



Reward sharpens orientation coding independently of attention

Stefano Baldassi^{1*} and Claudio Simoncini²

¹ Department of Psychology, University of Florence, Florence, Italy

² Team DyVA – Institut de Neurosciences Cognitives de la Méditerranée, Centre National de la Recherche Scientifique, Marseille, France

Edited by:

Julia Trommershäuser, New York University, USA

Reviewed by:

Justin L. Gardner, RIKEN Brain Science Institute, Japan

Leila Montaser-Kouhsari, Caltech, USA

*Correspondence:

Stefano Baldassi, Department of Psychology, University of Florence, Via di San Salvi, 12 Padiglione 26, 50135 Florence, Italy.
e-mail: stefano.baldassi@unifi.it

It has long been known that rewarding improves performance. However it is unclear whether this is due to high level modulations in the output modules of associated neural systems or due to low level mechanisms favoring more “generous” inputs? Some recent studies suggest that primary sensory areas, including V1 and A1, may form part of the circuitry of reward-based modulations, but there is no data indicating whether reward can be dissociated from attention or cross-trial forms of perceptual learning. Here we address this issue with a psychophysical dual task, to control attention, while perceptual performance on oriented targets associated with different levels of reward is assessed by measuring both orientation discrimination thresholds and behavioral tuning functions for tilt values near threshold. We found that reward, at any rate, improved performance. However, higher reward rates showed an improvement of orientation discrimination thresholds by about 50% across conditions and sharpened behavioral tuning functions. Data were unaffected by changing the attentional load and by dissociating the feature of the reward cue from the task-relevant feature. These results suggest that reward may act within the span of a single trial independently of attention by modulating the activity of early sensory stages through an improvement of the signal-to-noise ratio of task-relevant channels.

Keywords: reward, attention, orientation discrimination

INTRODUCTION

The activity of the visual channels, both at the neuronal and at the overall behavioral level, can be modulated by several sources of influence. Many modulatory activities depend on the global behavioral state of the organism, driven by cognitive, emotional or motivational factors. Since these states have a profound impact on the behavioral performance of the individuals, determining successes or failures of goal-directed behavior, their associated mechanisms of action have attracted the interests of psychologists, cognitive neuroscientists, and neurophysiologists for long time. Attention, learning, and reward are probably the most studied modulating factors of the sensory systems and of perceptual performance. In general, the idea of attention typically reflects fast, short-term modulation based on exogenous or endogenous cues to bias processing power toward specific spatial location or stimulus features. Perceptual learning instead reflects positive changes in the ability to detect or discriminate a stimulus as an effect of repeated presentations of the same stimulus (Gilbert et al., 2009). On the other hand, reward of specific actions or classes of stimuli is typically investigated assuming that it exerts long-term effects on sensory channels and that these effects would result in learning of specific stimuli, classes of stimuli, and/or specific responses.

Moreover, visual selective attention is mainly studied in its relations to changes of the early stages of the input-output flow of information processing, with a focus on the Visual Area V4 (McAdams and Maunsell, 1999; Reynolds et al., 1999; Ghose and Maunsell, 2008), V2 (Reynolds et al., 1999; Fang et al., 2009), V1 (Watanabe et al., 1998a,b; Kamitani and Tong, 2006), and as early as on the LGN (McAlonan et al., 2006, 2008). Instead, reward has

been widely studied as a variable affecting the later stages, closer to mechanisms related to visual-motor transformations (Schultz et al., 2000), to the decision-making modules (Glimcher and Rustichini, 2004; Hampton and O’doherly, 2007), and to the overt behavior (Behrens et al., 2007). More recently a number of studies have shifted the focus backward attempting to determine the effect of reward to purely sensory areas and opening new doors for re-framing the functional properties of the early visual modules (Shuler and Bear, 2006; Serences, 2008; Seitz et al., 2009). However, since reward constitutes a key tool in each neurophysiological paradigm of attention, studies on the early effect of reward cannot easily disentangle the effects of reward with those of attention (Maunsell, 2004). Indeed, recent proposals have raised the idea that perceptual performance can be modulated by reward through its action on the attentional system (Della Libera and Chelazzi, 2006; Peck et al., 2009), implying that attention has a monopoly over the modulation of perception. So, when dealing with early modulation of sensory coding, what are the functional relationships between reward, on one hand, and attention and learning, on the other hand? Is it possible to dissociate the modulatory effects of reward from those of attention and of learning? Here we try to answer these questions by investigating whether the probability of obtaining a given reward can yield a change in perceptual performance when attention is engaged in a concurrent task and learning is prevented by making the reward value associated to specific stimuli contingent on a trial-to-trial base. We have used a recently introduced psychophysical paradigm (Baldassi et al., 2006) to measure orientation discrimination acuity for a simple peripheral target (a task assumed to summon early mechanisms; Regan and Beverley, 1985; Bradley et al., 1987) and to

obtain at once a quantitative estimate of the observer's noisy internal response distributions for any physical value of the target, that in this article we will be referred to as behavioral tuning functions. Attention was controlled through the use of a concurrent task of varying difficulty, that has the key potential of showing independence of resources (Lavie, 2005; Alais et al., 2006), while learning could be excluded based on the fact that the same stimulus and the same response could be associated to one of two probabilities of obtaining reward (low reward probability, LRP, equal to 0.1 or high reward probability, HRP, equal to 0.9) unpredictably at each trial based on a precue (see **Figure 1**). We found that a higher likelihood of earning credit to obtain a Scratch-and-Win ticket, a highly efficient and effective reward even in non-gamblers, improved performance. In particular, higher reward rates produced finer orientation acuity, as revealed by lower thresholds (about 50% decrease), and this was possibly due to a significant change of the channel's signal-to-noise ratio (SNR), as revealed by sharper behavioral tuning functions when the reward was more likely to be achieved. The reward-based modulation of the peripheral target was unaffected by the difficulty, or load, of an interfering task at fixation. Moreover, the effect was dissociated from the nature of the cue, as it remained stable when the cue was modulated in the color domain and the task in the orientation domain. Our results are coherent with the possibility that reward may modulate perceptual performance independently of both attention and learning and offers novel insight for studying reward and attention by measuring their effects independently in the context of the same experimental paradigm.

MATERIALS AND METHODS

OBSERVERS

A total of six observers participated in this research. Two of them to the main experiment and the X-cue experiment, two to the feature-independent cue experiment, and two to all the experiments. They were undergraduate students of the Faculty of Psychology of the University of Florence, all naïve to the purpose of the study. Informed consent was obtained from all the subjects involved prior to the start of the experiment. They were also non-gamblers based on the criterion according to which subjects involved in gambling activities (including purchase of lottery tickets) more than once a month were excluded. They all had normal or correct-to-normal vision. Three of the subjects completed 600 trials for condition to reach a stable threshold measure, while three were selected to complete 2000 to 2400 trials per condition in order to achieve a reliable sample size to measure both thresholds and behavioral tuning functions. For the X-cue experiment we collected 600 trials per observer allowing only threshold analysis.

APPARATUS

Stimuli were created on a G4 Power Macintosh using the Psychophysics Toolbox v. 2.55 (Brainard, 1997; Pelli, 1997) and displayed on a 17" gamma-corrected CRT monitor (Mitsubishi Diamond Pro) with average luminance equal to 29 cd/m².

DESIGN AND STIMULI

Experimental trials in different experiments were structured combining five different segments: (A) a *foveal task* that consisted in the count of briefly flashed disks at fixation, devised to load

attention for the (B) *peripheral task*, an orientation task devised to probe the main effects of reward sought in the study, (C) a *reward cue*, actually shown at the beginning of the sequence, that informed the observers about the upcoming peripheral stimulus feature that yielded the highest reward rate, (D) a *response page(s)*, in which observers could give a response through the use of a mouse, and (E) a *feedback page* updating observers about the outcome of each trial and the accumulation of reward-based credit. **Figure 1A** schematizes a trial of the main experiment and its temporal structure.

Stimuli of the *foveal task* were disks with a diameter of 0.5° of visual angle flashed foveally for 150 ms. In order to vary the attentional load of the task, which in turn would impact the relative difficulty of the two main tasks, we administered two attentional conditions, a light load (LL) condition and a heavy load (HL) condition. The two conditions differed in the contrast of the disks, which was varied from a level of 80%, at which the stimulus was well visible, to a level ranging from 4 to 8% (adjusted in different subjects to match detection threshold), for the LL and the HL condition, respectively. The attention loading task required observers to count the number of a sequence of serially presented disks flashed a variable number of times (3–14 on a random base). In order to maintain attention foveally between consecutive flashes, the interval between two consecutive disks was jittered between 0.4 and 4 s to avoid predictability about the timing of the upcoming disk. Counting accuracy remained stable at around 95 and 55% respectively for the LL and HL condition, respectively. In order to ensure maximum attentional load with the central task, wrong counting voided the trial; for any voided trial a new trial was appended at the end of the block.

Stimuli of the *peripheral task* were Gabor patches (2 cpd sinusoidal gratings vignetted by a 2D Gaussian modulation of contrast with a space constant [σ] of 0.5°) displayed at a contrast of 80% at an eccentricity of 7° to the left or to the right of fixation. The peripheral patch was delivered for 150 ms in complete synchrony with one of the central disks. The disk containing the peripheral target was fully unpredictable. Only the first and last disk were excluded from the pool of disks that could be accompanied by the peripheral stimulus in order to maximize that attention was well focused on the central task and that it continued to be allocated foveally after the peripheral stimulus had appeared. In all the experiments subjects were asked to report the direction and the magnitude of a tilt offset of the Gabor patch. The tilt was given randomly clockwise (CW) or counterclockwise (CCW) and its amount varied in octaves from $\pm 2^\circ$ to $\pm 32^\circ$ in the main experiment and in the X-cue experiment (see later for details on the experimental conditions) and $\pm 0.5^\circ$ to $\pm 16^\circ$ in the feature-independent cue experiment to yield a complete psychometric function. The reference axis around which the target was tilted was $+45^\circ$ or -45° , randomly, in the main and the X-cue experiment, while it was always 0° (i.e., vertical) in the feature-independent cue experiment. The stimulus space is exemplified in **Figure 1B**.

The stimulus acting as *reward cue* varied in three different experiments but in all cases it consisted of one of two possible configurations. Note that this segment of the trial, when present, was always the very first stimulus displayed; we are explaining it after the two tasks only for the sake of clarity. In all conditions

the reward cue signaled that the match of its main feature with a key feature of the peripheral target implied a high probability of earning a reward (90%), identifying HRP trials, while the lack of match implied that the reward rate was as low as 10%, identifying LRP trials. In the *main experiment* the reward cue was a foveal oblique line subtending 3° of visual angle, visualized for 500 ms before the stimulus array (see later) and tilted either 45° CW or CCW from vertical. In this case the match had to be established between the cue and the axis of reference of the peripheral patch. In the *X-cue experiment* the cue consisted in a X made of two segments similar to that of the main experiment but being one white and one black. A positive match, thus a HRP trial, occurred if the main axis of the peripheral stimulus coincided with the orientation of the white segment of the X, while a match with its black segment corresponded with a LRP trial. In the *feature-independent cue experiment* the line was replaced by a Gabor patch equal to the target but modulated along the red-green (RG) or the blue-yellow (BY) axes, on a fully random base. The match that cued the reward rate was based on the Gabor's color, while the peripheral stimulus and task were still confined in the orientation domain, hence the feature of the reward cue and that of the reward effective, peripheral task, were dissociated.

Five hundred milliseconds after the offset of the last foveal disk two different *response pages* were shown in sequence, one for the foveal counting task and the other for the peripheral orientation identification and matching task. The counting response page, automatically shown 500 ms after the last disk, displayed the list of digits corresponding to the number of tracked flashes and observers were asked to click within the square patch containing the digit. The orientation response page allowed the orientation discrimination/identification response. It contained Gabor probes representing the entire set of CW and CCW tilts from both the -45° and the $+45^\circ$ axis, (5 tilts \times 2 directions \times 2 axes), and observers were asked to click on the probe that matched more closely the perceived tilt. In the feature-independent cue experiment only one line of CW and CCW probes modulated around vertical were shown.

Finally, at the end of the sequence of each trial the *feedback page* indicated the success of the trial and the accumulation of credit for obtaining the reward (a lottery Scratch-and-Win ticket was awarded for any 20 rewarded trials). The feedback page displayed two bars, a white bar that was elongated if the outcome of the trial led to reward and a black bar that was elongated in the presence of a wrong identification. The white bar was fully elongated, and a ticket donated, after any 20 correct identifications. Unrewarded trials (in the presence of correct discrimination) were signaled by no change in either bars. The change to the bars was clearly visible to each subject. When a Scratch-and-Win ticket was awarded both bars were reset to the initial position.

EXPERIMENTAL CONDITIONS AND PROCEDURE

We executed three main experiments and two baseline conditions. In the main experiment, reported in **Figure 1A**, we have used a reward cue consisting of a single line. The X-cue experiment was identical to the main experiment (in HL mode) except for the use of X-like cues made up of lines at opposite polarities. It was designed to exclude the effects of priming that the orientation of

the cue might exert on peripheral stimuli of close orientation. The feature-independent cue experiment differed in two aspects: the reward cue consisted in one of two combination of colors of the bars of a grating and the main axis of orientation was no vertical, not oblique. Accordingly, the orientation response page of this experiment contained only one set of probes (see **Figure 5**). The baseline conditions consisted in the measure of the peripheral threshold in the absence of counting task and in the main experiment without reward.

REWARD PATTERN

In the rewarded conditions the probability of achieving credit for the Scratch-and-Win ticket was equal to 0.9 (HRP condition) if the main axes of cue and stimulus (in the main experiment), if the orientation of the bright axis of an X (in the X-cue experiment) or if their color (in the feature-independent cue experiment) coincided. In the opposite case reward was granted with a probability of 0.1 (LRP condition). These reward probabilities were conditional to correct orientation discriminations; wrong discriminations gave no reward. In other words, we awarded subjects in the joint presence of (1) correct counting and, (2) correct identification of the direction of tilt off $+45^\circ$ or -45° (CW or CCW), implying no extra gain in the presence of identification of the exact probe in the response page. HRP and LRP trials were fully randomized, thus observers could not predict the reward probability until the peripheral target was shown. Data analysis. The orientation magnitude matching paradigm used here allowed us to analyze the data in two fundamental ways. The first is in terms of binary accuracies, with correct and wrong responses based on the direction of tilt of the clicked probe. For example, if the signal is sampled from the left side of any of the two "fans" of **Figure 1B** (a CCW signal) and the response is a click on any one of the response probes to the left, this will result in a correct response, while a click to any CW probe will result to a wrong response. This allowed us to provide standard psychometric measures, such as thresholds, out of conventional psychometric functions. The second scoring is based on the matching of each physical signal to individual response probes and is achieved by plotting the histogram representing the distribution of reported tilts for each physical signal displayed (Baldassi et al., 2006). We will call the two measures orientation discrimination and orientation identification, respectively. Trials in which counting was wrong were discarded for the main data analysis. Hence, data were analyzed separately for orientation discrimination and identification. Orientation discrimination data formed psychometric functions fitted by cumulative normal cdf. Each function was bootstrapped (Efron and Tibshirani, 1994) and refitted 100 times and the threshold was calculated (75% accuracy of the fitted function) for each bootstrap sample in order to have a reliable estimate of the threshold and its standard error. Orientation identification data fed the behavioral tuning functions (histograms representing the proportion of reported, or perceived tilt in the presence of a given physical tilt). We generated one such function for each physical angle used in the experiment and bootstrapped it 100 times to estimate the reliability of individual points. Each bootstrap sample was fitted with a normal pdf in order to provide a statistically reliable estimate of the Gaussian parameters (μ and σ).

RESULTS

Two important features of our paradigm should be highlighted here. The first is that, because the two reference axes were orthogonal and the least angular distance between the most CW tilt from -45° (i.e., the rightmost probe of the left “fan” of **Figure 1B**) and the more CCW tilt from $+45^\circ$ (i.e., the leftmost probe of the right “fan” of **Figure 1B**) was equal to 26° , there was no confusability between tilts around the two different axes. This was confirmed in all experiments. The second, related to the first, is that because the reward cues are set at neutral orientations, they therefore carry no task-relevant signal and do not provide any cue either for the discrimination or the identification task. It is worth noting that although the timing/counting feature of our experiment implies a broad range of intervals between the peripheral stimulus and the subsequent response, in pilot analyses we found no difference in counting nor in orientation performance when comparing the data of the lowest vs. the highest quartile of durations. In other words none of the results we present below can be attributed to memory effect.

ORIENTATION THRESHOLDS

We measured orientation discrimination thresholds and behavioral tuning functions for three reward levels, two attentional load levels and two different reward cues. **Figure 2** shows average thresholds, i.e., the orientation offset leading to 75% of correct discrimina-

tion responses, for the two types of trials, LRP (left gray points) and LRP (right black points), for the two attentional conditions, LL (circles) and HL (squares), and for the two type of cues, single line (filled symbols) and X-like cue (crossed squares). The two horizontal lines plot average thresholds the peripheral target was displayed without the counting task (lower gray line) and with the dual task but without reward (upper black line) and provide the basic demonstration that the two tasks used here share the same, limited-capacity system (t -test; $p < 0.001$). In all conditions orientation acuity was larger than in standard studies of orientation discrimination, where they typically span around $1\text{--}2^\circ$ (see also the feature-independent cue experiment below). This is simply due to the fact that the reference axes for the discrimination were tilted by 45° , reflecting the so-called *oblique effect* (Campbell et al., 1966), i.e., a rougher and noisier encoding of orientation relative to the horizontal and the vertical axis. All the reward rates, load conditions, and cue types showed significantly lower orientation discrimination thresholds than for the dual task without reward (black horizontal line) that were of about 9° . However, in the presence of reward, average thresholds decreased substantially, span-

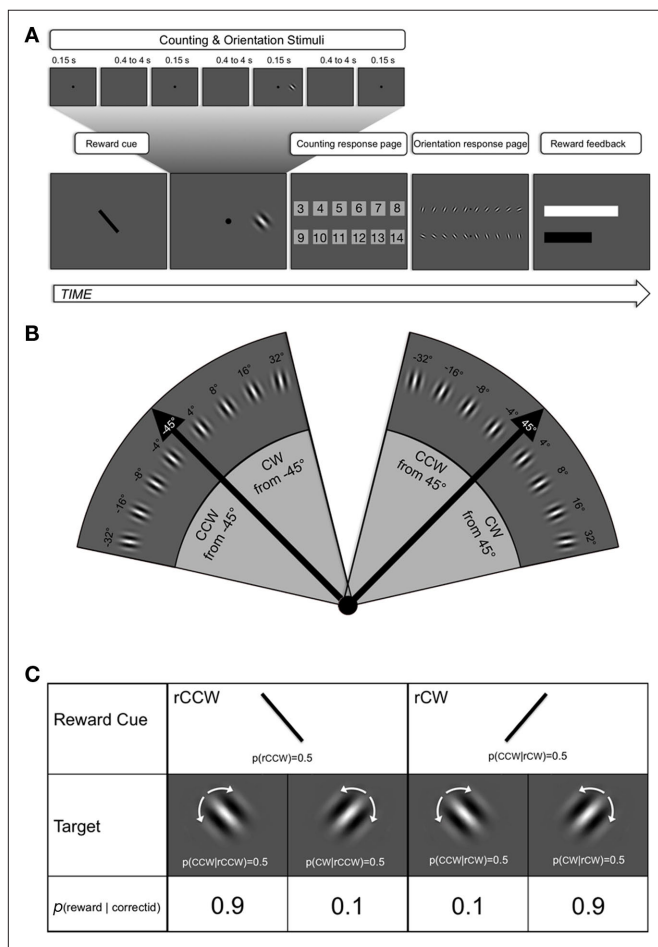


FIGURE 1 | Temporal structure of a trial (A), stimulus space (B), and reward patterns (C). (A) A trial began with a foveal line (subtending 3° of visual angle) displayed for 500 ms and tilted either 45° clockwise (CW) or counterclockwise (CCW) from vertical. Then the central attention loading task started. It consisted in a sequence of 3 to 14 flashes (100 ms, random) of a foveal disks with a random inter-disk interval of 0.4–4 s. Subjects were asked to track the exact number of flashes. During one of the flashes (excluding first and last), the target was shown 7° to the left or to the right of fixation in synchrony with the corresponding disk. It was tilted CW or CCW relative to either 45° or -45° . Then the first response page was displayed; it contained all the digits corresponding to the range of possible disk numbers and subjects had to report to the tracked number of disks with a mouse click. The following display contained the orientation identification and discrimination page. It contained 20 Gabor probes, one for each possible tilt around both the $+45^\circ$ (upper line) and the -45° axis (lower line). The five probes to the left, in each line, corresponded to CW tilts relative to the reference line, while the five to the right corresponded to CCW tilts. Observers had to click on the response probe that best matched the orientation of the peripheral target. After the orientation response, the last page of the trial sequence was shown. It contained a white and a black bar providing feedback about whether or not a trial led to reward, based on a visually salient size increase of the white or the black bar, respectively, and about the amount of rewarded identification needed to achieve another Scratch-and-Win ticket. The white bar was completed, and a ticket donated, after any 20 correct discriminations. (B) “Fan” diagram of the stimulus space. Peripheral targets were oriented Gabor patches whose exact orientation was determined by tilt offset around either a -45° , CCW reference axis (left fan) or a $+45^\circ$ CW reference axis (right fan). The two black arrows in each side represent the two reference axes as well as the two possible cues, one of which was randomly selected and displayed at the beginning of the trial. Notice that (1) the signal was never equal to the axis and, (2) the rightmost item of the left-hand fan is too tilted off-vertical to be confused with the leftmost tilt of the right-hand fan, implying independent coding of the two sets of signals. (C) Different lines of the table indicate, from top to bottom, the probability of each cue type, of each target type given the cue type, and the probability of earning reward given the combination of cues and targets. It has to be clear that: (1) there was an even probability (0.5/0.5) that any of the two cues were shown, (2) there was an even probability (0.5/0.5) that the target was tilted around the -45° or the $+45^\circ$ angle, and this in turn implies that there was no advantage whatsoever in biasing the response toward the cued axis; and finally (3) the probability of earning a reward depended on whether the main axes of cue and target matched or not, according to a 0.9 vs. 0.1 pattern, respectively. Consider that correct counting was the underlying condition for reward, as wrong counting voided the trial, making $p(\text{reward}) = 0$.

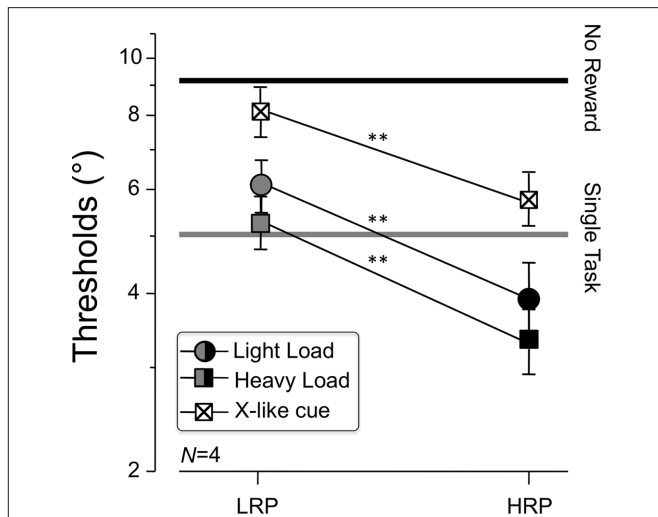


FIGURE 2 | Average orientation discrimination thresholds ($N=4$), corresponding to the 75% correct point of the psychometric function. The points represent the different reward rates (LRP, gray, and HRP, black) and different symbols represent different attentional conditions (light load, circles; heavy load, squares; X-cue, crossed squares). The straight horizontal lines marks the average orientation discrimination threshold for the peripheral target alone in the absence of attentional loading task (gray line, bottom) and for the dual task without reward (black line, top). Plotted data include only the analysis of trials in which the central task was successful (accurate counting). Error bars plot the SEM. The order of conditions (blocks) was shuffled throughout the experiment for all but the Heavy Load condition, executed later as a control experiment, which explains the slight (but not significant) reduction of thresholds (that leaves the pattern of results unaffected). Rewarding correct orientation discrimination responses, though as rarely as in 10% of the cases, sets performance of the main tasks to a level comparable to when there was no central task, whereas highly frequent rewards show an additional advantage of the same magnitude (about a factor of 1.5, $p < 0.01$). In the presence of an X-like cue thresholds are slightly higher for both reward rates, which may be due to a sub-optimal use of the cue. Importantly, the modulation of performance obtained by increasing the reward probability is of the same amount across attentional load and cue types, suggesting that the effect cannot be explained by the use of spare attentional resources allocated to the peripheral task. Notice that the absolute value of threshold is high as discrimination is performed around the oblique axes, where orientation coding is rougher (Campbell et al., 1966).

ning from about 8° – 6° for LRP trials to about 4° – 3° for HRP trials (Figure 2, left vs. right points). It is noticeable that introducing reward to the task, even in 10% of the trials, reduced thresholds substantially, but it is even more surprising that when the reward probability was as high as 90%, perceptual performance was lower than for the peripheral task alone (lower gray horizontal line) for both the LL and the HL condition. Again, differential learning cannot adequately explain these results as all the conditions (except the HL condition that was ran later, as a separate control experiment) were executed in the same block or in different blocks interleaved across conditions. Comparing the two reward rates of our experiment, orientation discrimination thresholds in LRP trials were about 50% higher than in HRP trials (t -test; $p < 0.01$ for LL and X-cue; $p < 0.001$ for HL). This difference was not affected by the attentional load devoted to the central counting task, as shown by the parallel functions of Figure 2, suggesting that the difference between reward rates could not be attributed to spare

attentional resources allocated peripherally in the HRP condition. Indeed, the counting performance did not depend at all on the reward rate, which remained stable at about 95% in the LL and 55% in the HL condition for both the LRP and HRP condition (t -test; $p = 0.769$), ruling out the possibility of response shifts *a posteriori*. Importantly, the 55% rate of correct counting shown in the HL, dual task coincided with the preliminary measures that we took in each observer for the counting task alone, in the absence of peripheral task, implying that this was an absolute limit introduced by the task and that the peripheral task did not shift resources, as otherwise counting performance should have worsened in the dual task. It is noteworthy that this effect was obtained when the reference axis of the peripheral target was tilted in the same direction of the cue, and that it worked also when the orthogonal axis (signaling a LRP trial) was physically part of the cue, in the X-like control experiment. This suggests that the higher likelihood of achieving a reward improved the representation of the cued axes according to a top-down mechanism.

BEHAVIORAL TUNING FUNCTIONS

We then inspected the difference between behavioral tuning functions obtained in different conditions to probe the nature of the mechanism solicited by higher reward rates. In particular, we compared the tuning functions obtained by two of the observers who collected a larger dataset for the purpose of the present analysis (CG and SM) for the target tilts of 4° and 8° , as they are near threshold and are more informative for containing identification errors (Baldassi et al., 2006). Each of the four panels of Figure 3 reports two pairs of behavioral tuning functions, for the LRP and the HRP condition, in gray and black respectively, and for the angle at 4° and 8° (pointed by the small gray arrows), to the left and to the right, respectively. The two observers are reported in the two columns, while the two attentional loads, light and heavy, are reported in the two rows. The bar plots inside each panel plot the σ of the functions according to the same color code and spatial arrangement of the main graphs. The points in each graph show the proportion of responses to each response probe for the physical tilt considered (4° to the left of each panel, 8° to the right), with positive angles reporting correct discrimination (i.e., CW for CW tilts and CCW for CCW tilts) and negative angles indicating wrong discriminations (CW when CCW and *vice versa*). The smooth curves are Gaussian fits to the data-points, continuous black and dashed gray for the HRP and the LRP condition, respectively; they were in all cases describing the data well, with R^2 values of the fit of 0.78 or higher. The main result, clearly evident across observers and conditions, is that a higher likelihood of earning a bonus makes all the curves narrower and sharper, indicating a more reliable representation of the physical angle at the perceptual level. In the LRP condition the range of confusability over the orientation domain was substantially broader, as indicated by the significant differences in the σ of the Gaussian fits (based on a Student's t -test on the bootstrap samples; $p < 0.01$ in all cases except for SM LL angle 8° and GC HL angle 4° , for which $p < 0.05$) observed for all conditions and observers. Importantly, this effect takes place with a comparable strength in both the LL and the HL condition, as confirmed by the bar plots embedded in Figure 3, confirming that we can reduce drastically the possibility that the peripheral

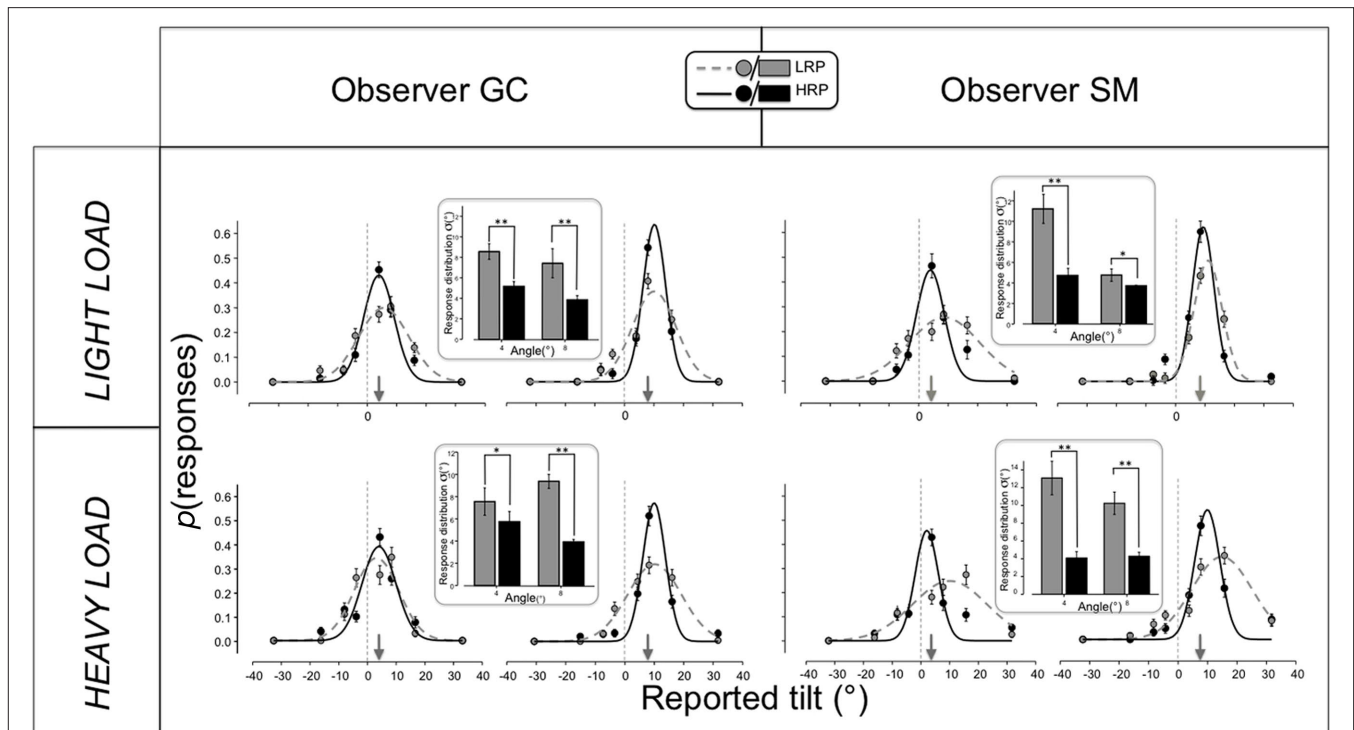


FIGURE 3 | Behavioral tuning functions obtained by two observers (CG, left, and SM, right) in the two attentional load conditions (light load, top, and heavy load, bottom). Each table cells reports two pairs of graphs, for the two physical angles around threshold (4° and 8° , indicated by the gray arrows), that is the histogram of reported angles given an angle of 4° and of 8° , to the left and to the right of each panel, respectively. Each graph plots the functions measured for HRP trials (black symbols) and LRP trials (gray symbols), fitted with Gaussian *pdfs* (straight black and dashed gray, respectively; $R^2 \geq 0.78$). The abscissae report negative angles for reported tilts yielding to errors, i.e., CCW identifications with CW signals and CW identifications for CCW signals collapsed together, and positive angles for correct discriminations. The error bars of each symbol represent the SEM of the estimate calculated by a bootstrap procedure (Efron and Tibshirani, 1994). The framed bar plots show the σ of the Gaussian fits with the error bars representing the SEM of the bootstrap estimates and the asterisks showing the significance level

(* $p < 0.05$; ** $p < 0.01$) of the Student's *t*-test comparing the distributions of bootstrap samples for LRP an HRP trial of individual observers and tilts (4° and 8°). The main effect, coherent with the threshold measurements shown in **Figure 2**, is that the width of the tuning functions of the HRP condition is considerably narrower than the LRP condition in all conditions, as directly shown by the embedded bar plots. This implies a more precise representation of the target's orientation when the task was more likely rewarded. The second effect is that lower reward rates shift the peaks of the functions toward tilt values larger than the physical angle, implying a general non-veridical representation of orientation (usually explained as off-orientation looking Solomon, 2002; Mareschal et al., 2006); however, higher reward rates restore the peaks to more veridical value close to the physical angle of the stimulus. The overall change in both the μ and the σ of the behavioral tuning functions indicates that reward sharpens significantly the internal representation of the orientation of a stimulus.

task can depend on “spare” attentional resources saved from the central task demand and allocated to the peripheral task. In fact, if the results of the LL condition were attributable to leaking of attentional resources, a full load to the central task would have annulled or strongly decreased any difference between LRP and HRP trials. Indeed, wrong counting made the $p(\text{reward}) = 0$, and the counting performance was around 55% for all observers in the HL task (with no difference across reward rates whatsoever); therefore, as confirmed by personal reports, they always had to put a great attentional effort to keep their counting performance as high as they could. The suggestion that reward makes perception more veridical is confirmed, at a visual inspection, by the position of the means (peaks, μ) of the behavioral tuning functions. In the HRP condition, this parameter matches more closely the physical tilt of the stimulus in all cases, but more clearly (and more reliably from a quantitative analysis) in observer SM. The mispositioning of the distribution peaks to tilt values higher than the actual stimulus for the discrimination, is well known in literature as “off-orientation

looking” (Solomon, 2002; Mareschal et al., 2006), i.e., the strategy of relying on orientation channel more tilted than the stimulus to optimize performance in orientation discrimination tasks.

MODEL

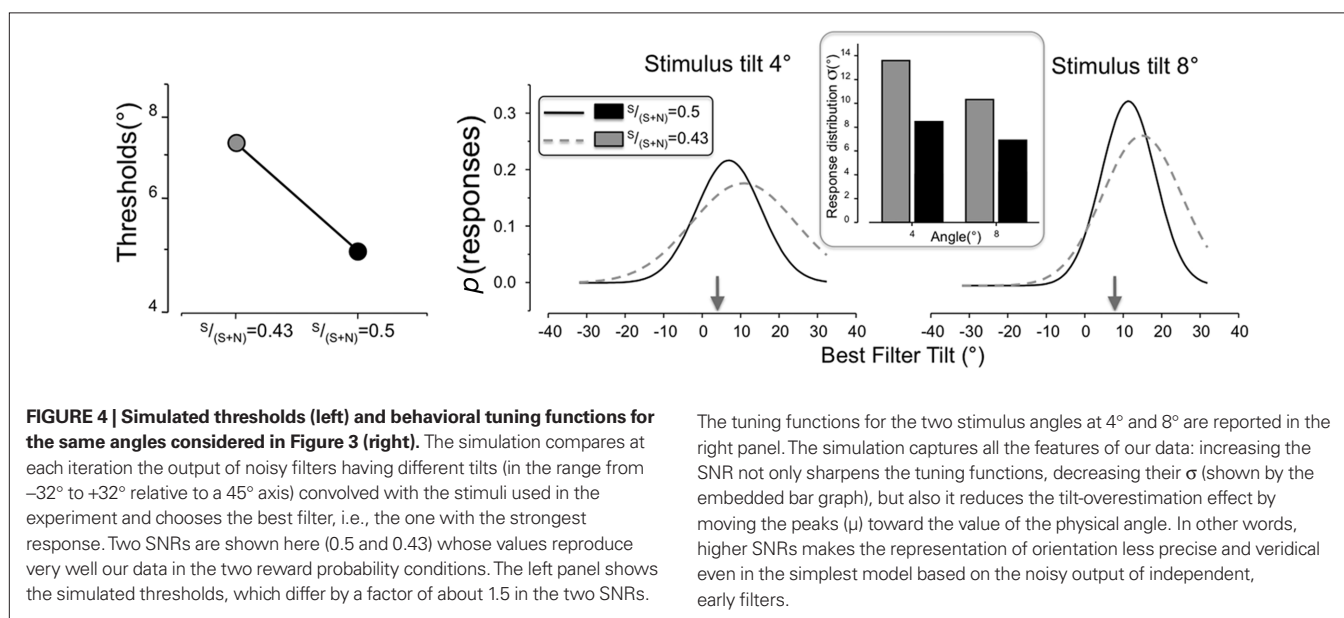
In order to verify the possibility that the mechanism supporting the reward-based modulation of orientation discrimination was a reduction of SNR at an early level, we ran a Monte Carlo simulation using the same stimuli of our experiment that, at each trial, were convolved with a bank of noisy filters of optimal spatial frequency and phase. The filters' set was formed by selecting all the orientations that were used as stimuli and that could be selected in the response page (i.e., 10 tilts from -32° to 32°). Each filter was perturbed by an independent source of noise that was recalculated at each iteration (trial) and whose amount was modulated in different runs. The sum of the squares of each pixel of the convolution matrix was taken as a measure of response of each filter. The filter yielding maximum output in each iteration was taken as the

magnitude matching probe selected at each trial of the real experiment and was used to determine correct and wrong responses in the simulated-orientation discrimination task. If, for example, in a given iteration a stimulus of 4° produced the maximum output in the -8° filter, the latter angle was counted for generating the tuning function and the discrimination response was wrong. We ran 2000 trial for each of the 10 angles and used four SNRs (calculated as $S/S + N$), from 0.5 to 0.35 (where lower numbers imply stronger noise). We reasoned that if our simple SNR hypothesis was correct, then we should be able to reproduce the results of our experiment, i.e., the difference between the LRP and HRP trials could be reproduced by finding two appropriately different SNRs. **Figure 4** shows that this simple simulation reproduced very closely the entire pattern of results, both qualitatively and quantitatively. Thresholds increased from 4.5° to 6.8° when the SNR moved from 0.5 to 0.43. More importantly, the two noise levels reproduced very well the behavioral tuning functions found empirically: decreasing SNR not only increased the σ of the distribution, but it also moved its peak in both the 4° and the 8° angle to tilt values larger than the stimulus tilt. This simulated form of off-orientation looking (Solomon, 2002; Mareschal et al., 2006), is simply due to the fact that in the presence of higher level of noise, it is computationally favorable to solve similar binary tasks by using channels with larger deviations from the reference. Thus, the entire pattern of results of our experiment are well explained by the behavior of a simple model of orientation discrimination/identification whose decision rule is based on the maximum output of a bank of linear, noisy filters tuned to the possible signals.

FEATURE-INDEPENDENT CUE EXPERIMENT

In the main experiment we modulated the attentional load by summoning the observers' attentional resources to a central task with two levels of difficulty and found that different attentional loads did not alter quantitatively nor qualitatively the results. This may imply that the modulatory channels of reward and attention

act independently, even at early processing stages. However, our data could be alternatively explained as an effect of some sort of feature cueing dependent on a priming effect of the reward cue to the subsequent orientation task, independently on the central counting task. In other word, the presence of a cue line tilted at $+45^\circ$ or -45° might have enhanced the representation of angles around that value at the expense of the orthogonal tilts. Thus, in order to rule out more directly this effect, we decided to carry a different experiment in which the feature of the cue and that of the task were independent. In this experiment we cued the reward probability using the association between reward cue and target on color, while the task still required an orientation judgment. The reward cues were Gabor patches modulated around two independent color axes, BY or RG. The structure of the trial matched that of the main experiment and is summarized in **Figure 5A** (see Materials and Methods for details). **Figure 5B** reports average thresholds of four observers (two of which new to the experiment) and shows clearly that even if the cue did not contain any information to prime the processing of orientation signals, a color coincidence between cue and target improved threshold by about 50%, which is in strict consistence with the results of the main experiment. As expected, orientation sensitivity measured around the main axis was much finer than around the oblique axis, because it discounted the oblique effect (Campbell et al., 1966) and, for the same reason, the behavioral tuning functions were sharper. **Figure 5C** reports the σ of the tuning functions, plotted in **Figure 5D**, for two of the four observers (one new to the experiment). A Student's *t*-test comparing the two distributions of bootstrapped functions for LRP (gray bars) and HRP trials showed significant difference in both observers. Again, tuning functions were sharper and the mean was more veridical when the chance of obtaining reward increased, in HRP trials, suggesting that the effect of reward found in this study is not an epi-phenomenon of feature-based attention. However, it has been known that when observers deploy attention to specific features of a visual object all the unattended features of the same



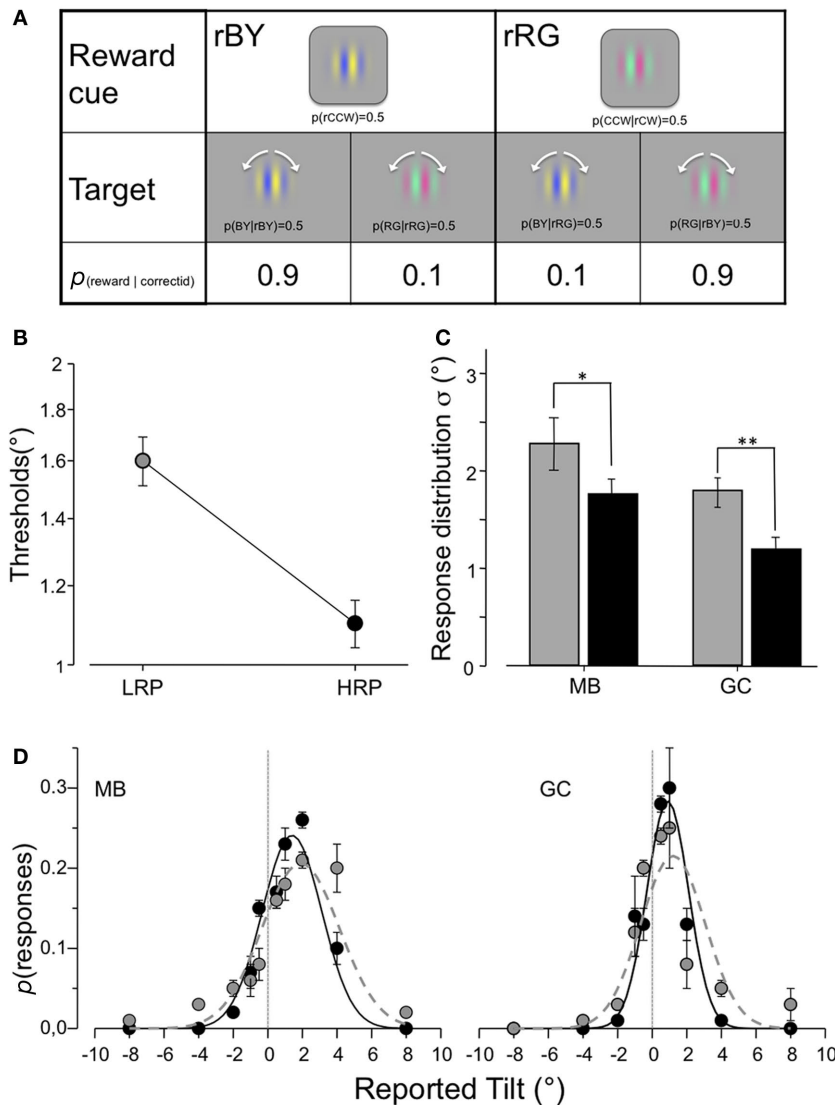


FIGURE 5 | Feature-independent cue experiment procedure (A) and results (B–D). (A) The reward cue was now a Gabor patch modulated either along the blue-yellow (rBY) or along the red-green (rRG) color axis. The task remained an orientation discrimination/identification task of either a color matching (HRP condition) or of a differently colored patch (LRP condition). For both colors, we asked to judge tilt offsets from a unique, vertical reference axis. **(B)** Orientation discrimination thresholds of four observers follow closely the pattern shown by

the main experiment and the model, decreasing by a factor of about 1.5 when the color of the cue and the target gratings matched ($p < 0.001$). **(C,D)** The behavioral tuning functions obtained by two observers (MB and GC) by collapsing the two near-threshold tilts (1° and 2°) confirm the results of the previous experiments showing narrower σ and more veridical μ in the HRP condition. Error bars and reliability of the effects are based on a bootstrap procedure (see Materials and Methods).

object may enjoy attentional priority (e.g., Melcher et al., 2005) and, even though the experimental setup is different, this might explain the advantage observed here. Therefore, we have made a final control in which the reward cue was drawn with both axes ($+45^\circ$ and -45°), resulting in an X-like cue in which one oblique bar was black while the other was white and we have instructed the observers that an HRP trial was signaled by a match of the stimulus axis with the white line of the X-cue, whereas an LRP trial was signaled by a match of the black line of the cue with the stimulus. The polarity of the two axes was randomly established and the results confirmed completely the trend obtained by visually showing only one of the axes, with threshold in LRP trials that were 1.6 to 1.8

times higher than HRP trials in different observers. We think the results of these two control experiments rules out convincingly the idea that the effects measured in this study were due to some sort of implicit priming provided by the tilted cue.

DISCUSSION

In this study we present converging measures to show that the precision of orientation judgment is modulated by the probability that a positive response leads to a reward (in the form of offbeat and cost-efficient Scratch-and-Win lottery tickets). This occurs independently of whether or not attention is engaged elsewhere and even occurs when the reward cue provides absolutely no information

about the response. Perceptual learning or associations that extend over the span of a single trial cannot explain our results as the same stimuli and the same responses could be associated unpredictably with HRP or LRP trials. We have modulated the “motivational” state (Kawagoe et al., 1998) on a trial-to-trial basis and found quick modulations of the perceptual representation of features encoded at an early stage, such as orientation. A possible leakage of attention to the peripheral locations cannot explain the results. Indeed, even though we took measures at different central loads and in different setups (i.e., the control experiments), the positive effect of higher reward rates remained stable at about 50%. If the effect were due to leakage of attention, the LL condition should have shown a much greater effect, but this was not the case.

The reward cue we have used is different from any previous cues used in attentional literature, in that it is not a predictive cue. In other words, it is always neutral, uninformative with respect to the feature that leads to a correct response in the peripheral task. Within the context of Bayesian models of visual performance, it does not affect the prior as feature and location cues do. In the main experiment there are residual possibilities that the cue enhanced the representation and/or the decisional weighting of the class of orientations around the cued axis, leading to better performance. This does not hold for the feature-independent cue, in which the axis of reference is the same for HRP and LRP trials. Another factor that is unlikely to explain the present results is arousal, as the two reward schemes were interleaved within each trial and observers needed to keep their alertness high at least until the peripheral target appeared, as it implicitly signaled the level of reward probability of each given trial. We can also exclude the effect of memory on our data, that is the potential reward-based difference between trials in which there was a long wait between stimuli and responses (as the time between peripheral stimulus and response could be longer than 20 s) and trials with shorter waits. Indeed, preliminary analyses showed that neither HRP nor LRP trials performance showed correlation with the temporal distance between the stimulus and response. All our observers reported that they made a decision on the probe to click on at the time of the presentation of the peripheral stimulus, not at the response page. Moreover, the data of the LRP condition show performances comparable to those obtained in the absence of the central attention task (horizontal line of **Figure 2**). Finally, the counting performance was unaffected by the reward probability (i.e., the same number of counting errors were done within HRP with LRP trials) and dependent only on the central disk contrast, suggesting that alertness was constant across conditions. Importantly, the data were convincingly fit by a model based on the modulation of SNRs of early linear, noisy filters whose individual output is compared with a max rule to make a decision at each trial.

In summary, the results suggest that the reward likelihood may affect the SNR of individual orientation-selective channels at early stages of the visual system independently of attention to the rewarded task or the stimuli. We have used a psychophysical measure that probes orientation coding, an elementary visual function. The primary visual cortex (V1) is a good candidate for such an effect, as most of its cells have orientation-tuned receptive fields (Hubel and Wiesel, 1968; Hubel et al., 1978), it has been evoked to account for psychophysical orientation discrimination

(Regan and Beverley, 1985; Ringach, 1998) and it has been recently found to be modulated by the reward rate in animals (Shuler and Bear, 2006) as well as by the reward history in human observers (Seitz et al., 2009). This is consistent with recent accounts of perceptual learning in psychophysical hyperacuity tasks, which is explained by the action of feedback mechanisms acting on the receptive fields properties of V1 neurons (Fahle, 2004). We have not studied the interocular transfer (as the study by Fahle did), but we have examined the concept of orientation channels in a way consistent with the properties of orientation tuning within the primary visual cortex. Interestingly, a recent, elegant behavioral study that has found effects of training and reward on orientation processing even in unconsciously processed stimuli (Figure 1D in Seitz et al., 2009), the control condition testing orientation processing in untrained eyes and unrewarded orientations exhibited a sensibly higher variability of the psychometric functions, possibly implying lower SNRs. The two sets of results are difficult to compare directly, but in their study this may imply sharper coding for the trained eye and/or rewarded orientations being reflected in the narrower confidence limits of their psychometric functions. If this were true, then we may have tapped into similar reward-dependent early mechanisms of sensory coding. Platt and Glimcher (1999) have observed LIP neurons, with projections that feedback to V1 (Barone et al., 2000), whose levels of activity are positively correlated with the reward value of different stimuli independent of motor factors. The reward value biases also caudate neurons speeding up saccadic latencies (Lauwereyns et al., 2002). It is hence plausible to speculate that similar structures may be involved in our results. However, while these experiments either set a constant association between each stimulus and the amount of reward associated or involve many trials before changing such associations, our experiment overcame this by showing reward effects based on a trial-by-trial, unpredictable coincidence between a cue and the target stimulus. As such the present findings are novel and may open many questions for further investigations on the physiological mechanisms and anatomical circuitries of reward, that until very recently were not assumed to involve primary sensory areas at all (Schultz, 2000). The distinctive feature of our task of relying on trial-wide effects makes it different from recent studies showing reward-based modulation in V1 (Shuler and Bear, 2006; Serences, 2008) or A1 (Beitel et al., 2003). In those cases the modulation depends on the reward history associated to each stimulus, while in our experiment integrating past trials does not provide any additional cue to succeed in the task and earn reward. The direct involvement of early sensory stages within the network of reward-related neuromodulatory activities, and in particular the involvement of dopaminergic activity in our results, may fit with the presence of D1 receptors in the striate cortex (Eickhoff et al., 2007). Fast, phasic responses by dopamine neurons have been found for reward probabilities lower than one, but not when the reward was always acknowledged (Mirenowicz and Schultz, 1994; Schultz, 2000). Further research using similar behavioral paradigms in animals may shed light on this question.

These results provide insight into the basic computations performed by the elementary visual channels involved in such tasks, but some important points will need to be addressed and expanded in future studies. A point to resolve would be to discern whether the modulation of the SNR is due to some form of gain control

of the signal (Carrasco et al., 2004; Reynolds and Heeger, 2009) or to a mechanism of noise reduction (Lu et al., 2002). Our simulation cannot distinguish between these two possibilities, as what we change is the ratio of the signal to the noise. We are currently running new experiments to draw such a distinction seeing how reward affects the contrast (Carrasco et al., 2004) and how external noise impacts on performance across reward and attentional conditions. An additional point that deserves further consideration is whether we actually probed a change at the sensory level, as we proposed earlier in this paper, or whether differential weighting at decisional stages may explain the same effect. Changes of the relative weighting of inputs at the decisional level have successfully explained a number of attentional phenomena (see Eckstein et al., 2009 for a review) in the context of studies relying of predictive cues and tasks to be performed on one out of N signals (with $N > 1$). In this study we use a cue that is not predictive (or, better, it predicts only the rate of reward given correct responses) and a unique peripheral signal. Yet, especially for the main experiment, there is a possibility that the two populations of channels coding orientation values around the two oblique axes may have been weighted (or monitored) differentially. This possibility is less plausible for the feature-independent cue modulation. Another interesting finding lies in the reduction of the “off-orientation looking” effect of orientation discrimination (Solomon, 2002; Mareschal et al., 2006) with high reward rates. It seems that the reward-based modulation makes orientation discrimination more efficient by allowing the use of matched filters (i.e., filters with orientation tuning more ideal for the physical signal) that in neutral conditions would be performing less efficiently because of a negative trade-off between signal and noise associated with this specific task. In other words, we may assume that similar tasks are based on the discrimination between two directions of orientation (CW and CCW) relative to a reference axis are accomplished by comparing the output of channels tuned to tilt in one direction (i.e., CW) with channels tuned to the opposite direction (i.e., CCW; Baldassi and Verghese, 2002). Off-orientation looking would then occur when lower SNRs would cause the behavioral tuning function to “invade” the negative side corresponding to wrong discrimination. If this occurs too often, then the system

mediates by using a channel that is less optimal but more certain about the tilt side. When similar top-down modulations intervene by reducing the spread of the response to the given signal (that is increasing the SNR), the system recognizes the improvement and selects the best matching filter for the orientation discrimination task. It has been argued that most of the findings on perceptual and decisional modulations by reward are contaminated by some form of visual attention, and that reward and attention cannot be easily disentangled empirically (Maunsell, 2004). Attention has also been found to spread to task-irrelevant features if bound to task-relevant features (Melcher et al., 2005), possibly explaining our feature-independent reward experiment, but the two tasks are very different and any feature-binding effect would not rule out our main conclusion. However, as long as attention is operationally defined as the limited amount of resources available to process task-relevant information, being thus withdrawn by more primary tasks (such as our counting task), our study provides novel insight into the mechanisms of reward-based modulation as well as exemplifying a useful methodological template for both single neuron and brain imaging studies aimed at disentangling the two behavioral factors.

CONCLUSION

What are the broad implications of these findings? At a more general level we found that when one’s performance is rewarded, this will not only affect the output of goal-directed behavior, as one would intuitively expect, but it will also improve the quality of the signals on which motor responses are based. To use an analogy, the archer’s shot will succeed not only because of a superior adjustment to his aim, but also because the target is better seen. This in turn has implications for training and education in numerous areas, in particular for competitive sport, where sensory-based performance is fundamental, but momentary motivation may be variable.

ACKNOWLEDGMENTS

This research was supported by the Italian Ministry of Universities and Research (PRIN) and by EC Project “STANIB” (FP7 ERC). We are grateful to David Murphy and Nicoletta Berardi for several suggestions on this manuscript.

REFERENCES

- Alais, D., Morrone, C., and Burr, D. (2006). Separate attentional resources for vision and audition. *Proc. Biol. Sci.* 273, 1339–1345.
- Baldassi, S., Megna, N., and Burr, D. (2006). Visual clutter causes high-magnitude errors. *PLoS Biol.* 4, e56. doi: 10.1371/journal.pbio.0040056
- Baldassi, S., and Verghese, P. (2002). Comparing integration rules in visual search. *J. Vis.* 2, 559–570.
- Barone, P., Batardiere, A., Knoblauch, K., and Kennedy, H. (2000). Laminar distribution of neurons in extrastriate areas projecting to visual areas V1 and V4 correlates with the hierarchical rank and indicates the operation of a distance rule. *J. Neurosci.* 20, 3263–3281.
- Behrens, T., Woolrich, M., Walton, M., and Rushworth, M. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.
- Beitel, R. E., Schreiner, C. E., Cheung, S. W., Wang, X., and Merzenich, M. M. (2003). Reward-dependent plasticity in the primary auditory cortex of adult monkeys trained to discriminate temporally modulated signals. *Proc. Natl. Acad. Sci. U.S.A.* 100, 11070–11075.
- Bradley, A., Skottun, B., Ohzawa, I., Sclar, G., and Freeman, R. (1987). Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior. *J. Neurophysiol.* 57, 755–772.
- Brainard, D. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436.
- Campbell, F. W., Kulikowski, J. J., and Levinson, J. (1966). The effect of orientation on the visual resolution of gratings. *J. Physiol.* 187, 427–436.
- Carrasco, M., Ling, S., and Read, S. (2004). Attention alters appearance. *Nat. Neurosci.* 7, 308–313.
- Della Libera, C., and Chelazzi, L. (2006). Visual selective attention and the effects of monetary rewards. *Psychol. Sci.* 17, 222–227.
- Eckstein, M. P., Peterson, M. F., Pham, B. T., and Droll, J. A. (2009). Statistical decision theory to relate neurons to behavior in the study of covert visual attention. *Vision Res.* 49, 1097–1128.
- Efron, B., and Tibshirani, R. J. (1994). *An Introduction to the Bootstrap* (Monographs on Statistics and Applied Probability). New York: Chapman and Hall/CRC.
- Eickhoff, S. B., Rottschy, C., and Zilles, K. (2007). Laminar distribution and co-distribution of neurotransmitter receptors in early human visual cortex. *Brain Struct. Funct.* 212, 255–267.
- Fahle, M. (2004). Perceptual learning: a case for early selection. *J. Vis.* 4, 879–890.
- Fang, F., Boyaci, H., and Kersten, D. (2009). Border ownership selectivity in human early visual cortex and its modulation by attention. *J. Neurosci.* 29, 460–465.
- Ghose, G., and Maunsell, J. (2008). Spatial summation can explain the attentional modulation of neuronal responses to multiple stimuli in area V4. *J. Neurosci.* 28, 5115–5126.

- Gilbert, C. D., Li, W., and Piech, V. (2009). Perceptual learning and adult cortical plasticity. *J. Physiol.* 587(Pt 12), 2743–2751.
- Glimcher, P., and Rustichini, A. (2004). Neuroeconomics: the consilience of brain and decision. *Science* 306, 447–452.
- Hampton, A., and O’Doherty, J. (2007). Decoding the neural substrates of reward-related decision making with functional MRI. *Proc. Natl. Acad. Sci. U.S.A.* 104, 1377–1382.
- Hubel, D. H., and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195, 215–243.
- Hubel, D. H., Wiesel, T. N., and Stryker, M. P. (1978). Anatomical demonstration of orientation columns in macaque monkey. *J. Comp. Neurol.* 177, 361–380.
- Kamitani, Y., and Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Curr. Biol.* 16, 1096–1102.
- Kawagoe, R., Takikawa, Y., and Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nat. Neurosci.* 1, 411–416.
- Lauwereyns, J., Watanabe, K., Coe, B., and Hikosaka, O. (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature* 418, 413–417.
- Lavie, N. (2005). Distracted and confused? selective attention under load. *Trends Cogn. Sci.* 9, 75–82.
- Lu, Z., Lesmes, L., and Doshier, B. (2002). Spatial attention excludes external noise at the target location. *J. Vis.* 2, 312–323.
- Mareschal, I., Dakin, S., and Bex, P. (2006). Dynamic properties of orientation discrimination assessed by using classification images. *Proc. Natl. Acad. Sci. U.S.A.* 103, 5131–5136.
- Maunsell, J. (2004). Neuronal representations of cognitive state: reward or attention? *Trends Cogn. Sci.* 8, 261–265.
- McAdams, C., and Maunsell, J. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* 19, 431–441.
- McAlonan, K., Cavanaugh, J., and Wurtz, R. H. (2006). Attentional modulation of thalamic reticular neurons. *J. Neurosci.* 26, 4444–4450.
- McAlonan, K., Cavanaugh, J., and Wurtz, R. H. (2008). Guarding the gateway to cortex with attention in visual thalamus. *Nature* 456, 391–394.
- Melcher, D., Papathomas, T. V., and Vidnyanszky, Z. (2005). Implicit attentional selection of bound visual features. *Neuron* 46, 723–729.
- Mirenowicz, J., and Schultz, W. (1994). Importance of unpredictability for reward responses in primate dopamine neurons. *J. Neurophysiol.* 72, 1024–1027.
- Peck, C. J., Jangraw, D. C., Suzuki, M., Efeim, R., and Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex. *J. Neurosci.* 29, 11182–11191.
- Pelli, D. (1997). The Video Toolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442.
- Platt, M. L., and Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Regan, D., and Beverley, K. (1985). Postadaptation orientation discrimination. *J. Opt. Soc. Am. A* 2, 147–155.
- Reynolds, J. H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *J. Neurosci.* 19, 1736–1753.
- Reynolds, J. H., and Heeger, D. J. (2009). The normalization model of attention. *Neuron* 61, 165–185.
- Ringach, D. (1998). Tuning of orientation detectors in human vision. *Vision Res.* 38, 963–972.
- Schultz, W. (2000). Multiple reward signals in the brain. *Nat. Rev. Neurosci.* 1, 199–207.
- Schultz, W., Tremblay, L., and Hollerman, J. (2000). Reward processing in primate orbitofrontal cortex and basal ganglia. *Cereb. Cortex* 10, 272–284.
- Seitz, A. R., Kim, D., and Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron* 61, 700–707.
- Serences, J. T. (2008). Value-based modulations in human visual cortex. *Neuron* 60, 1169–1181.
- Shuler, M., and Bear, M. (2006). Reward timing in the primary visual cortex. *Science* 311, 1606–1609.
- Solomon, J. (2002). Noise reveals visual mechanisms of detection and discrimination. *J. Vis.* 2, 105–120.
- Watanabe, T., Harner, A. M., Miyauchi, S., Sasaki, Y., Nielsen, M., Palomo, D., and Mukai, I. (1998a). Task-dependent influences of attention on the activation of human primary visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* 95, 11489–11492.
- Watanabe, T., Sasaki, Y., Miyauchi, S., Putz, B., Fujimaki, N., Nielsen, M., Takino, R., and Miyakawa, S. (1998b). Attention-regulated activity in human primary visual cortex. *J. Neurophysiol.* 79, 2218–2221.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 13 December 2010; accepted: 21 January 2011; published online: 08 February 2011.

Citation: Baldassi S and Simoncini C (2011) Reward sharpens orientation coding independently of attention. *Front. Neurosci.* 5:13. doi: 10.3389/fnins.2011.00013

This article was submitted to *Frontiers in Decision Neuroscience*, a specialty of *Frontiers in Neuroscience*.

Copyright © 2011 Baldassi and Simoncini. This is an open-access article subject to an exclusive license agreement between the authors and Frontiers Media SA, which permits unrestricted use, distribution, and reproduction in any medium, provided the original authors and source are credited.