

Rewards Evoke Learning of Unconsciously Processed Visual Stimuli in Adult Humans

Aaron R. Seitz,^{1,2,3,*} Dongho Kim,^{1,3} and Takeo Watanabe¹

¹Department of Psychology, Boston University, 64 Cummington Street, Boston, MA 02215, USA

²Department of Psychology, University of California, Riverside, 900 University Avenue, Riverside, CA 92521, USA

³These authors contributed equally to this work

*Correspondence: aseitz@ucr.edu

DOI 10.1016/j.neuron.2009.01.016

SUMMARY

The study of human learning is complicated by the myriad of processing elements involved in conducting any behavioral task. In the case of visual perceptual learning, there has been significant controversy regarding the task processes that guide the formation of this learning. However, there is a developing consensus that top-down, task-related factors are required for such learning to take place. Here we challenge this idea by use of a novel procedure in which human participants, who were deprived of food and water, passively viewed visual stimuli while receiving occasional drops of water as rewards. Visual orientation stimuli, which were temporally paired with the liquid rewards, were viewed monocularly and rendered imperceptible by continuously flashing contour-rich patterns to the other eye. Results show that visual learning can be formed in human adults through stimulus-reward pairing in the absence of a task and without awareness of the stimulus presentation or reward contingencies.

INTRODUCTION

A difficulty in understanding the processes of skill learning is that skills are typically learned in the context of behavioral training paradigms that involve a variety of perceptual, decisional, and motor processes. In the field of visual perceptual learning, researchers are often concerned with studying aspects of plasticity that are thought to take place within the early visual system, but observations along these lines are typically confounded by the contribution of high-level factors in this learning. A simple reason for this can be found in a typical definition of perceptual learning as “performance improvements on a perceptual task after training.” This “training” almost uniformly involves performing a task that is the same as or resembles that used to evaluating the learning. Given this constraint there have been great debates regarding the processes that gate perceptual learning.

A prevailing hypothesis was that for a feature to be learned, focused attention must be directed to that feature. Along these lines, a number of studies have found perceptual learning for task-relevant features (i.e., features that must be attended to

successfully perform the task), whereas exposed task-irrelevant features (i.e., features that convey no useful information to that task) showed no or only a very limited amount of sensitivity change (Ahissar and Hochstein, 1993; Schoups et al., 2001; Shiu and Pashler, 1992). These results have led to the conclusion that perceptual learning will not occur to stimuli that are task irrelevant and unattended.

However, recent research of task-irrelevant perceptual learning (TIPL) has demonstrated that task-irrelevant stimulus features can be learned when they are presented in temporal conjunction with task-relevant features (Ludwig and Skrandies, 2002; Nishina et al., 2007; Seitz et al., 2005a, 2005b, 2005c, 2006; Seitz and Watanabe, 2003, 2005; Watanabe et al., 2001, 2002). These studies demonstrate that directed attention is not necessary for perceptual learning to occur and have argued that perceptual learning is gated by reinforcement processes (for reviews see Seitz and Watanabe, 2005; Seitz and Dinse, 2007).

However, a commonality of studies of attentional learning and those of TIPL is that recently, a seeming consensus has been reached in the conclusion that learning is gated by top-down, task-related factors (Ahissar and Hochstein, 1993; Doshier and Lu, 1998; Dupuis-Roy and Gosselin, 2007; Li et al., 2004; Polley et al., 2006; Seitz et al., 2005a; Seitz and Watanabe, 2005; Shiu and Pashler, 1992). For example, even in the case of TIPL, learning occurs for stimuli that are correlated with task targets (Seitz and Watanabe, 2003), but does not occur when targets are not fully processed (such as in the attentional blink) (Seitz et al., 2005a). Similarly, studies of auditory learning have concluded that reward-based learning also requires high-level gating (Polley et al., 2006). The main difference between TIPL and attentional learning frameworks is that in the case of attentional learning, high-level gating processes select *which* features are to be learned, whereas TIPL gating selects *when* learning should occur (Seitz and Watanabe, 2005).

These studies raise the question of whether visual skill learning requires an active, goal-directed process or whether learning can occur automatically without any task, stimulus awareness, or goal-directed behavior. While this may seem like an extreme question, it is worth investigating given that in other forms of learning, such as of classical conditioning, learning can take place simply through pairing stimuli with reinforcers, even without awareness (Morris et al., 1998; Pessiglione et al., 2008). The model of perceptual learning proposed by Seitz and Watanabe suggests that successful performance of a task works

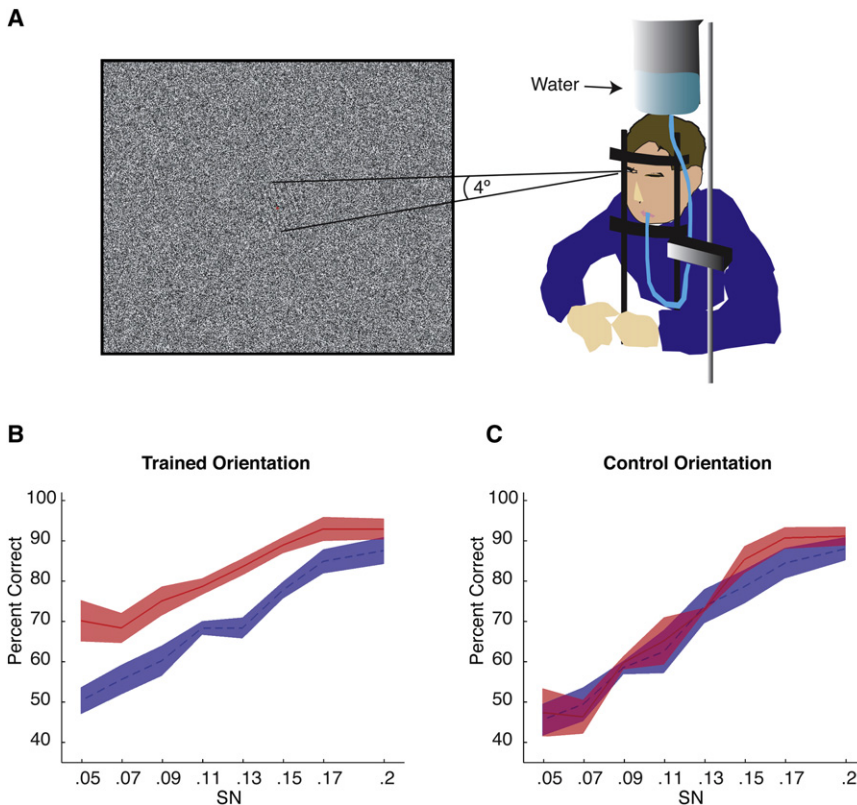


Figure 1. General Procedure and Results of Experiment 1

(A) Participants received occasional drops of water as rewards while passively viewing visual orientation stimuli (20% signal; 2 cycles/°; 4° diameter) without any task (see Figure S2 for task schematics). (B) Results for trained orientation. Blue curve indicates psychometric function from the first sensitivity test, and the red curve, that of the second test; shaded regions indicate standard error. A clear learning effect is evident by the separation of the curves. (C) Results for untrained orientation indicate no reliable change in performance.

as an internal reward and triggers neuromodulatory learning signals such as acetylcholine, norepinephrine, and dopamine, which gate learning and thus restrict sensory plasticity (Seitz and Watanabe, 2005). The learning signals provide a likely mechanism of learning since they are triggered by important events in behavior and are released relatively diffusively throughout the brain (Dalley et al., 2001; Schultz, 2000). However, studies show that these same learning signals are released outside the context of a task, including at times of unpredicted rewards and in the context of classical conditioning (Dayan and Balleine, 2002; O'Doherty et al., 2006; Pearce and Bouton, 2001; Schultz et al., 1997; Thiel et al., 2002; Yu and Dayan, 2005). Accordingly, one might predict that visual learning should occur for stimuli that are temporally paired with rewards, even outside the context of a task.

Here we take the task out of perceptual learning by use of a classical conditioning procedure in which human subjects, who were deprived of food and water, passively viewed visual stimuli while receiving occasional drops of water as rewards (Dorris and Glimcher, 2004; Lauwereyns et al., 2002; Leon and Shadlen, 1999). The advantage of this procedure is that it allows us to address the mechanisms that gate visual learning and reduce the contribution of decision-stage and response learning effects. In the first experiment, subjects passively viewed visual orientation stimuli that were temporally paired with the liquid rewards. In the second experiment, these orientation stimuli were viewed monocularly and rendered imperceptible by continuously flashing contour-rich patterns to the other eye (Tsuchiya and Koch, 2005). Results show that visual learning can be

formed in human adults through stimulus-reward pairing in the absence of a task and without awareness of the stimulus presentation or reward contingencies. Furthermore, these learning effects were specific to the eye to which the stimuli were presented, a hallmark of early visual processing. These results demonstrate that visual learning in humans can be driven by reward signals and that this learning can take place

RESULTS

in an automatic way and can affect early stages of visual processing.

The basic design of this study is typical of procedures used to study visual learning. Human participants were first tested on their sensitivity to the oriented sinusoidal gratings (the stimulus for which we are evaluating learning), then underwent a “training stage” in which a specific angle of orientation was “trained,” and finally underwent a second testing session that measured possible changes in sensitivity for the trained and control (i.e., untrained) orientations. Changes in performance between the test sessions that are specific to the trained orientation are the signature of visual learning (Fahle, 2005). The key innovation in the present study is that the participants were given no task during training. Instead, they were instructed to passively view the computer monitor, maintain gaze on a central fixation spot, and enjoy the occasional drop of water that was delivered through a tube that was placed in their mouths (Figure 1). The “trained orientation” refers to the orientation that was paired in a consistent temporal relationship (i.e., preceded and partially overlapped) with the reward (i.e., the drop of water) and the “control orientation” was a second orientation that was presented equally often as the trained orientation, but was never temporally paired with reward. To ensure that the drops of water were in fact rewarding, participants were required to refrain from eating or drinking for 5 hr prior to each experimental session. The use of this procedure allowed us to take the task




Figure 2. Training Procedure for Experiment 2

Alternating blocks of 15 s duration CFS stimuli were presented to each eye. The CFS stimuli consisted of a sequence of full-screen textured pattern images that were presented at a rate of 10 Hz. For the trained eye a sequence of 2 Hz noise and orientation patterns were presented while CFS was shown to the other eye. For the untrained eye, a gray screen was presented while CFS was presented to the trained eye. After every 5 min, participants took a 3 min break. These sequences repeated eight times in each of the 20 training sessions.

out of perceptual learning and to examine the specific hypothesis that reward-related learning signals are sufficient to cause improvements in visual sensitivity for visual stimuli paired with rewards.

In the first experiment, participants underwent 9 days of this training procedure (see [Experimental Procedures](#)). Given that there was no training task, participants' performance is evaluated through a comparison of the ability of participants to discriminate the trained and control orientations in the sensitivity tests conducted before and after training. To make discrimination challenging the orientation patterns were degraded with spatial noise, and we generated a psychometric function ([Figure 1B](#)) that describes participants' discrimination performance as a function of the signal-to-noise ratio (SN). With this type of graph, learning can be determined by whether the psychometric function changes between the sensitivity tests. Such a learning effect can be seen in [Figure 1B](#), where significant learning was found for the trained orientation that was paired with reward [$F(1,3) = 41.1, p < 0.001$, two-way repeated-measures ANOVA]. However, no significant learning effect was observed for the control orientation [[Figure 1C](#), right; $F(1,3) = 2.3, p = 0.23$]; results are similar for d' and reaction time ([Figure S3](#) available online).

These results demonstrate that visual learning can occur through reward pairing in the absence of a task and are consistent with the hypothesis that visual learning occurs automatically based upon stimulus-reward contingencies of which the subjects were not aware. However, the results do not rule out

the possibility that learning was due to participants directing more attention to the trained than the control orientations. During the training phase, the orientation stimuli were presented at the SN level of 0.2, which was subtle, but for which the orientation stimuli can be reliably discriminated. Participants may have determined the orientation of the trained and control orientations and discovered the stimulus-reward contingencies. Given that directed attention is well recognized as a factor that drives visual learning ([Ahissar and Hochstein, 1993](#); [Seitz and Watanabe, 2005](#); [Seitz and Dinse, 2007](#))

we designed a new experiment in which orientation stimuli were made imperceptible to ensure that participants could not discover the stimulus-reward contingencies.

In this second experiment, we used the technique of continuous flash suppression (CFS) ([Tsuchiya and Koch, 2005](#)) to render the orientation stimuli imperceptible throughout the training, and the orientation stimuli were only presented in a single eye (trained eye). CFS is a robust type of binocular rivalry in which a series of bright, contour-rich patterns are continuously flashed to one eye while a less salient image is presented to the other eye ([Figure 2](#)). In this circumstance, where the inputs to the two eyes do not match, the eye receiving the stronger input (here the contour-rich patterns) dominates perception and suppresses the perception of the image presented to the eye receiving the weaker input. Thus by using CFS we were able to eliminate participants' awareness of the stimuli (and thereby the stimulus-reward contingencies) during training while largely preserving the responsiveness of visual areas to the orientation stimuli ([Leopold and Logothetis, 1996](#); [Logothetis et al., 1996](#)).

In Experiment 2, participants underwent 20 days of training, and a comparison of performance between the first and last sensitivity tests revealed a similar pattern of learning to that found in the first experiment. A significant effect of learning was found for the trained orientation [[Figure 3A](#); $F(1,3) = 18.6, p < 0.01$]. However, no learning was observed for control orientation [[Figure 3B](#); $F(1,3) = 0.30, p = 0.62$; see [Figure S5](#) for d' and reaction time data]. These results confirm that perceptual learning can be

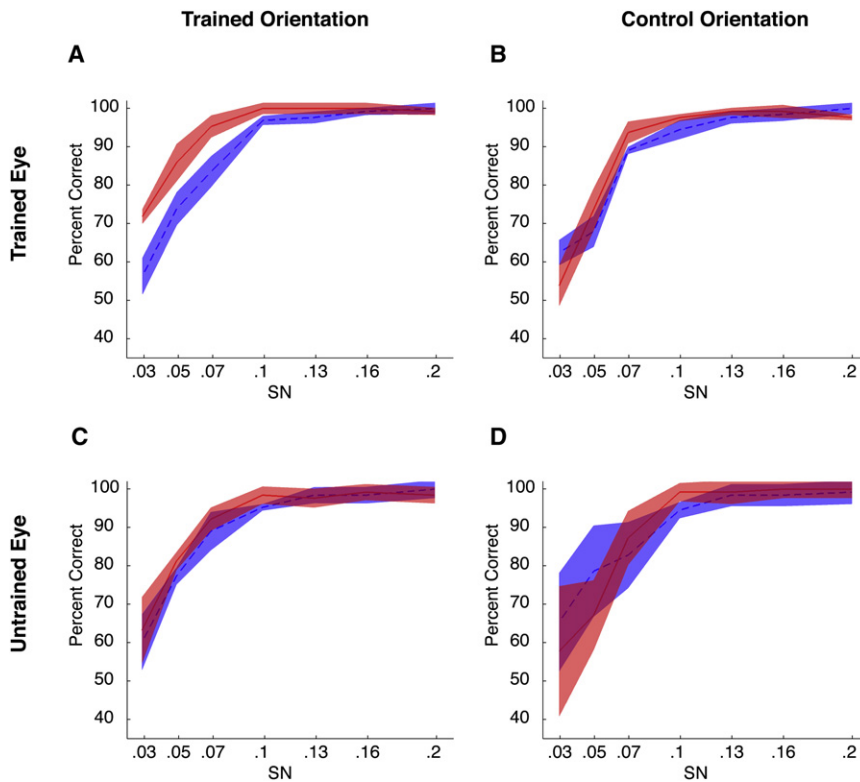


Figure 3. Results of Experiment 2

(A) Results for trained orientation. Blue curve indicates psychometric function from the first sensitivity test, and the red curve, that of the final test; shaded regions indicate standard error. A clear learning effect is evident by the separation of the curves. (B) Results for untrained orientation indicate no reliable change in performance. (C and D) Results for untrained eye showed no reliable change in performance for either the trained (C) or control (D) orientations.

subjects from being aware of the orientation stimuli during the training phase of the experiment.

An important aspect of this study is that participants were trained with a stimulus that was only presented to one of their eyes (Figures 3C and 3D). This enabled us to examine whether learning in the trained eye transferred to sensitivity to the same stimuli presented to the untrained eye. The specificity of the learning effects can provide important clues to what stages of visual processing underlie the observed learning effects;

driven by reward signals, even in the absence of awareness of the learned stimuli and stimulus-reward contingencies.

To evaluate whether participants may have been aware of the presentation of the orientation stimuli or of the stimulus-reward contingencies during training, we conducted a test of stimulus awareness in a separate session after the final sensitivity test was complete. In this test, participants viewed the same stimulus sequence that had been used in the training stage and were asked to report the orientation of the stimuli presented to the suppressed eye. The results showed that participants failed to respond to the vast majority of the stimulus presentations, and that when participants did respond, there were no significant differences between the number of correct ($11.9\% \pm 6.0\%$) and incorrect ($10.1\% \pm 5.0\%$) responses. This posttest should be a conservative estimate of suppression, given that perceptual switches under condition of binocular rivalry occur with increasing frequency with long-term exposure (Suzuki and Graboweky, 2007). However, we ran another awareness test in a group of eight naive subjects who hadn't undergone training. In this awareness test, we presented rewards in some trials and in other trials presented only noise patterns. Subjects were asked to report when they observed orientation stimuli. Results showed no differences of response patterns between trials containing an Orientation + Reward ($4.9\% \pm 4.5\%$), Orientation + no-Reward ($5.2\% \pm 4.8\%$), Reward + no-Orientation ($5.1\% \pm 5.0\%$), or no-Reward + no-Orientation ($5.1\% \pm 4.5\%$). Furthermore, in debriefings, subjects from the training sessions showed no indication of having noticed the orientation stimuli or the stimulus-reward contingency during the training sessions. These results indicate that the CFS procedure was effective in preventing

ocular-specific learning has been argued to be evidence of plasticity of early, monocular stages of visual processing (Fahle et al., 1995; Karni and Sagi, 1991; Lu et al., 2005). Analysis of the sensitivity tests for the untrained eye showed that there were no significant performance changes for either the trained orientation [$F(1,3) = 3.4$, $p = 0.16$] or the control orientation [$F(1,3) = 0.34$, $p = 0.60$] and no effects in either d' or reaction time (Figure S6). Given that there was a slight trend of a learning effect for the trained direction in the untrained eye, we calculated an index of interocular transfer ($100 \times [\text{untrained eye learning}]/[\text{trained eye learning}]$) to evaluate more directly the degree of ocular specificity of the learning. This index showed that for the four subjects, 53.9%, -34.4% , 23.5%, and 28.3% (respectively) of the learning effect for the trained orientation in the trained eye transferred to the untrained eye; note that the negative value in subject 2 indicates that performance in the untrained eye was worse in the posttest than it was in the pretest. The fact that the learning effect is largely specific to both the stimulus orientation and the eye of training is indicative of plasticity involving an early, monocular stage of visual processing (Nishida et al., 1994; Paradiso et al., 1989; von der Heydt et al., 1984), although some learning may be occurring at later processing stages as well (Ahissar and Hochstein, 1997).

While these experiments suggest that a reward process is involved in this learning, it is possible that the results are due to mechanisms other than reward. To test this we ran two new experiments. In one experiment, we asked subjects to rate the pleasantness (on a scale of -2 [very unpleasant] to 2 [very pleasant]) of the water delivered during the experiment in one session after they had been deprived of food and water and in

another session in which they had not been deprived (order of sessions counterbalanced across subjects). For the 14 subjects who participated in this test, we found that subjects rated water to be pleasant (1.4 ± 0.23) after deprivation, but did not do so without deprivation (0.07 ± 0.16), when asked at the end of an hour-long session that resembled the training task from Experiment 2; this difference in favorability between the deprived and nondeprived conditions was significant ($p < 0.01$, two-tailed paired *t* test). In a second control experiment, we had eight new subjects participate in a replication of Experiment 2; however, these subjects were not deprived of food and water. Given that we did not find water to be rewarding without deprivation, we did not expect to find perceptual learning to result from this experiment. As hypothesized we failed to find learning for the orientation that was paired with water in the absence of deprivation [$F(1,7) = 2.2$, $p = 0.20$]. While we acknowledge that with a null effect we cannot prove that no learning occurred, it is notable that we failed to find learning even with twice the power (i.e., eight subjects rather than four) of Experiments 1 and 2. And although we cannot say that no learning occurred in the nondeprived condition, at least the results indicate that a higher reward value (such as water with deprivation) induces stronger learning. These results demonstrate that water was rewarding to subjects after deprivation and that learning did not occur to the same extent in a condition where the water delivery was not rewarding. Together, these results confirm our hypothesis that the learning is indeed related to reward.

Another question regarding the reported results is how reward impacts the processing of the grating stimuli during training. Did the water delivery reward the presentation of the gratings, or did it interrupt the effectiveness of the masks that immediately follow the gratings? To discriminate between these possibilities, we ran an additional control experiment in which a group of 20 new subjects were water deprived and then tested on grating orientation discrimination under two (interleaved) trial types (Reward and no-Reward; in Reward trials a drop of water was given 400 ms after the presentation of the grating). One group of 10 subjects participated in the Mask trials, where subjects reported the orientation of the gratings when the gratings were preceded and followed by a noise mask (as was done in the tests of Experiments 1 and 2). The other group of 10 subjects participated in the no-Mask condition in which a gray screen was presented instead of the noise mask that normally followed the grating presentation. Performance was highly similar among the four conditions (Mask-no-Reward = $80.2\% \pm 1.3\%$; no-Mask-no-Reward = $80.4\% \pm 1.5\%$; Mask-Reward = $78.7\% \pm 2.7\%$; no-Mask-Reward = $77.4\% \pm 2.3\%$). An ANOVA showed no main effects of Reward [$F(1,9) = 0.42$, $p = 0.53$] or Masking [$F(1,9) = 0.71$, $p = 0.41$], nor an interaction between them [$F(1,9) = 0.90$, $p = 0.36$]. These results indicate that there was no significant effect for the mask. This may be because the grating stimuli were presented for 500 ms before the mask onset, which allowed plenty of time for them to be processed before the onset of the mask.

DISCUSSION

Our results challenge current theories of perceptual learning that presuppose that learning requires top-down gating processes

such as attention (Ahissar and Hochstein, 1993) or task-related reinforcement signals (Herzog and Fahle, 1997; Petrov et al., 2006; Roelfsema and van Ooyen, 2005; Seitz et al., 2005a; Seitz and Watanabe, 2005). We show that stimulus-reward pairing is sufficient to cause learning even in the absence of awareness of the learned stimuli or stimulus-reward contingencies. The results are consistent with a process that gates learning that may originate from subcortical reward systems (Schultz et al., 1997; Yu and Dayan, 2005). We suggest that visual skill learning is generally an automatic (i.e. unconscious) process and that goal-directed factors such as directed attention serve mostly to bias how learning takes place rather than actually gating the learning process. We acknowledge that this conjecture is contentious and further studies will be required to more fully test this.

While the methods used in this experiment are similar to those used in conditioning, it is important to realize that the aspect of learning that is being described here is different from that discussed in the typical studies of conditioning. In contrast to classical conditioning, learning in perceptual processes does not measure a stimulus-response association but instead measures a change in sensitivity to the stimulus. It is thus important to verify that the observed learning effects are resultant from perpetual learning rather than a learning of decision strategies or response patterns. While our procedure was designed to minimize the contribution of task-related learning due to the training task (we had no training task), it is still important to account for the possibility that the observed learning effects are resultant from a learned bias. For example, we have previously demonstrated that a component of TIPL can be explained by a learned perceptual bias, which could be considered a conditioned visual response for preferentially seeing the trained stimulus (Seitz et al., 2005b). We argue that this is an unlikely explanation of the present results given that our d' analysis (see Figures S3 and S5) shows that after training, subjects are more sensitive in discriminating between the two orientations (this rules out a response bias to choose a trained orientation or a simple perceptual bias to be more likely to see the trained orientation). Furthermore, the fact that the increase of sensitivity was accompanied by an increase in accuracy specific for the rewarded orientation is consistent with our suggestion that the learning effect is specific to the trained orientation. Also, the fact that learning in Experiment 2 is ocular specific further demonstrates that these results are not simply due to a bias of the trained orientation, which should be equally evident for both eyes. We also note that a bias to report the trained orientation is in line with a previous study of TIPL in which we found both bias and sensitivity changes for the trained stimulus; however, that bias was best explained as a perceptual effect and was consistent with a conditioned response to "see" the paired stimulus even when none was presented (Seitz et al., 2005b). Thus, while we cannot rule out some contribution of bias, all together, it is very reasonable to conclude that our results reflect that there is plasticity in the visual processing of the trained orientation.

The fact that we found a high degree of ocular specificity in the masking experiment is not surprising given the design of this experiment. The CFS procedure is specifically designed so

that the activity of most binocular neurons is dominated by the contour-rich patterns rather than the orientation patterns. Given this, little correlation is expected between the activity of binocular neurons and the delivery of the liquid reinforcers. On the other hand, cells that are monocular will remain responsive to the perceptually suppressed orientation stimuli and show a more consistent pattern of responses in correlation with reward delivery. While it is intriguing to conclude that the learning effect is taking place in V1, the fact that some degree of ocular specificity remains in higher visual areas (although at lower incidence; Uka et al., 2000) makes this conclusion premature. However, physiological studies conducted with a similar procedure in awake, behaving macaques indicate that the learning effect is taking place at V4 or earlier (E. Franko, A. Seitz, and R. Vogels, 2006, Soc. for Neurosci., abstract; and E. Franko, A. Seitz, and R. Vogels, 2007, Soc. for Neurosci., abstract). Also, it is hard to predict the degree of ocular specificity that would have been expected without suppression. Still, our argument that early-stage, ocular-specific neurons are impacted by reinforcement signals remains valid.

While we discuss the observed learning effects as resulting from a reinforcement process, we use this term loosely and we note that reward, and more specifically dopamine, is only one candidate mechanism that may underlie the learning. We have previously discussed that a variety of neuromodulators (such as acetylcholine, norepinephrine, dopamine, etc.) have the general properties that we expect in a learning signal; they are released relatively diffusely throughout the brain in correlation with behaviorally relevant events (Dayan and Balleine, 2002; O'Doherty et al., 2006; Pearce and Bouton, 2001; Schultz et al., 1997; Seitz and Watanabe, 2005; Thiel et al., 2002; Yu and Dayan, 2005). Additionally, while we rule out the possibility that attention is directed to the learned stimuli during training, we acknowledge that the observed learning effects are consistent with types of attention that are not stimulus directed (Fan et al., 2002; Posner and Petersen, 1990). Within this framework there can be a commonality between the learning signals hypothesized to underlie perceptual learning and those thought to underlie some aspects of attention (Seitz and Watanabe, 2005).

Our reward-learning technique is a novel methodology for studying human visual learning and helps to overcome limitations of previous studies that have relied on explanations of learning based upon complicated, and not well understood, psychological processes (Ahissar and Hochstein, 1993; Herzog and Fahle, 1999; Seitz and Watanabe, 2005; Seitz and Dinse, 2007). The benefit of a liquid reward is that it yields a well-characterized physiological response (Mitz, 2005; Schultz, 2006) and allows for a close comparison with physiological studies in primates, which also rely upon liquid rewards. Furthermore, an increasing number of studies are characterizing effects of liquid rewards/punishment on human brain processing (Gottfried et al., 2003; O'Doherty et al., 2002, 2006). The research presented here provides important clarity regarding the mechanisms that guide learning in the adult brain and introduces a methodology by which processing components involved in learning can be more easily dissociated from the myriad of processing elements involved in conducting a behavioral task.

EXPERIMENTAL PROCEDURES

Participants

A total of eight participants (aged between 19 and 35) were used in the main experiments; four participants (two male and two female) in Experiment 1 and four participants (two male and two female) in Experiment 2. The experiments were conducted in accordance with the IRB approved by the Committee on Human Research of Boston University and the Declaration of Helsinki.

Experiment 1

Apparatus

The stimuli were presented using Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) for MATLAB® (The MathWorks, Natick, MA) on a Macintosh G5 computer. The stimuli appeared on a 19" CRT monitor with a resolution of 1024 by 768 pixels and a refresh rate of 85 Hz. The viewing distance was 25 cm. A chin rest was used to maintain the participants' head position. The participants used a computer keyboard to make responses. Water was delivered using a ValveLink®8.2 system made by Automate Scientific, Inc.

Stimuli

Oriented sinusoidal gratings (112.5° or 22.5° ; 2 cycle/°; 4° diameter; luminance $\sim 50 \pm 50$ cd/m²) were presented at the center of the screen and were spatially masked by noise. In the sensitivity tests, the SN varied from trial to trial (0.05, 0.07, 0.09, 0.11, 0.13, 0.15, 0.17, and 0.2; see Figure S1 for example of stimuli). In the training sessions a constant SN of 0.2 was used. The variable SN stimuli were created by randomly choosing the SN proportion of pixels from the orientation image and 1 – SN proportion of pixels from a noise image. The screen area outside of the 4° stimulus region was all from the noise image (see Figure 1). The noise was generated from a sinusoidal luminance distribution. In this way, the statistics of the luminance distributions were preserved between the orientation stimulus and the background and there were no texture elements that could distinguish the orientation patch from the noise field when the SN of the orientation stimulus was brought to 0.

Sensitivity Tests

The first test stage was conducted at least 1 day before the beginning of the training stage, and the second test was conducted at least one day after the end of the training stage (see Figure 2A for schematic). Participants gazed upon the small fixation point on the center of the screen and viewed a stream of noise images that changed at 2 Hz. Every 4000 ms the fixation point changed from red to green to indicate the presence of an orientation stimulus to which the participants had 2000 ms to respond with a key press to indicate which of the 112.5° or 22.5° orientations was presented. The two orientations and the eight SN levels were pseudorandomly interleaved with 56 repetitions per condition, yielding a 1 hr long session with 896 trials.

Training Sessions

Participants were asked to refrain from eating or drinking prior to each of the nine training sessions. In these sessions they were given no task other than to maintain fixation at a red dot in the center of the screen. The stimulus on the screen changed at a rate of 2 Hz, and, at random intervals, a sinusoidal grating (SN = 0.2) was presented for 500 ms (see Figure 1 for schematic). If the grating was of the trained orientation (randomly chosen to be either 112.5° or 22.5° for each participant), ~ 1 ml of water was delivered to the participants through a tube that was positioned in their mouths; no water was delivered for the control orientation. The water delivery overlapped the last 100 ms of the presentation of the trained orientation. This timing was chosen so that reward could lead to learning both because it was predicted by the trained orientation and because reward-related processing would be concurrent with stimulus processing of the trained orientation. The stimulus sequence was random with the exception that the minimal time between orientation presentations was 3000 ms. Each training session contained of ~ 350 presentations of each orientation (exact number varied because stimulus sequence was random in this experiment) and lasted 1 hr, and participants consumed between 300–350 ml of water.

Experiment 2

Apparatus

The apparatus was similar to that used in Experiment 1 with the exception that participants viewed stimuli from two monitors through a haploscope. Two 19"

CRT monitors with resolution of 1024 by 768 pixels and a refresh rate of 120 Hz were used. The viewing distance was 82 cm.

Stimuli

The orientation stimuli were similar to those of Experiment 1; however, some changes were made to optimize the stimuli for use with CFS. Sinusoidal gratings and noise images were presented at only 10% contrast (luminance 44 ± 7 cd/m²) and the stimulus field subtended a total of 8° diameter with the rest of the screen in a uniform gray; to prevent hard edges at the stimulus boundary, the outer 4°–8° annulus, consisting of pure noise, was blurred with a Gaussian profile ($\sigma = 0.5^\circ$). In tests the SN varied from trial to trial (0.03, 0.05, 0.07, 0.1, 0.13, 0.16, and 0.2). In the training sessions a constant SN of 0.2 was used.

Sensitivity Tests

These sessions were the same as in Experiment 1, with the following exceptions. A total of three sensitivity tests were conducted: the first test stage before the training stage, the second after 10 days of training, and the third after 20 days of training. For each 4000 ms trial the orientation and noise stimulus sequence was presented to one eye and a gray screen was presented to the other eye. The two orientations, the eight SN levels, and the two eyes-of-presentation were pseudorandomly interleaved with 28 repetitions per condition, yielding a 1 hr long session with 896 trials.

Training Sessions

The training session was similar to that for Experiment 1, but was adapted to promote CFS conditions that would ensure that the orientation stimuli were suppressed during the entire duration of the 20 day training period. To accomplish this, we presented alternating blocks of 15 s CFS stimuli to each eye. The CFS stimuli consisted of a sequence of full-screen textured pattern images that were presented at a rate of 10 Hz (see Figure 3A). The texture pattern consisted of 300 randomly placed, physically overlapping rectangles that varied in size with dimensions between 0.5° and 5°, had random angles of orientation, and consisted of saturated colors (0 or 100 cd/m²). In addition to the textured patterns, spatially sparse, colored, pixel noise covered a total area of 50% of the screen; we found that the addition of this noise was necessary to reliably suppress the high-frequency content of the training stimuli. For the trained eye a sequence of 2 Hz noise and orientation patterns were presented while CFS was shown to the other eye. For the untrained eye, a gray screen was presented while CFS was presented to the trained eye. After every 5 min, participants took a 3 min break. These sequences repeated eight times per day. In each day of training, 160 trials of each of the trained and untrained orientations were presented and participants drank ~150 ml of water.

Awareness Tests

The awareness test was conducted on a day following the final sensitivity test. In these sessions participants viewed the identical stimulus sequence as was shown during the training phase, and they were asked to report if they saw orientation stimuli similar to those that were presented during the sensitivity tests. After conducting this session, participants were asked to fill out a survey querying them on whether they noticed orientation stimuli or reward contingencies during the training sessions, and if so, which orientation was paired with reward.

SUPPLEMENTAL DATA

The supplemental data for this article include six supplemental figures and can be found at [http://www.neuron.org/supplemental/S0896-6273\(09\)00083-X](http://www.neuron.org/supplemental/S0896-6273(09)00083-X).

ACKNOWLEDGMENTS

We would like to thank John O'Doherty, Mitsuo Kawato, and Rufin Vogels for helpful discussions regarding this work; Ladan Shams, Rachel Miller, and Pablo Bartfield for comments on the manuscript; and Jonathan Dobres for his artistic contributions. This study was funded by grants from NIH (R01 EY015980 and R21 EY017737), NSF (BCS-0345746, BCS-0549036, and BCS-PR04-137 Center of Excellence for Learning in Education, Science, and Technology), and the Human Frontier Science Program Organization (RGP18/2004).

Accepted: January 23, 2009

Published: March 11, 2009

REFERENCES

- Ahissar, M., and Hochstein, S. (1993). Attentional control of early perceptual learning. *Proc. Natl. Acad. Sci. USA* 90, 5718–5722.
- Ahissar, M., and Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature* 387, 401–406.
- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spat. Vis.* 10, 433–436.
- Dalley, J.W., McGaughy, J., O'Connell, M.T., Cardinal, R.N., Levita, L., and Robbins, T.W. (2001). Distinct changes in cortical acetylcholine and noradrenaline efflux during contingent and noncontingent performance of a visual attentional task. *J. Neurosci.* 21, 4908–4914.
- Dayan, P., and Balleine, B.W. (2002). Reward, motivation, and reinforcement learning. *Neuron* 36, 285–298.
- Dorris, M.C., and Glimcher, P.W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44, 365–378.
- Dosher, B.A., and Lu, Z.L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proc. Natl. Acad. Sci. USA* 95, 13988–13993.
- Dupuis-Roy, N., and Gosselin, F. (2007). Perceptual learning without signal. *Vision Res.* 47, 349–356.
- Fahle, M. (2005). Perceptual learning: specificity versus generalization. *Curr. Opin. Neurobiol.* 15, 154–160.
- Fahle, M., Edelman, S., and Poggio, T. (1995). Fast perceptual learning in hyperacuity. *Vision Res.* 35, 3003–3013.
- Fan, J., McCandliss, B.D., Sommer, T., Raz, A., and Posner, M.I. (2002). Testing the efficiency and independence of attentional networks. *J. Cogn. Neurosci.* 14, 340–347.
- Gottfried, J.A., O'Doherty, J., and Dolan, R.J. (2003). Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* 301, 1104–1107.
- Herzog, M.H., and Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vision Res.* 37, 2133–2141.
- Herzog, M.H., and Fahle, M. (1999). Effects of biased feedback on learning and deciding in a vernier discrimination task. *Vision Res.* 39, 4232–4243.
- Karni, A., and Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc. Natl. Acad. Sci. USA* 88, 4966–4970.
- Lauwereyns, J., Takikawa, Y., Kawagoe, R., Kobayashi, S., Koizumi, M., Coe, B., Sakagami, M., and Hikosaka, O. (2002). Feature-based anticipation of cues that predict reward in monkey caudate nucleus. *Neuron* 33, 463–473.
- Leon, M.I., and Shadlen, M.N. (1999). Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron* 24, 415–425.
- Leopold, D.A., and Logothetis, N.K. (1996). Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 379, 549–553.
- Li, W., Piech, V., and Gilbert, C.D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nat. Neurosci.* 7, 651–657.
- Logothetis, N.K., Leopold, D.A., and Sheinberg, D.L. (1996). What is rivaling during binocular rivalry? *Nature* 380, 621–624.
- Lu, Z.L., Chu, W., Dosher, B.A., and Lee, S. (2005). Independent perceptual learning in monocular and binocular motion systems. *Proc. Natl. Acad. Sci. USA* 102, 5624–5629.
- Ludwig, I., and Skrandies, W. (2002). Human perceptual learning in the peripheral visual field: sensory thresholds and neurophysiological correlates. *Biol. Psychol.* 59, 187–206.
- Mitz, A.R. (2005). A liquid-delivery device that provides precise reward control for neurophysiological and behavioral experiments. *J. Neurosci. Methods* 148, 19–25.
- Morris, J.S., Ohman, A., and Dolan, R.J. (1998). Conscious and unconscious emotional learning in the human amygdala. *Nature* 393, 467–470.

- Nishida, S., Ashida, H., and Sato, T. (1994). Complete interocular transfer of motion aftereffect with flickering test. *Vision Res.* 34, 2707–2716.
- Nishina, S., Seitz, A., Kawato, M., and Watanabe, T. (2007). Effect of spatial distance to the task stimulus on task-irrelevant perceptual learning of static Gabors. *J. Vis.* 7, 1–10.
- O'Doherty, J.P., Deichmann, R., Critchley, H.D., and Dolan, R.J. (2002). Neural responses during anticipation of a primary taste reward. *Neuron* 33, 815–826.
- O'Doherty, J.P., Buchanan, T.W., Seymour, B., and Dolan, R.J. (2006). Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and ventral striatum. *Neuron* 49, 157–166.
- Paradiso, M.A., Shimojo, S., and Nakayama, K. (1989). Subjective contours, tilt aftereffects, and visual cortical organization. *Vision Res.* 29, 1205–1213.
- Pearce, J.M., and Bouton, M.E. (2001). Theories of associative learning in animals. *Annu. Rev. Psychol.* 52, 111–139.
- Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442.
- Pessiglione, M., Petrovic, P., Daunizeau, J., Palminteri, S., Dolan, R.J., and Frith, C.D. (2008). Subliminal instrumental conditioning demonstrated in the human brain. *Neuron* 59, 561–567.
- Petrov, A., Doshier, A., and Lu, Z. (2006). Perceptual learning without feedback in non-stationary contexts: data and model. *Vision Res.* 46, 3177–3197.
- Polley, D.B., Steinberg, E.E., and Merzenich, M.M. (2006). Perceptual learning directs auditory cortical map reorganization through top-down influences. *J. Neurosci.* 26, 4970–4982.
- Posner, M.I., and Petersen, S.E. (1990). The attention system of the human brain. *Annu. Rev. Neurosci.* 13, 25–42.
- Roelfsema, P.R., and van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Comput.* 17, 2176–2214.
- Schoups, A., Vogels, R., Qian, N., and Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature* 412, 549–553.
- Schultz, W. (2000). Multiple reward signals in the brain. *Nat. Rev. Neurosci.* 1, 199–207.
- Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annu. Rev. Psychol.* 57, 87–115.
- Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Seitz, A.R., and Watanabe, T. (2003). Psychophysics: Is subliminal learning really passive? *Nature* 422, 36.
- Seitz, A., and Watanabe, T. (2005). A unified model for perceptual learning. *Trends Cogn. Sci.* 9, 329–334.
- Seitz, A.R., and Dinse, H.R. (2007). A common framework for perceptual learning. *Curr. Opin. Neurobiol.* 17, 148–153.
- Seitz, A., Lefebvre, C., Watanabe, T., and Jolicoeur, P. (2005a). Requirement for high-level processing in subliminal learning. *Curr. Biol.* 15, R753–R755.
- Seitz, A.R., Nanez, J.E., Holloway, S.R., Koyama, S., and Watanabe, T. (2005b). Seeing what is not there shows the costs of perceptual learning. *Proc. Natl. Acad. Sci. USA* 102, 9080–9085.
- Seitz, A.R., Nanez, J.E., Holloway, S.R., and Watanabe, T. (2005c). Visual experience can substantially alter critical flicker fusion thresholds. *Hum. Psychopharmacol.* 20, 55–60.
- Seitz, A.R., Nanez, J.E., Holloway, S.R., and Watanabe, T. (2006). Perceptual learning of motion leads to faster flicker perception. *PLoS ONE* 1, e28.
- Shiu, L.P., and Pashler, H. (1992). Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Percept. Psychophys.* 52, 582–588.
- Suzuki, S., and Grabowecky, M. (2007). Long-term speeding in perceptual switches mediated by attention-dependent plasticity in cortical visual processing. *Neuron* 56, 741–753.
- Thiel, C.M., Friston, K.J., and Dolan, R.J. (2002). Cholinergic modulation of experience-dependent plasticity in human auditory cortex. *Neuron* 35, 567–574.
- Tsuchiya, N., and Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nat. Neurosci.* 8, 1096–1101.
- Uka, T., Tanaka, H., Yoshiyama, K., Kato, M., and Fujita, I. (2000). Disparity selectivity of neurons in monkey inferior temporal cortex. *J. Neurophysiol.* 84, 120–132.
- von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science* 224, 1260–1262.
- Watanabe, T., Nanez, J.E., and Sasaki, Y. (2001). Perceptual learning without perception. *Nature* 413, 844–848.
- Watanabe, T., Nanez, J.E., Sr., Koyama, S., Mukai, I., Liederman, J., and Sasaki, Y. (2002). Greater plasticity in lower-level than higher-level visual motion processing in a passive perceptual learning task. *Nat. Neurosci.* 5, 1003–1009.
- Yu, A.J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692.