from a variety of angles, and could be of any make or model. The photos contained a variety of objects and textures in addition to cars, trees, and people.

Our primary analyses of RT and accuracy focused on two factors: whether the current scene contained a distractor of the category that had acted as target in the immediately

preceding trial, and whether that preceding trial had garnered high-magnitude or low-magnitude reward. The first trial from each block was discarded from analysis, as were all trials where the target category had repeated, where the preceding trial had been incorrectly completed, where the target was absent, or where RT fell more than 3 standard deviations from the per-subject mean. Before the rejection of outliers and inaccurately performed trials this design garnered about 32 trials where the prior target was present in the scene as a distractor where in 16 cases this was preceded by reward.

We focused on results observed when the target category in trial n did not appear as either target or distractor category in trial $n-1$ (see example in Figure 1b). Because all scenes contained either 2 or 3 categories, the analysis therefore included transitions from 2-category to 3-category scenes (with both the previous target and distractor repeated as distractors) versus transitions from 2-category to 2-category scenes (with only the distractor repeated), as illustrated in Table 1.

*** INSERT FIG 2 ABOUT HERE ***

Results

Results from Experiment 1 are illustrated in Figure 2. Response appears slow when the current scene contained one or more distractor items that were exemplars of the same category as the target in the immediately preceding trial, but only when the preceding trial had garnered reward (Figure 2a). Accuracy also appears to decrease under these circumstances (Figure 2b).

Statistical analysis began with two separate repeated-measures analyses of variance (RANOVAs) - one for RT and one for accuracy - with factors for prior reward (reward vs. no-reward) and distractor presence (exemplar from prior target category present as distractor in current scene vs. no exemplar from prior target category present). Analysis of RT revealed an interaction ($F(1,19) = 5.53, p = 0.030$), but no main effects (prior reward: $F < 1$; distractor presence: $F(1,19) = 2.90, p = 0.105$). Analysis of arcsin transformed accuracy also revealed an interaction ($F(1,19) = 4.73, p = 0.043$) and a main effect of reward ($F(1,19) = 5.46, p = 0.031$; distractor presence: $F < 1$). Follow-up t-tests demonstrated that RT was reliably slower ($t(19) = 2.14, p = 0.046$) and accuracy lower ($t(19) = 3.36, p = 0.003$) when the prior target acted as distractor following reward as compared to when no reward was received. Reward had no impact on performance when the target category was repeated between trials (not illustrated; RT: 468 ms vs. 470 ms; $t(19) = 0.152, p = 0.880$; accuracy: 89% vs. 91%; $t(19) = 0.982, p = 0.339$).

We conducted an additional analysis in order to determine if the physical presence of a target in the rewarded trial was necessary to create the priming effect. To this end we analyzed trials that had been preceded by a trial where participants had correctly reported the absence of the target. There was no evidence of the interactive pattern characterizing the reward-priming effect in these results (RT interaction: $F(1,19) = 1.23, p = 0.282$; accuracy interaction: $F(1,19) = 1.53, p = 0.231$). A statistical comparison of RT results garnered by this analysis to those described above, where the target was present in the preceding trial, identified a significant 3-way interaction ($F(1,19) = 6.13, p = 0.023$). The interactive pattern illustrated in

Figure 2a was thus reliably different from the pattern observed when the target had been absent in the preceding trial.

Discussion

Results from Experiment 1 suggest that good outcome following visual selection of a category exemplar causes other members of that category to become salient and attention-drawing. As a result, search for examples of a different target category was disrupted when a member of the reward-associated category was present in the scene as a distractor.

Though our interpretation of these results rests on the idea that reward primes target representations, we did not find that reward impacted target response when the target category was repeated between trials. This suggests the presence of a ceiling effect: the semantic cue, which was always valid, gave participants the opportunity to fully establish a top-down attentional set for the target category. Performance was very good under these circumstances and there was little opportunity for the association of reward to further improve speed or accuracy.

In this first experiment feedback was overlaid on the search scene for 2 s. at the end of each trial. This aspect of the design was motivated by recent theory and modelling work suggesting that the coincidence of stimulation and reward feedback might be important to the creation of visual plasticity (Roelfsema & van Ooyen, 2006; Roelfsema, van Ooyen, & Watanabe, 2010). By this, perceptual learning will occur when a neural representation is activated by attention and diffuse neuromodulatory signals encoding reward feedback are

concurrently introduced to the cellular environment. Thus we wanted to have the stimuli be present at the moment of reward feedback. But the 2 s. interval following reward feedback may have had the adverse effect of allowing participants time to extensively reconsider the scene with the outcome in mind. If reward aroused participants and motivated them to process the scene differently during this interval, such re-examination may have changed processing of the target category in the next trial. To test whether the duration of reward feedback contributed to the effects we obtained - and to confirm our basic finding - we conducted a second experiment. Here reward feedback was presented for a much shorter period before the mask reappeared.

Experiment 2

Method

Participants

Seventeen neurologically typical students of the VU University Amsterdam took part in Experiment 2. All gave informed consent prior to participation, reported normal or corrected-to-normal vision, and received financial compensation. Two had false alarm rates of 50% or greater and were discarded from analysis. Eight of the remaining 15 were male (mean age: 21.5 +/- 2.8 years SD).

Experimental stimuli and procedure

In Experiment 2 feedback was overlaid on the scene for 250 ms before the mask reappeared for 1250 ms. All other design and analysis parameters were as in Experiment 1.

*** INSERT FIG 3 ABOUT HERE ***

Results

Statistical analysis again took the form of two RANOVAs. Analysis of RT revealed an interaction ($F(1,14) = 5.82, p = 0.030$) but no main effects ($F_s < 1$). RT results show an apparent difference after no-reward trials as a function of prior target category presence that was not evident in Experiment 1, with response speeding when an exemplar of the prior target category is present in the scene (Figure 3a, broken trace), but this was not reliable ($t(14) = 1.41, p = 0.180$). A follow-up t-test demonstrated that RT was marginally slower when the prior target acted as distractor following reward as compared to when no reward was received ($t(14) = 2.14, p = 0.050$). Analysis of arcsin transformed accuracy revealed a main effect of distractor presence ($F(1,14) = 13.19, p = 0.003$) but no other effects (reward: $F(1,14) = 3.17, p = 0.097$; interaction: $F < 1$). Reward had no impact on performance when the target category was repeated between trials (not illustrated; RT: 541 ms vs. 539 ms; $t(14) = 0.124, p = 0.903$; accuracy: 88% vs. 89%; $t(14) = 0.519, p = 0.612$).

As in Experiment 1, there was no evidence of the interactive pattern in RT when analysis was constrained to conditions where the preceding trial had resulted in a correct target absent response (interaction: $F < 1$). Statistical analysis of results observed when the preceding target had been present versus absent garnered a marginally significant three-way interaction ($F(1,14) = 3.88, p = 0.069$), again suggesting that the physical presence of the

target in trial $n-1$ was necessary for the interactive pattern that characterizes the reward-priming effect.

Discussion

Experiment 2 reproduces the RT effect identified in Experiment 1 with a roughly equivalent effect size, demonstrating that rewarded selection of a category example causes members of that category to become salient and attention-drawing in the next trial. A corresponding effect on accuracy was observed in Experiment 1, but failed to appear in Experiment 2, suggesting a difference in response criterion or speed-accuracy trade-off between the samples. Consistent with this, RT was generally slower in Experiment 2. In Experiment 2 reward feedback was overlaid over the scene for a much shorter interval than was the case in Experiment 1 and the scene was subsequently masked. This reduced the possibility that participants might extensively re-examine the scene following reward feedback, and thus the possibility that such rumination might underlie the observed effect.

What does reward actually prime? Object category exemplars in Experiments 1 and 2 varied in terms of perspective, size, shape, and location, and were often partly occluded by other objects. Furthermore, objects were embedded in cluttered visual scenes with a number of other objects that shared the low-level features of the target category. It is therefore unlikely that the effect is caused by the priming of low-level features like orientation. One alternative is that reward might prime object representations of intermediate complexity. These mid-level features, like a person's arm or a car's tire, may be highly diagnostic of a specific category (eg.

Evans & Treisman, 2005; Ullman et al., 2002; Treisman, 2006; Reeder & Peelen, 2013). Such a priming effect could accordingly impact detection of category exemplars that have the same mid-level features even if they share few low-level characteristics.

If reward is priming the mid-level features physically present in a category exemplar it should have no impact on the visual processing of subsequent category exemplars characterized by other prototypical category features. To test whether reward priming can act to potentiate the set of features that define a category, rather than the specific features that were physically present in the display, we developed a new variation of the paradigm. As in Experiments 1 and 2, participants detected examples of real-world categories in natural scenes with the target category cued in each trial. However, we limited the design such that only two categories could act as target - cars and people - and we employed a new set of images (see Figure 4). These were selected or generated such that they contained two types of person exemplar. In one case only the head and shoulders were visible, with the rest of the body occluded or falling outside the frame of the image. In the other, only the legs and torso were visible, again with the rest of the body occluded or outside the frame (see Figure 4a). A head-and-shoulders scene thus contained a person, as did a legs-and-torso scene, but these shared no diagnostic visual features that could be used to determine the presence of a person.

If the influence of reward identified in Experiments 1 and 2 involved category-level priming our expectation was that rewarded selection of a person would cause subsequent images of people to disrupt search, even when the specific features defining this object changed. However, we considered the possibility that this influence might rely on the specific strength of relationship between a feature and its corresponding category and accordingly

conducted a pilot study to guide the development of Experiment 3. In this pilot we had 12 people search through the new scenes, giving them random-magnitude reward after each correctly performed trial. Results showed that a head-and-shoulders distractor tended to disrupt search when it followed rewarded selection of a legs-and-torso target, expressing as a marginal interaction between prior reward and distractor presence ($F(1,11) = 2.58$, $p = 0.136$). The corresponding effect for the alternate order, where rewarded selection of a head-and-shoulders target was followed by a legs-and-torso distractor, was less reliable ($F < 1$) and the difference between these interactions itself trended toward significance ($F(1,11) = 3.86$, $p = 0.075$). While we do not want to draw strong conclusions from these marginal results, the directionality of this pattern is consistent with results from the face-processing literature. Here consideration of body features is found to impact the perceptual representation of a subsequent face (Ghuman, McDaniel, & Martin, 2010; see also Brandman & Yovel, 2012; Lai, Oruc, & Barton, 2011), but there is no evidence in the literature of the reverse relationship.

Given that prior work suggests that consideration of a body has a particularly strong impact on the perceptual processing of a face, and given the trend in our pilot data, we designed Experiment 3 to test the specific possibility that rewarded selection of a legs-and-torso target would prime a head-and-shoulders distractor. Such a pattern would clearly demonstrate that reward-priming of visual categories can occur under circumstances where this could not reflect the priming of mid- or low-level visual features. To this end, Experiment 3 employed three key scene types: those that contained head-and-shoulders people but no cars, those that contained both head-and-shoulders people and cars, and those that contained legs-and-torso people but no cars. We were specifically interested in results observed when

rewarded selection of a legs-and-torso target preceded a trial containing a head-and-shoulders distractor (see Table 1b).

*** INSERT FIG 4 ABOUT HERE ***

Experiment 3

Method

Participants

Based on an estimate of effect size derived from our pilot results we substantially increased the sample size of Experiment 3. Fifty neurologically typical participants recruited at the Center for Mind/Brain Sciences at the University of Trento took part. All gave informed consent prior to participation, reported normal or corrected to normal vision, and received financial compensation. Two participants had false alarm rates of 50% or greater and were discarded from analysis. 12 of the remaining 48 participants were male (mean age: 23.4 +/- 4.7 years SD).

Experimental stimuli and procedure

As in prior experiments, participants were presented with black-and-white photos of natural scenes (15° x 11.25° visual angle) on a CRT monitor and asked to report the presence or absence of exemplars of a target category via keyboard button presses. Participants in

Experiment 3 responded with their dominant hand, pressing the right arrow key when targets were present and the left arrow key when they were absent. The cued target category was more commonly people (72% of trials) than cars (28% of trials), targets were present in 66% of trials, and scenes were presented for 83 ms before being replaced by a mask that remained onscreen until response. Immediately following response the mask disappeared and the scene was presented alone for 1500 ms. Reward feedback was subsequently overlaid at fixation for 250 ms before the mask reappeared for 1250 ms. Reward feedback had the same visual characteristics as in prior experiments and reward magnitude was determined in the same manner. A point had an approximate cash value of €0.0013. Each participant completed 624 trials (divided into four blocks) and mean pay was €11.49 +/- €0.37 SD. Participants were informed of total accumulated points at the end of every block, were instructed to try and maximize their earnings through accurate performance, and were explicitly told that the reward magnitude for each correctly-performed trial was random.

Three types of target-present scene were employed: in 25 cases the scene contained head-and-shoulders people but no cars, in 25 it contained legs-and-torso people but no cars, and in 25 it contained head-and-shoulders people as well as cars. A set of control images containing no people or cars were used for target-absent trials. Approximately half of the scenes containing head-and-shoulders or legs-and-torso people were artificially generated. This involved the use of image editing software to cut and move visual objects such that parts of people were occluded or to crop images to the same purpose.

The first trial from each block was excluded from analysis, as were all trials where the target category had repeated, where the target was absent, where the preceding response had

been an error, and where RT fell more than 3 SDs from the mean. Critical trials were those where the preceding trial had resulted in correct detection of a legs-and-torso target and the scene in the current trial contained both a car target and head-and-shoulders distractor. To increase the number of these critical trials, scenes containing a car target and head-and-shoulders distractor were always preceded by trials where the cued target was people, the scene did not contain a car distractor, and either a legs-and-torso target was present (66% of cases) or there was no example of a person in the scene (33% of cases). Note that with this confine, a critical trial followed a legs-and-torso target scene in only 12.5% of cases, and thus a legs-and-torso target scene was not a reliable predictor of subsequent trial characteristics. Before the rejection of outliers and incorrect trials this design garnered about 44 trials per participant in which a head-and-shoulders distractor followed a legs-and-torso target, where in 22 cases this was preceded by reward.

*** INSERT FIG 5 ABOUT HERE ***

Results

Results are illustrated in Figure 5. Initial examination of the across-participant RT results identified rightward skew (Shapiro-Wilk test of normality, $W = 0.946$, $p = 0.028$) and analysis of RT was accordingly based on log-transformed data. As in prior experiments, there is an RT cost when the scene contained an example of the prior target that was exacerbated by prior reward. A RANOVA of RT with factors for distractor presence and reward identified a main effect of distractor presence ($F(1,47) = 22.65$, $p < 10^{-4}$) and a critical interaction ($F(1,47) =$

5.20, $p = 0.027$; prior reward: $F(1,47) = 1.96$, $p = 0.168$). A follow-up t-test demonstrated that RT was slower when the prior target acted as distractor following reward as compared to when no reward was received ($t(47) = 2.68$, $p = 0.010$). Analysis of arcsin transformed accuracy revealed a main effect of distractor presence ($F(1,47) = 5.37$, $p = 0.025$), but no other effects ($F_s < 1$).

General Discussion

Results from three experiments demonstrate that good outcome following selection of a category exemplar causes other members of that category to become salient and attention-drawing. As a result, search for examples of a different target category will be disrupted when a member of the reward-associated category is present in a visual scene as a distractor. Reward thus primes the visual processing of categories of visually heterogeneous objects embedded in cluttered real-world scenes.

The reward priming identified here does not appear to rely on top-down attentional set. Participants in the study knew that reward was randomly determined and that the category indicated by the semantic cue at the beginning of each trial was the only relevant category on that trial. There was no reason for them to establish a strategic set for the target category that had garnered reward previously, but there was clear motivation to establish a set for the cued target category. In spite of this, exemplars taken from the recently-rewarded category disrupted search for the target. This is consistent with the idea that reward's impact on categorical salience can be automatic and cognitively impenetrable, as has been suggested of

reward's influence on the salience of low-level visual features (eg. Anderson & Yantis, 2013; Hickey et al. 2010a; Hickey & van Zoest, 2012).

Prior work investigating low-level reward priming has shown that stimuli with features that recently characterized rewarding targets become salient and attention-drawing (eg. Hickey et al., 2010; Anderson, Laurent, & Yantis, 2011). In contrast, in Experiment 3 we find that rewarded selection of a person will prime subsequent examples of people even when the two instances are characterized by a very different set of visual characteristics. There is a remote possibility that this reflects an impact of reward on a very small set of features that were shared across head-and-shoulders and legs-and-torso category examples. For example, both head-and-shoulders and legs-and-torso people were characterized by clothing, and the texture of this clothing may have repeated across trials. However, our image set was large and composed of photos taken under different environmental conditions, each one containing people of varying race wearing very different clothes and costumes. The possibility of such coincidental feature repetition was therefore remote, and even were this to occur in a handful of trials this would not be adequate to drive the RT effects evident in our results. We accordingly interpret our results as evidence that reward not only primes the features that are physically present in a stimulus, but also the set of characteristics that are not currently present but define the category to which this object belongs.

We think this might be a reflection of the nature of categorical representation. Categories of knowledge have long been modelled as hierarchical networks of conceptual nodes (e.g. Quillian, 1962) with activation of one node spreading to related concepts (e.g. Collins & Loftus, 1975). Strongly connected nodes reflect strong conceptual relationships, and

a node with such connections will be sensitive to the activation of its neighbours. Visual selection of a stimulus that activated one node of a network - a human leg, for example - may cause a spread of activity such that prototypical features and characteristics of the category - like faces - become activated even when physically absent from the evoking stimulus. This activation may render such representations sensitive to the diffuse reinforcement signals in visual cortex thought to underlie visual plasticity (Roelfsema & van Ooyen, 2005; Roelfsema, van Ooyen, & Watanabe, 2010). However, this idea requires a caveat. In computational models spreading activation is necessarily a small effect that influences only directly adjacent nodes in the network and lasts only for a very short period of time. It is not immediately apparent how a physiologically subtle effect of this nature would produce the rather large change in overt behaviour we observe in the current results.

A complementary possibility is that reward directly primes higher-level object representations that are activated by multiple object parts and object views (Marr & Nishihara, 1978). Evidence for such representations comes from monkey electrophysiology, where cells in temporal cortex have been observed to respond equally to visually heterogeneous parts of the same object (e.g. the head and body; Wachsmuth et al. 1994) or to different objects that are strong visual associates (Messinger et al, 2001; Sakai & Miyashita, 1991). Priming at this representational stage would be effective when individual category examples do not overlap in terms of low- or mid-level visual features.

While we emphasize the novel evidence of categorical priming provided in this study, this is not meant to negate the importance or efficacy of the direct reward-priming of visual features described in prior work (eg. Chelazzi & Della Libera, 2009; Hickey et al. 2010;

Anderson et al. 2011). Under the vast majority of circumstances exemplars of an object category will at least share mid-level features: trees tend to have trunks, cars tend to have wheels, and people tend to have arms. Under constrained circumstances they will even share low-level features: red berries, vertical street signs, long pointy knives in x-rayed suitcases. The hierarchical network model of reward-priming we describe above has the appealing quality of accounting not only for indirect priming, via spreading activation, but also for direct potentiation of the visual features physically present in a stimulus. Under the vast majority of circumstances observation of reward-priming in vision will reflect a combination of such direct effects on low- and mid-level features alongside indirect effects instantiating the priming of a conceptual category.

In this context, it is also important to emphasize that reward-priming of categories appears to remain very much a visual effect, requiring experience with visual stimuli to be elicited. Evidence for this is provided by analysis of experimental results under circumstances where the target was absent in trial $n-1$, when participants received reward for correctly reporting the absence of this stimulus. If the mere act of establishing an attentional set for a category of stimuli were sufficient to create sensitivity to a reward signal, then these trials should have shown the same sort of priming effects identified in conditions where the target had been physically present in the display. In fact, we found no evidence of priming following target-absent trials and a reliable difference between results observed under these circumstances and results observed when the prior target was present in the scene. Being rewarded for searching for a car does not appear sufficient to prime car-like features; actual experience of a category exemplar seems required for priming to occur.

As noted in the introduction, the visual features defining an example of a target category change drastically as a function of environmental variables like perspective and distance, and individual category exemplars often share very few low-level characteristics. If reward priming relied solely on the potentiation of the specific visual features present in a reward-associated object it would be frequently ineffective. Here we demonstrate that reward can additionally potentiate visual processing of categories of visually heterogeneous real-world objects, in this way guiding vision through cluttered naturalistic environments.

References

- Anderson, B.A., Laurent, P.A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences USA*, 108, 10367-10371.
- Anderson, B.A. & Yantis, S. (2013). Persistence of value-driven attentional capture. *Journal of Experimental Psychology: Human Perception & Performance*, 39(1), 6-9.
- Baldassi, S. & Simoncini, C. (2011). Reward sharpens orientation coding independently of attention. *Frontiers in Neuroscience*, 8, 5-13.
- 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), 309-69.
- Brandman, T., & Yovel, G. (2012). A face inversion effect without a face. *Cognition*, 125(3), 365-372.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological review*, 82(6), 407.
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methodology in Psychology*, 1(1), 4-45.
- Della Libera, C., & Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychological Science*, 17(3), 222-7.
- Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, 20(6), 778-784.
- Evans, K. & Treisman, A. (2005). Perception of objects in natural scenes: Is it really attention free? *Journal of Experimental Psychology: Human Perception and Performance*, 31(6), 1476-1492.
- Franken, I. H. A., Booij, J. & van den Brink, W. (2005). The role of dopamine in human addiction: from reward to motivated attention. *European journal of pharmacology*, 526, 199-206.
- Ghuman, A. S., McDaniel, J. R., & Martin, A. (2010). Face adaptation without a face. *Current Biology*, 20(1), 32-36.
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010a). Reward guides attention via anterior cingulate cortex. *Journal of Neuroscience*, 30(33), 11096 - 11106.
- 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.
- 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(1), 117-128.

Hickey, C. & van Zoest, W. (2012). Reward creates oculomotor salience. *Current Biology*, 22(7), R219-220.

Hickey, C., & van Zoest, W. (2013). Reward-associated stimuli capture the eyes in spite of strategic attentional set. *Vision research*, 92, 67-74.

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.

Kiss M, Driver J, Eimer M (2009) Reward priority of visual target singletons modulates ERP signatures of attentional selection. *Psychological Science*, 20, 245–251.

Lai, M., Oruç, I., & Barton, J. J. (2012). Facial age after-effects show partial identity invariance and transfer from hands to faces. *Cortex*, 48(4), 477-486.

Li, F., VanRullen, R., Koch, C., & Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences*, 99(14), 9596-9601.

Messinger, A., Squire, L. R., Zola, S. M., & Albright, T. D. (2001). Neuronal representations of stimulus associations develop in the temporal lobe during learning. *Proceedings of the National Academy of Sciences*, 98(21), 12239-12244.

Peelen, M. V., & Kastner, S. (2014). Attention in the real world: toward understanding its neural basis. *Trends in cognitive sciences*, 18(5), 242-250.

Potter, M. C. (1999). Understanding sentences and scenes: The role of conceptual short-term memory. In Veronika Coltheart (Ed.), *Fleeting Memories: Cognition of Brief Visual Stimuli* (13-46). MIT Press.

Quillian, M. R. (1962). A revised design for an understanding machine. *Mechanical translation*, 7(1), 17-29.

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(8), 981-988.

Reeder, R. R., & Peelen, M. V. (2013). The contents of the search template for category-level search in natural scenes. *Journal of vision*, 13(3), 1-13.

Roelfsema, P. R., & van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural computation*, 17(10), 2176-2214.

Roelfsema, P., van Ooyen, A., & Watanabe, T. (2010). Perceptual learning rules based on reinforcers and attention. *Trends in Cognitive Science*, 14(2), 64-71.

Russel, B., Torralba, A., Murphy, K., & Freeman, W. (2008). LabelMe: A database and web-based tool for image annotation. *International Journal of Computer Vision*, 77, 157-173.

- Sakai, K., & Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. *Nature*, 354(6349), 152-155.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, 36(2), 241-263.
- Seitz, A., Kim, D., & Watanabe T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*, 61, 700-707.
- Seitz, A. & Watanabe, T. (2003). Is subliminal learning really passive? *Nature*, 422, 36.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582), 522-522.
- Toates, F. (1986). *Motivational Systems*, Cambridge Univ. Press, Cambridge, 1986.
- Treisman, A. (2006). How the deployment of attention determines what we see. *Visual Cognition*, 14, 411-443.
- Ullman, S., Vidal-Naquet, M., & Sali, E. (2002). Visual features of intermediate complexity and their use in classification. *Nature Neuroscience*, 5(7), 682-687.
- Wachsmuth, E., Oram, M. W., & Perrett, D. I. (1994). Recognition of objects and their component parts: responses of single units in the temporal cortex of the macaque. *Cerebral Cortex*, 4(5), 509-522.
- Wolfe, J., Vo, M., Evans, K., & Greene, M. (2011). Visual search in scenes involves selective and nonselective pathways. *Trends in Cognitive Sciences*, 15(2), 77-84.
- Wyble, B., Folk, C., & Potter, M. C. (2013). Contingent attentional capture by conceptually relevant images. *Journal of experimental psychology: human perception and performance*, 39(3), 861-871.

Acknowledgment

This work was supported by a Netherlands Organization for Scientific Research (NWO) VENI grant to CH (016-125-283) and the Autonomous Province of Trento, Call 'Grandi Progetti 2012', project 'Characterizing and improving brain mechanisms of attention – ATTEND'.

Tables

A. Experiments 1 and 2

Trial n-1	Trial n <i>Distractor present</i>	Trial n <i>Distractor absent</i>
<u>cars</u> , trees	cars, trees, <u>people</u>	trees, <u>people</u>
cars, <u>trees</u>	cars, trees, <u>people</u>	cars, <u>people</u>
<u>cars</u> , people	cars, <u>trees</u> , people	<u>trees</u> , people
cars, <u>people</u>	cars, <u>trees</u> , people	cars, <u>trees</u>
<u>trees</u> , people	<u>cars</u> , trees, people	<u>cars</u> , people
trees, <u>people</u>	<u>cars</u> , trees, people	<u>cars</u> , trees

B. Experiment 3

Trial n-1	Trial n <i>Distractor present</i>	Trial n <i>Distractor absent</i>
<u>legs-and-torso</u>	<u>car</u> , head-and-shoulders	<u>car</u>

Table 1. A. Overview of target-present trial transitions included in the main analyses of Experiments 1-2. All scenes shown in Experiments 1 and 2 contained either 2 or 3 object categories. Underlined categories indicate the target category on that trial. B. Overview of target-present trial transitions in the main analysis of Experiment 3. Scenes in Experiment 3 could contain 0, 1 or 2 categories.

Figure Captions

Figure 1 - a.) General paradigm as employed in Experiment 1. Participants searched through images of natural scenes for examples of a cued target category. Performance on a given trial (trial n) was analyzed as a function of characteristics of the prior trial (trial $n-1$). b.) Factors of interest: whether the target category from trial $n-1$ was present in the scene as a distractor and whether correct performance in that preceding trial had garnered reward. Note that cues are illustrated in English here but were in Dutch in Experiments 1 and 2 and Italian in Experiment 3.

Figure 2 - a.) Reaction time and b.) accuracy results from Experiment 1. Reaction times here and in subsequent figures reflect the conditional median of correct trials, accuracy reflects the conditional mean, and error bars represent within-subject standard error (Cousineau, 2005).

Figure 3 - a.) Reaction time and b.) accuracy results from Experiment 2.

Figure 4 - a.) Examples of scene stimuli employed in Experiment 3. The first six images contain examples of people defined by presence of legs and torso; the second six contain examples of people defined by presence of head and shoulders. b.) Conditional design of Experiment 3. The critical condition occurred when participants detected a legs-and-torso target in trial $n-1$ before searching for a car in a scene containing a head-and-shoulders distractor in trial n .

Figure 5 - a.) Reaction time and b.) accuracy results from Experiment 3.

Figure 1

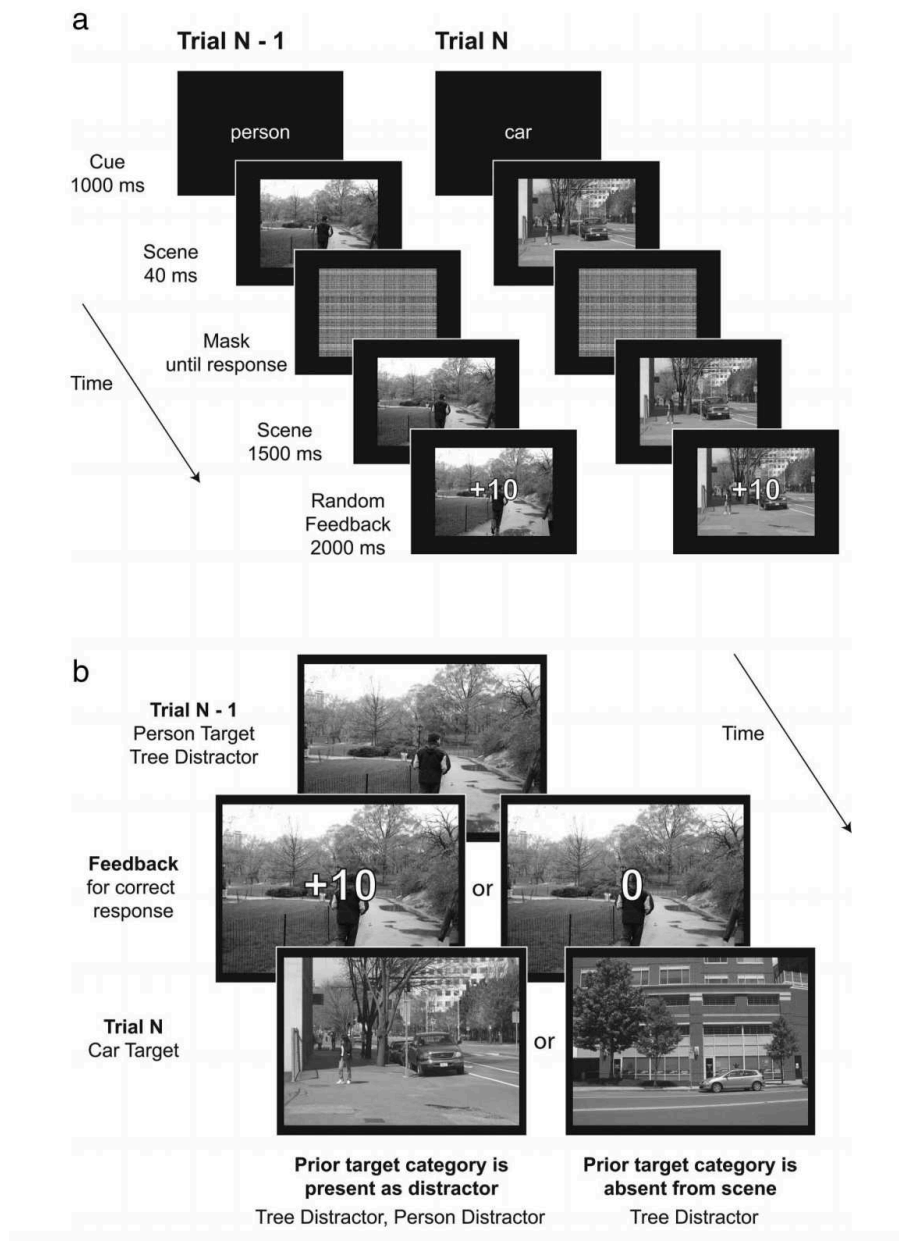


Figure 2

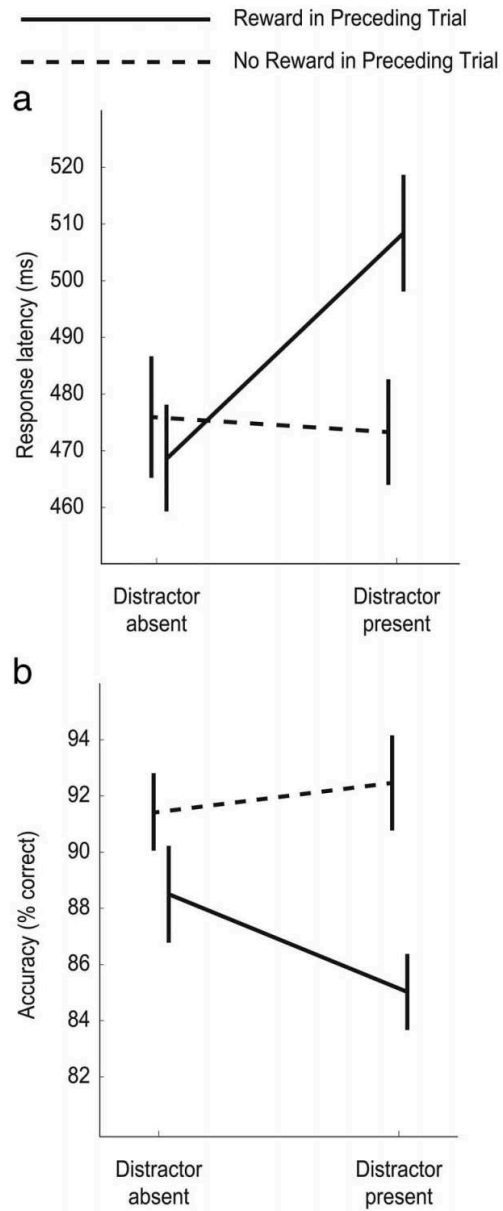


Figure 3

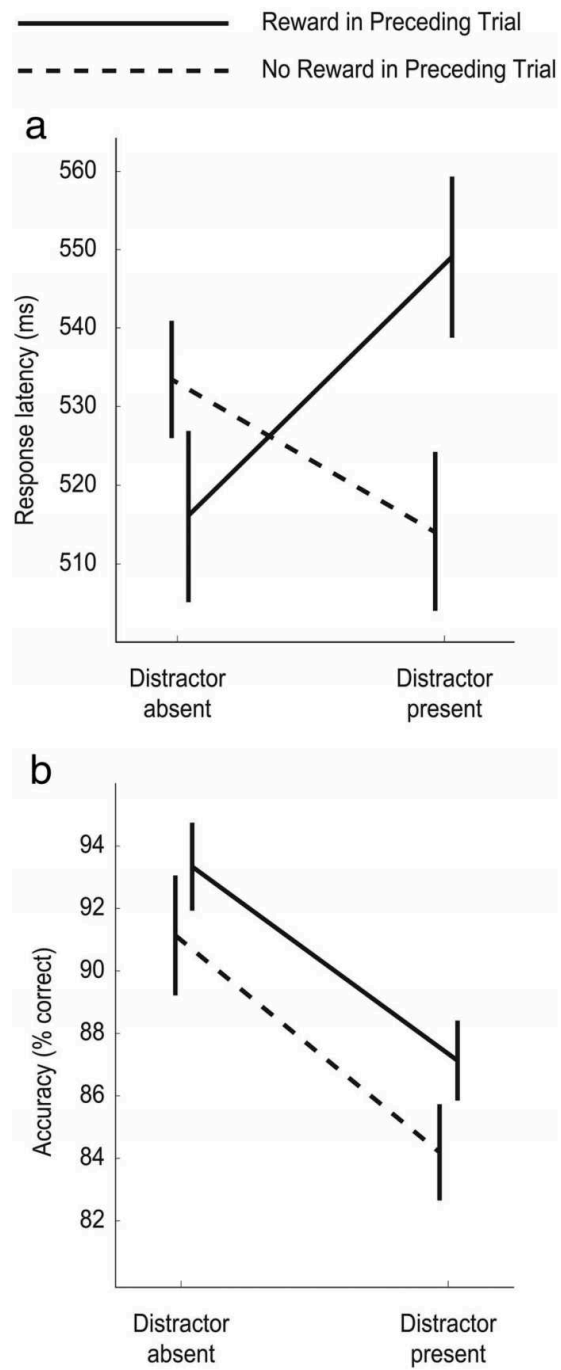


Figure 4

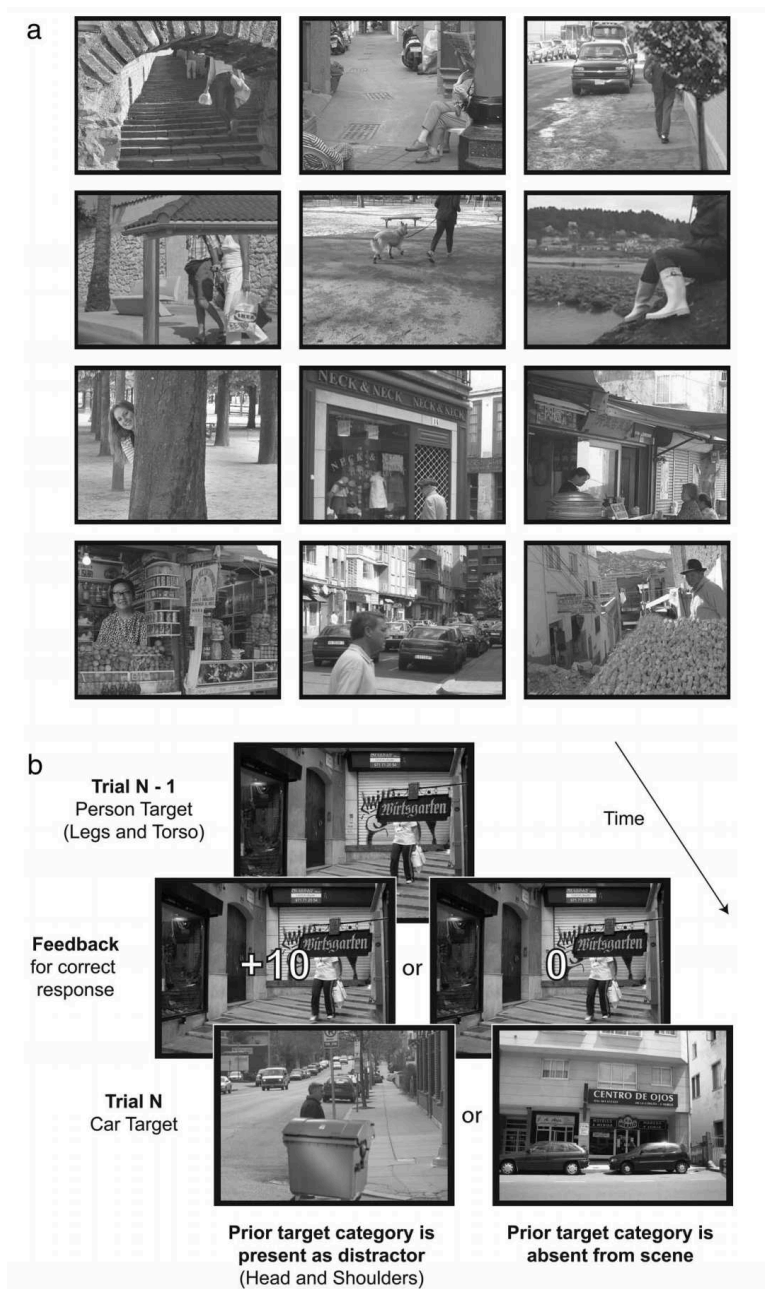


Figure 5

