

Dissociable Neural Effects of Long-term Stimulus–Reward Pairing in Macaque Visual Cortex

Edit Frankó¹, Aaron R. Seitz², and Rufin Vogels¹

Abstract

■ It has been proposed that perceptual learning may occur through a reinforcement process, in which consistently pairing stimuli with reward is sufficient for learning. We tested whether stimulus–reward pairing is sufficient to increase the sensorial representation of a stimulus by recording local field potentials (LFPs) in macaque extrastriate area V4 with chronically implanted electrodes. Two oriented gratings were repeatedly presented; one was paired with a fluid reward, whereas no reward was given at any other time. During the course of conditioning the LFP increased for the rewarded compared to the unrewarded orientation. The time course of the effect of

stimulus–reward pairing and its reversal differed between an early and late interval of the LFP response: a fast change in the later part of the neural response that was dissociated from a slower change in the early part of the response. The fast change of the late interval LFP suggests that this late LFP change is related to enhanced attention during the presentation of the rewarded stimulus. The slower time course of the early interval response suggests an effect of sensorial learning. Thus, simple stimulus–reward pairing is sufficient to strengthen stimulus representations in visual cortex and does this by means of two dissociable mechanisms. ■

INTRODUCTION

Many studies in the last century have shown that experience influences the perception of a visual stimulus: Training in the detection or discrimination of a stimulus improves the ability to detect or discriminate the stimuli (Goldstone, 1998). Intriguingly, a recent human psychophysical study has shown that such perceptual learning occurs even for subliminally presented stimuli, that is, stimuli that the subject is unaware of (Watanabe, Nanez, & Sasaki, 2001), and which are task-irrelevant. Subsequent research showed that this subliminal learning only occurs when the stimuli are paired with successful execution on a concurrent task or a reward (Seitz & Watanabe, 2003). Based on these psychophysical data, it has been proposed that simple stimulus–reward pairing is sufficient to obtain learning. This reinforcement-based theory of perceptual learning proposes that stimulus–reward pairing enhances neural processing of the stimulus paired with the reward but not of stimuli that were not paired with the reward (Seitz & Watanabe, 2005). Such a mechanism would prevent subjects from automatically learning all stimuli that they are exposed to, which is highly inefficient, but, instead, increases neural processing only of significant stimuli as tagged by a concurrent reward. It should be noted that such a mere stimulus–reward pairing mechanism might be supplemented by stimulus-selective attentional

or task-related mechanisms (Yotsumoto & Watanabe, 2008; Seitz & Dinse, 2007; Blake, Heiser, Caywood, & Merzenich, 2006; Polley, Steinberg, & Merzenich, 2006; Roelfsema & van Ooyen, 2005; Ahissar & Hochstein, 1993).

A critical assumption of the reinforcement-based theory of perceptual learning is that mere stimulus–reward pairing is sufficient to boost selectively the response of the reward paired stimulus in the brain regions that change during learning. Previous electrophysiological studies in nonhuman primates have shown that extensive training in visual discrimination tasks enhances the selectivity of primary visual (V1; Li, Piech, & Gilbert, 2004; Crist, Li, & Gilbert, 2001; Schoups, Vogels, Qian, & Orban, 2001; Ito, Westheimer, & Gilbert, 1998) and extrastriate V4 neurons (Raiguel, Vogels, Mysore, & Orban, 2006; Rainer, Lee, & Logothetis, 2004). Some results also indicate that V4 shows stronger learning effects than V1 neurons (Raiguel et al., 2006; Yang & Maunsell, 2004; Ghose, Yang, & Maunsell, 2002).

Given the work suggesting that V4 is involved in perceptual learning, we tested the assumption that mere pairing of a stimulus with a reward is sufficient to selectively increase the neural response in V4 to that stimulus. V4 neurons are known to be sensitive to simple local form features, such as orientation (Desimone, Schein, Moran, & Ungerleider, 1985) or curvature (Pasupathy & Connor, 1999, 2001), and to wavelength (Schein & Desimone, 1990; Zeki, 1978), which allows a manipulation of simple, controllable stimulus parameters. Thus, we chose to stimulate V4 with two oriented gratings, which differed 90° in

¹K.U. Leuven Medical School, Leuven, Belgium, ²University of California, Riverside

orientation, and we consistently presented one of these two grating orientations in conjunction with a liquid reward. To prevent response saturation and to increase the possibility of observing a reward effect, we reduced the visibility of the gratings by replacing 80% of the gratings' pixels with noise and embedded them in a background noise pattern. The animals were required only to fixate while the sequence of stimuli was shown. Reward was delivered only when two conditions were met: fixation of the fixation target and presentation of one particular orientation. This protocol is essentially the same as used in discriminative classical (Pavlovian) conditioning.

Both animals were chronically implanted with a multi-electrode array in area V4, allowing measurements of neural activity during the conditioning on successive days, and thus, to document the time courses of the reward induced changes in the neural response.

We recorded local field potentials (LFPs) from each of the 96 electrodes during stimulus presentation. The LFP measures synaptic activity including other types of slow activity such as spike afterpotentials and voltage-dependent membrane oscillations and reflects the input and intracortical processing of a large population of neurons around the tip of the electrode (Logothetis, 2003). In this study, we tested whether the LFP of the rewarded stimulus relative to that of the unrewarded one would be increased by consistent stimulus–reward pairing.

METHODS

Subjects

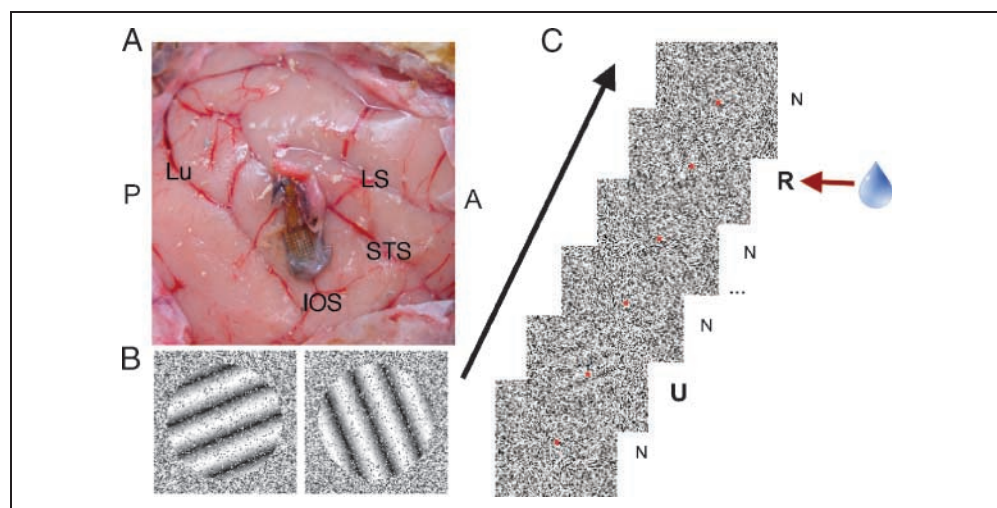
Two adult rhesus monkeys (M1: male, 3 kg; M2: female, 4.3 kg) with normal (M2) or corrected-to-normal vision (M1 both eyes correction: $-3D$) were extensively trained to fixate a central fixation point for long durations. Subjects

were equipped with polymer headposts to fix the head to the primate chair during experiments. A 96-electrode array (see below for details) was implanted into prelunate gyrus (dorsal area V4) in the right hemisphere of M1 and in the left hemisphere of M2. The position of the array during the surgery was determined using sulcal landmarks. All surgical procedures were performed under general anesthesia (isoflurane/nitrous oxide) and sterile conditions. In Monkey M1 we have postmortem verification of the array being in area V4 (Figure 1A). In M2, because it is still alive, we visualized the array position with MRI (3-T Siemens; MP-RAGE sequence; custom-designed coil; 0.6-mm slices): Artifacts of the array were observed in prelunate gyrus dorsal to inferior occipital sulcus and posterior to superior temporal sulcus, which indirectly confirm accurate positioning of the array. Animal care and experimental procedures were approved by the ethical committee of the Katholieke Universiteit Leuven Medical School.

Recordings and Apparatus

We used the Utah array (Cyberkinetics, Neurotechnology Systems, Foxboro, MA) for the recordings. This 4×4 mm array has ninety-six 1-mm-long electrodes, metalized with platinum, in a square configuration with $400 \mu\text{m}$ interelectrode distance. Impedances, measured with a 100-nA, 1-kHz current in vivo during the course of the study, ranged 1.0–2.5 M Ω and 0.1–0.7 M Ω for M1 and M2, respectively. Reference was one of two subdural wires with stripped ends placed several centimeters frontal to the recording region. The same reference wire was used throughout the training but we confirmed that the two reference wires, which were placed a few centimeters far from each other, produced similar LFPs. We observed occasionally spikes from single cells or multiunits in M1 but these were unresponsive to

Figure 1. Electrode array location, stimuli, and task. (A) Position of the array. A 96-electrode array (Utah array) was implanted into the dorsal part of area V4 as shown in this autopsy photograph of the brain of M1. LS = lateral sulcus; STS = superior temporal sulcus; Lu = lunate sulcus; IOS = inferior occipital sulcus; A = anterior; P = posterior. (B) The stimuli. The stimuli were sinusoidal gratings spatially masked by 80% sinusoidal noise (for clarity, the gratings in the figure are displayed with only 20% noise). The orientations were 22.5° and 112.5° for M1, and 67.5° and 157.5° for M2. (C) Paradigm. The monkeys fixated while full-screen noise (N) was changed every 500 msec, and at pseudorandom intervals one of the two oriented gratings (U = unrewarded; R = rewarded) was presented. The delivery of the fluid reward overlapped the last 100 msec of the rewarded orientation.



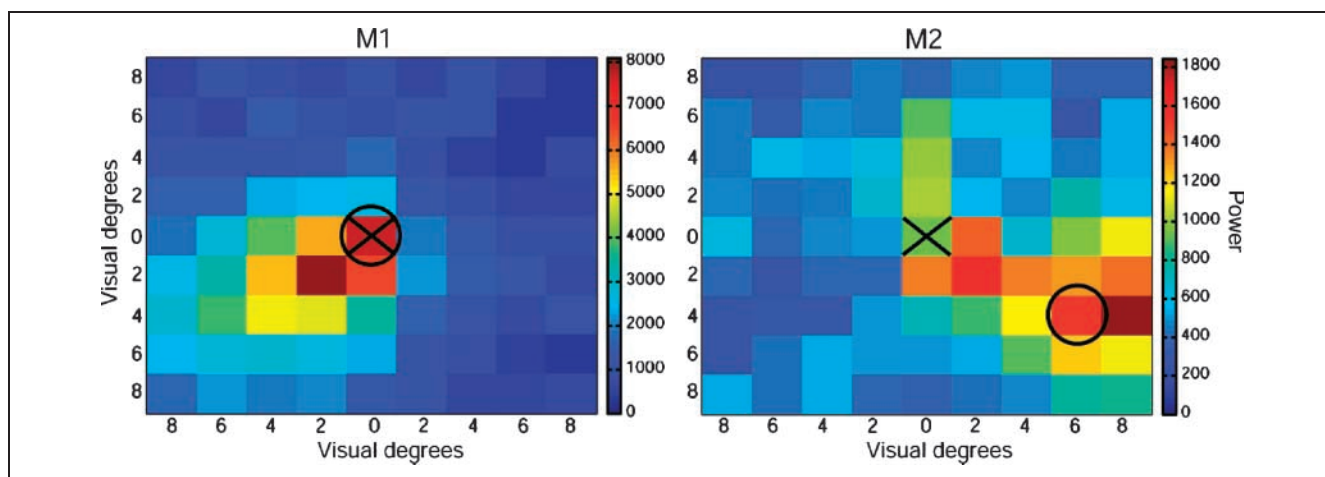


Figure 2. Receptive field mapping. Averaged power across channels for the 81 stimulus positions ($18^\circ \times 18^\circ$) for monkey M1 and M2. The cross indicates the foveal position, the circle indicates the stimulus position used during the training.

the oriented gratings. During the training, no spikes could be recorded in M2. To obtain stable recordings from a settled array, we conducted the present study 6 and 7 months after array implantation in M1 and M2, respectively.

We employed the Cerebus system (Cyberkinetics, Neurotechnology Systems) to record neural signals from the array as well as stimulus and behavioral events. The input impedance of the Cerebus preamplifier is $>10^{12} \Omega$. LFPs were filtered on-line between 0.3 and 250 Hz, and were sampled at 1 kHz. Subjects' eye movements were monitored using a noninvasive infrared eye-tracking device (Iscan, Burlington, MA). A PC running custom software controlled stimulus presentation, ran the task, and sampled eye position signals using DSP boards.

Stimuli

The stimuli were small, circular patches 2° and 4° in diameter (respectively for M1 and M2) containing a sinusoidal grating of 2 cycles/degree, with orientation at 22.5° and 112.5° for M1 and 67.5° and 157.5° for M2 (Figure 1B). This patch was displayed against a noise background that had the same sinusoidal luminance distribution as the grating. Signal strength was manipulated by replacing a proportion of the pixels of the grating by sinusoidal noise. During stimulus–reward pairing conditioning, this proportion was fixed to 80% (20% signal). Grating phase and the masking noise were randomized across presentations. The stimuli were presented gamma-corrected on a 20-inch display monitor with 100% contrast.

The position of the grating was based on a retinotopic mapping of the LFP responses in each animal, conducted before the conditioning study. To determine the stimulus position that evoked the strongest LFP, a black star (size of $2^\circ \times 2^\circ$) was presented at 81 positions (steps: 2°) in an $18^\circ \times 18^\circ$ square grid centered on the fixation point. The stimulus was presented for 300 msec during fixation of a

small target. Each stimulus was presented about 50 and 40 times in monkeys M1 and M2, respectively. To quantify the LFP response to the different stimulus positions, we computed the baseline-corrected squared amplitude of the signal obtained during stimulus presentation. We subtracted from each LFP of a stimulus position the mean LFP value computed from 0 till 300 msec before stimulus onset at that position. After this baseline correction, we computed the squared amplitude of the LFP during stimulus presentation. We chose that position for the conditioning experiment which generated a strong LFP signal on the majority of the electrodes. The LFP retinotopic maps (Figure 2) differed between animals, being more eccentric in M2 than in M1, which is very likely due to a different array location in V4. Accordingly, the selected stimulus position for the conditioning experiment differed between the monkeys: foveal position in M1 and in the lower visual field, 4° below the horizontal and 6° lateral to the vertical meridian in M2.

Task and Conditioning Protocol

The subjects initiated a trial by fixating a red spot (size: 0.19°) in the center of the screen. Every 500 msec, a new noise background filled the screen (Figure 1C). One of the two oriented gratings appeared for 500 msec after one to six noise background presentations (Figure 1B and C). One orientation was consistently paired with a fluid reward (rewarded stimulus), which was delivered 400 msec after stimulus onset. The other orientation was never followed by a reward (unrewarded stimulus). Reward was delivered only when two conditions met: fixation of the fixation target and presentation of one particular orientation. Thus, the juice delivery was always preceded by, and partially overlapped, the presentation of one particular orientation (the rewarded stimulus), whereas the presentation of the other orientation was never followed by a

reward (the unrewarded stimulus). The stimulus presentation continued as long as the monkey maintained gaze within a small fixation window (full size: 1.25–1.5 visual degrees), on average, for 11 and 7.5 sec in M1 and M2, respectively, and until the sum of the presentations of the rewarded and unrewarded gratings equated 6. The presentation of the gratings followed a pseudorandomized schedule in which the order of rewarded and unrewarded grating presentations was random with the constraint that the maximal difference in proportion of rewarded versus unrewarded gratings could be only 33% in a completed trial. The actual number of grating presentations during the experiment was similar for the rewarded and unrewarded stimuli [average number of presentations per day: M1, rewarded = 440 ($SD = 201$), unrewarded = 419 ($SD = 186$); M2, rewarded = 943 ($SD = 190$), unrewarded = 818 ($SD = 192$)]. In M1 the rewarded and the unrewarded stimulus occurred, on average, at the same time within a trial (M1: median rewarded = 12th position, $Q_1 = 7$, $Q_3 = 19$; median unrewarded = 12th position, $Q_1 = 6$, $Q_3 = 18$), whereas in M2 there was a difference of one position (500 msec) between the average position of the two stimuli (M2: median rewarded = 9th position, $Q_1 = 5$, $Q_3 = 15$; median unrewarded = 8th position, $Q_1 = 5$, $Q_3 = 14$). Given the range of position differences, these small differences between the within-trial position of the rewarded and unrewarded stimuli are negligible. The animals showed a slightly greater tendency to break fixation to the rewarded versus the unrewarded stimulus in the first phase of the training. This tendency was greater in M2 than in M1 and decreased consistently, for both animals, throughout the experiment. However, after reversal of the stimulus–reward pairing (see below), the sign of the effect differed between animals, with M1 showing more aborts for the unrewarded orientation and M2 showing a slight tendency for more aborts to the rewarded orientation (data not shown).

The rewarded orientation was chosen from the *sensitivity test* conducted before the start of the conditioning. In this test, we measured the LFPs for four orientations (22.5°, 67.5°, 112.5°, 157.5°) each at four signal-to-noise ratio levels (10%, 20%, 40%, 80%). The gratings were shown interleaved with presentations of full-screen scenes and full-screen noise. Stimuli were presented for 300 msec, preceded by a 300–1000 msec noise background, and 400 msec after stimulus offset fluid reward was delivered in 50% of the presentations. Only oblique orientations were used as these might be more susceptible to training effects than the principal orientations (Vogels & Orban, 1985). The orientation which evoked the smallest LFP response was chosen as the rewarded stimulus, and the orthogonal one as the unrewarded with the constraint that different orientations were chosen in M1 and M2. Thus, we started the conditioning with 112.5° and 67.5° being the rewarded orientations in M1 and M2, respectively. The difference between the rewarded and unrewarded stimulus was on the first day of the training, $-7.5 \mu\text{V}$ (the valley-

peak amplitude for the rewarded grating: $12.9 \mu\text{V}$; unrewarded grating: $20.4 \mu\text{V}$) and $-1.8 \mu\text{V}$ (rewarded grating: $37.6 \mu\text{V}$; unrewarded grating: $39.4 \mu\text{V}$) in M1 and M2, respectively. The stimulus–reward pairing was reversed after 55 and 37 days of conditioning for M1 and M2, respectively. During the reversal, the previously rewarded stimulus became unrewarded and the previously unrewarded one became rewarded. We continued the conditioning with this reversed pairing for 54 and 20 days for M1 and M2, respectively. The whole training, including sensitivity tests, lasted, with interruptions, 9 months for M1 and 3 months for M2.

Analyses of LFPs

For LFP signal analysis, we employed the ELAN software (Mental Processes and Brain Activation Lab, Inserm Unit 280, Lyon, France) and custom scripts written in MATLAB. Six of the 96 and 16/96 of the electrodes in M1 and M2, respectively, were excluded because of their high noise level or high impedance ($>14 \text{ M}\Omega$). Stimulus presentations with LFP amplitudes exceeding a $400\text{-}\mu\text{V}$ (M1) or $700\text{-}\mu\text{V}$ (M2) threshold, or followed by a saccade outside the fixation window (abort) within 500 msec poststimulus onset, were rejected from the analyses. For the analyses, we included only LFP responses to grating presentations that were preceded by at least three background noise patterns (1500 msec). We used this criterion because in the LFP stimulus-related fluctuation can be seen even few seconds after the stimulus presentation. To have similar baselines to the stimuli without strong effect of the previous grating, we excluded stimulus presentations that were preceded by a grating within 1500 msec. We performed baseline correction on the LFP responses using the 300-msec interval before stimulus onset to be able to compare the amplitude of the responses.

To compare the effect of the conditioning for the rewarded and unrewarded stimuli, we averaged the baseline-corrected LFPs across the channels, and computed the signed difference in LFP waveform each training day for both animals. For statistical comparisons of the LFP waveforms between the rewarded and unrewarded stimuli, we employed the Kruskal–Wallis test using a sliding window of 10 msec width in steps of 5 msec for each channel (Type I error level: $p < .0001$; this small p value was used to correct for multiple comparisons).

Time–frequency power analysis of the LFP signal of each channel was performed using convolution of the signal for each selected stimulus presentation with Morlet wavelets (center frequency: $f_0/\sigma_f = 7$; Tallon-Baudry & Bertrand, 1999) in a 1500-msec-long window ranging from -500 to 1000 msec relative to stimulus onset. The time–frequency power analysis was performed between 2 and 150 Hz in steps of 1. The trial-by-trial power was averaged across the different presentations of the same stimulus, and was normalized for each frequency dividing the values by the mean power computed in the baseline period 300 msec before grating onset. Then the normalized power was

averaged across electrodes within a time window in which we saw the strongest change in the LFP amplitude. To quantify the changes in power, we computed the mean normalized power for five “conventional” frequency bands ranging between 7 and 150 Hz: alpha (7–14 Hz), beta (15–29 Hz), low gamma (30–59 Hz), middle gamma (60–99 Hz), and high gamma (100–150 Hz). In each band, we computed a power index, defined as (power firstly rewarded – power firstly unrewarded)/(power firstly rewarded + power firstly unrewarded). We want to stress that our analyses focus on the difference between the responses to the rewarded and unrewarded stimuli. If the baseline power changed between days, it should have a similar effect on the normalized power for both stimuli.

Eye Movement Analysis

We performed several analyses of the eye movements recorded for those stimulus presentations that entered the LFP analysis (see above). In the first analysis, we computed the standard deviation of the horizontal and vertical eye positions during the rewarded and unrewarded stimulus presentations. Standard deviations of eye position were computed for different poststimulus temporal intervals (see Results), for each stimulus presentation, and then averaged across the presentations of a stimulus for each daily session.

The standard deviation of eye position captures all sorts of eye movements, such as tremor, drifts, and microsaccades. To isolate saccades from drifts and tremors, we performed an analysis of microsaccades. For this we employed the automatic microsaccade detection algorithm of Engbert and Kliegl (2003). An eye movement inside the fixation window of which the speed exceeded 3 standard deviations of the eye speed computed for each stimulus presentation was classified as a microsaccade. The number of microsaccades and their amplitude were averaged for different temporal intervals across stimulus presentations for each day.

RESULTS

During the course of the training, the LFP, averaged across electrodes, increased in both animals for both the rewarded and unrewarded orientations (Figure 3A, B, and C) before the stimulus–reward reversal (vertical lines in panels of Figure 3). Importantly, during this phase of the conditioning, the difference between the LFP amplitudes for the two grating orientations increased (Figure 3D): Initially, the responses to the two orientations were similar, while later on the LFP amplitude evoked by the rewarded stimulus became larger than that by the unrewarded one.

Closer examination of the LFP waveforms for the rewarded and unrewarded orientation revealed distinct temporal intervals of activation induced by the stimuli. Following a first orientation unselective peak at around

50 msec (negative in M1; positive in M2; see Figure 3B and C), both monkeys showed a selective training-induced increase in the amplitude of the LFP. Based on visual inspection of Figure 3D, we defined in both monkeys the interval in which this first orientation-selective LFP peak occurred. These intervals (Int1) ranged from 110 to 200 msec and from 85 to 125 msec poststimulus onset in M1 and M2, respectively, as indicated in Figure 3C (I). For each day, we determined within this interval the maximum difference in mean LFP using 5-msec-long windows. For each day, the (signed) LFP differences were averaged within a range of 45 msec centered on the maximum of that day. These averaged LFP amplitude differences are plotted in Figure 4. This quantification of the LFP amplitude in the early selective response interval showed that, during the course of training, the difference in LFP amplitude between the rewarded and unrewarded orientation increased significantly (correlation between LFP difference and training days before stimulus–reward reversal: M1, $r = .57, p < .05, n = 55$; M2, $r = .83, p < .05, n = 37$; Figure 4A and B).

Following this early response interval, a second, biphasic response interval was evident. The later phase (Int2) of this response interval, roughly from 250 msec to reward onset, was more prominent and was the focus of our analysis of this second interval (however, both phases showed qualitatively similar effects). We defined in each monkey an interval in which this later orientation-selective activity occurred: Int2 ranged from 270 to 330 msec and from 240 to 350 msec in M1 and M2, respectively. For each day, we determined within Int2 the absolute maximum difference in mean LFP using 5-msec-long windows, and then computed the daily LFP differences within a range of 45 msec centered on the maximum of that day (same procedure as for Int1; see above). As shown in Figure 4C and D, during the course of training, in Int2 the mean difference in amplitude between the rewarded and unrewarded orientations changed preceding the stimulus–reward reversal. This effect of conditioning in Int2 was significant in both animals (M1: $r = -.65, p < .05, n = 55$; M2: $r = .82, p < .05, n = 37$), but differed in sign between the animals. However, in Int2 both animals showed an increase in absolute LFP difference between the rewarded and unrewarded orientations.

These orientation-selective changes in Int1 and Int2 show that stimulus–reward pairing increased the amplitude of the LFP signal for the rewarded compared to the unrewarded stimulus in a visual area, and these changes occurred well before reward onset. To better understand this pairing effect, we reversed the orientation–reward pairing after 55 and 37 days of conditioning in M1 and M2, respectively. We found that the effect of reversing the stimulus–reward pairings differed greatly between the early and the late response intervals. In Int1, the difference between the rewarded and the unrewarded orientation decreased slowly after reversal (M1: $r = -.82, p < .05, n = 54$; M2: $r = -.77, p < .05, n = 20$; Figure 4A and B).

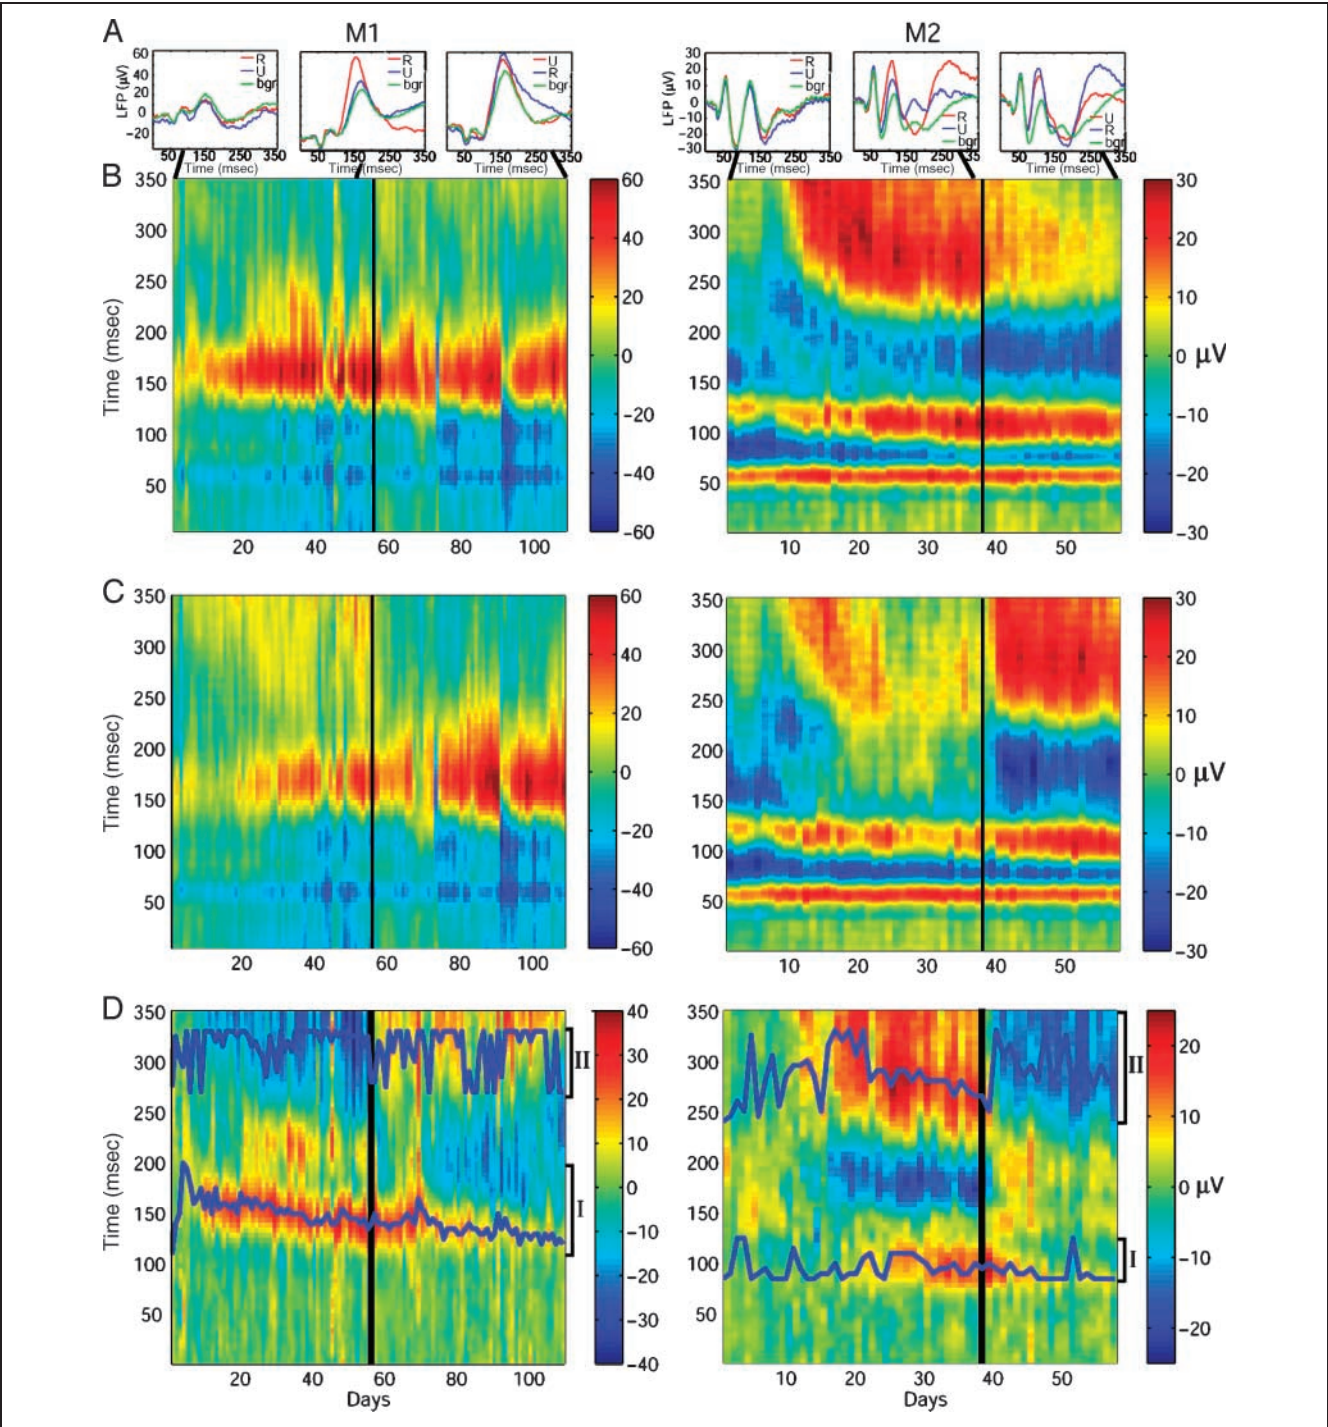



Figure 3. LFP waveforms to the rewarded and unrewarded orientations. (A) LFP response to the gratings and to the background noise stimulus on the first day of the training, on the 50th (M1) and on the 36th (M2) day, and on the last day of the training. Note that on the last day of training of M1 (after the reversal), the second interval amplitude change occurred relatively late. (B and C) Color-coded LFP amplitude from stimulus onset until 350 msec after stimulus onset as a function of the training days for the firstly rewarded grating (B, 112.5° for M1, 67.5° for M2), and for the firstly unrewarded one (C, 22.5° for M1, 157.5° for M2). (D) LFP differences. Color-coded difference in LFP amplitude between the two gratings (112.5°–22.5° for M1 and 67.5°–157.5° for M2) from stimulus onset until 350 msec as a function of the training days. The vertical black line shows the time of the reversal. The blue lines are the maximum LFP differences in the first interval [Int1 (I): 110–200 msec (M1), 85–125 msec (M2)], and the absolute maximum LFP difference in the second interval [Int2 (II): 270–330 msec (M1), 240–350 msec (M2)].

Figure 4. LFP differences. (A to D) The curves represent the difference in LFP amplitude between the two gratings (112.5° – 22.5° for M1 and 67.5° – 157.5° for M2) averaged within a 45-msec window centered on the maximum in the first interval (A and B), and on the absolute maximum in the second interval (C and D; Figure 3 blue lines). The vertical black lines indicate the time of the reversal.



Underlying this diminution of the response differences was a gradual increase in the LFP amplitude for the newly rewarded orientation; however, the response to the firstly rewarded orientation decreased a little (Figure 3B and C). This is in strong contrast to the effect of reversal in Int2: After changing the stimulus–reward pairing, the LFP difference between the two orientations reversed completely to favor the newly rewarded orientation within 3 days (Figure 4C and D; M1: $r = -.10$, ns , $n = 54$; M2: $r = -.59$, $p < .05$, $n = 20$; note that the nonsignificant effect of stimulus–reward reversal in M1 when restricting the analysis to the data obtained after the reversal is due to the fast rate of the reversal effect in this animal). Reversal of stimulus–reward contingencies demonstrated that the stimulus–reward pairing had distinct effects with different time courses: a relatively fast change in the later part of the neural response that was dissociated from a slower change in the early part of the response.

For M2, the total number of stimulus presentations during the entire training was 51,331 for the firstly rewarded grating (67.5°) and 48,650 for the firstly unrewarded grating (157.5°). If this small, 5%, difference in stimulus presentation number would be responsible for the changes in LFP with training, the LFP response to the grating with 67.5° orientation should be higher at the end of the training, but this was not the case in this animal. The total number of stimulus presentations in M1 was more similar (<1%): for the firstly rewarded grating (112.5°), 47,189 presentations, and for firstly unrewarded grating (22.5°), 47,542 presentations. Thus, the reward contingency, and not mere stimulus exposure, determined the orientation-dependent changes in LFP.

In the analysis above, we averaged the LFPs across electrodes. This absence of electrode selection provided an

unbiased assessment of the effect of stimulus–reward pairing on the LFP. We also analyzed the effect of stimulus–reward pairing on individual electrodes of both monkeys. Figure 5 shows the mean LFP difference between the two gratings (firstly rewarded – firstly unrewarded) averaged in a window of 45 msec for both the first and second intervals separately for each electrode as a function of the conditioning days. It is clear that the effects observed when averaging across electrodes were also seen on individual electrodes. Note that in both animals the sign of the LFP differences was consistent across electrodes. The size of the effect differed a bit among electrodes in M1, but much less so in M2. This is probably related to the smaller stimulus and foveal position of the stimulus in M1. The average pairwise correlation among the LFPs of the different electrodes was indeed smaller in M1 than in M2: for M1 on the first day of training (mean \pm standard deviation), $r = .8 \pm .14$, and on the last day of training, $r = .94 \pm .16$; for M2 on the first day of training, $r = .98 \pm .27$, and on the last day of training, $r = .98 \pm .21$.

Although the animals were restricted to keep fixation within a small window (size: <1.5 visual degrees), we examined whether eye movements differed between rewarded and unrewarded presentations. In Int1 the mean standard deviations of eye position, averaged across stimulus presentations, differed little between rewarded and unrewarded orientations ($<0.01^\circ$; Figure 6A and B). For this interval, the difference in mean standard deviation of the eye position (firstly rewarded – firstly unrewarded) did not correlate with the difference in LFP amplitude (x : $r = -.08$, ns and y : $r = -.01$, ns for M1; x : $r = -.15$, ns and y : $r = -.25$, ns for M2). We also measured the frequency of microsaccades during the stimulus presentation. It was similar for both the rewarded and unrewarded

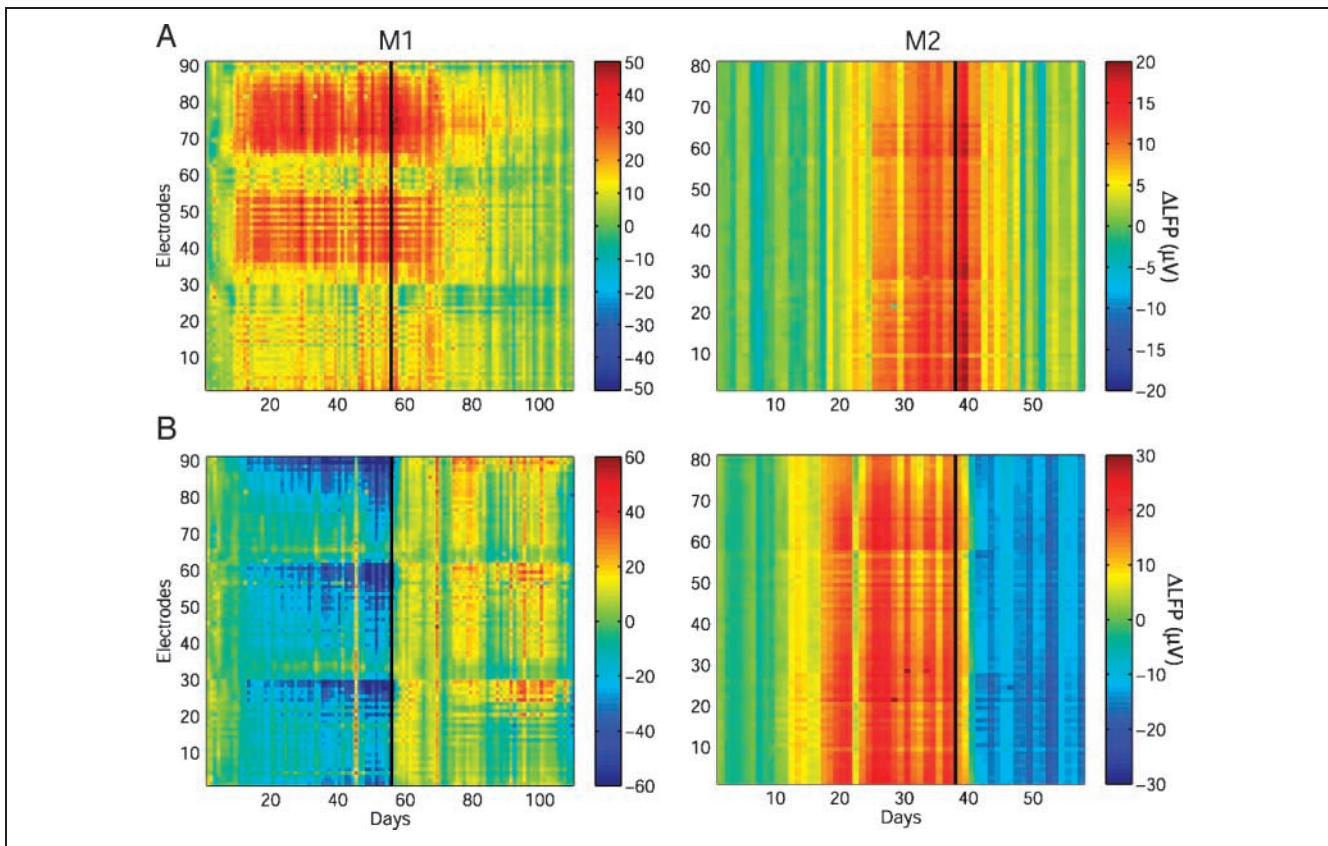
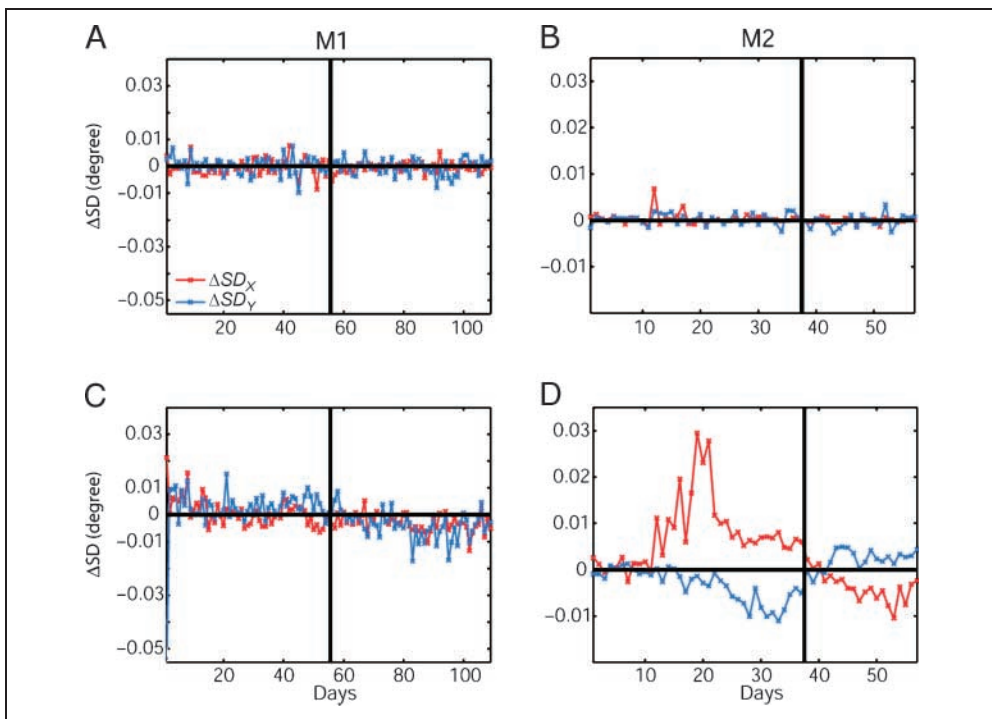


Figure 5. The effect of the reward for individual electrodes of monkey M1 and M2. Color-coded mean LFP difference between the two gratings (firstly rewarded – firstly unrewarded) averaged in a window of 45 msec in the first interval (A) and in the second interval (B), separately for each electrode as a function of the training days. The vertical black line indicates the time of the reversal.

Figure 6. Eye positions. (A to D) The red and the blue curves represent the differences in the average standard deviations of the horizontal (red) and the vertical (blue) eye positions between the two gratings in the first and second intervals (e.g., ΔSD_x grating 1 – ΔSD_x grating 2). The vertical black lines indicate the time of the reversal.



stimuli (M1: rewarded = 1.58 Hz, unrewarded = 1.53 Hz; M2: rewarded = 1.92 Hz, unrewarded = 1.92 Hz). Neither the number nor amplitude of microsaccades between rewarded and unrewarded conditions correlated significantly ($p < .05$) with the LFP difference for this early interval [M1: number of microsaccades ($r = -.12$, ns), amplitude ($r = .09$, ns); M2: number of microsaccades ($r = .18$, ns), amplitude ($r = .04$, ns)]. Hence, for both monkeys, the learning effect observed in the LFP in Int1 cannot be explained by eye movements, which is not surprising given the short latency of Int1. However, for Int2 (Figure 6C and D), small eye position differences between the rewarded and unrewarded stimuli correlated with the LFP differences [correlation of difference in eye position standard deviation and LFP difference (Int2): M1 ($r = -.11$, ns for x ; $r = -.22$, $p < .05$ for y); M2 ($r = .73$, $p < .05$ for x ; $r = -.80$, $p < .05$ for y)]. In M2, the difference in LFP in Int2 between the rewarded and unrewarded stimuli correlated significantly with the difference in number ($r = .7$, $p < .05$) and amplitude ($r = -.61$, $p < .05$) of microsaccades across days in this interval. No such correlation between microsaccade metrics and LFP differences were present in M1 (number: $r = -.11$, ns ; amplitude: $r = -.16$, ns), who made less eye movements than M2 in this late interval (probably due to the foveal position of the grating for M1). Close inspection of the small changes in eye movements and the LFP after the reversal showed that the eye movements did not cause the differences observed in Int2 LFP responses. This can be clearly seen in the first two sessions after reversal in M2, where the LFP effect remained similar to that just before reversal, whereas there

was no corresponding difference in eye movements between the two orientations.

We also examined differences in the power of the LFP response to the two orientations in different frequency bands (time–frequency analysis) as a function of training days (Figure 7). For the time–frequency analysis, we used the same intervals as for the LFP amplitude analyses, except for a 15-msec longer time window for Int1 in M2 to have a better estimation of the power in the low frequencies [(Int1) M1: 110–200 msec, M2: 85–140 msec; (Int2) M1: 270–330 msec, M2: 240–350 msec]. Before the reversal, the difference in power between the rewarded and unrewarded orientations in Int1 increased both in the low (7–29 Hz) and in the high frequencies (50–150 Hz) with training.

We computed the mean power for five frequency bands: alpha (7–14 Hz), beta (15–29 Hz), low gamma (30–59 Hz), middle gamma (60–99 Hz), and high gamma (100–150 Hz). Figure 7 shows how the power index (see Methods) changes with the training (averaged within the bands, shown only for those bands for which the correlation of power index and days was significant; $p < .05$). For these selected significant bands for Int1, the correlations of the power indices and training days were positive before reversal and negative after reversal (correlation coefficients are listed in Table 1). The power changes in Int2 differed strikingly from those in Int1: Stimulus–reward pairing increased the power for the high frequencies and decreased the power for the low frequencies, and this effect switched after reversing the orientation–reward pairing (Figure 7B; Table 1). Similar trends were

Figure 7. Differences in power indices. (A and B) Power indices averaged within frequency bands in the first (A) and in the second (B) intervals for monkey M1 and M2. The vertical black lines indicate the time of the reversal. The mean changes in power index were averaged within the following frequency bands: alpha (7–14 Hz), beta (15–29 Hz), low gamma (30–59 Hz), middle gamma (60–99 Hz), high gamma (100–150 Hz). Only those bands for which the power index had a significant correlation with training days (Table 1) are plotted.

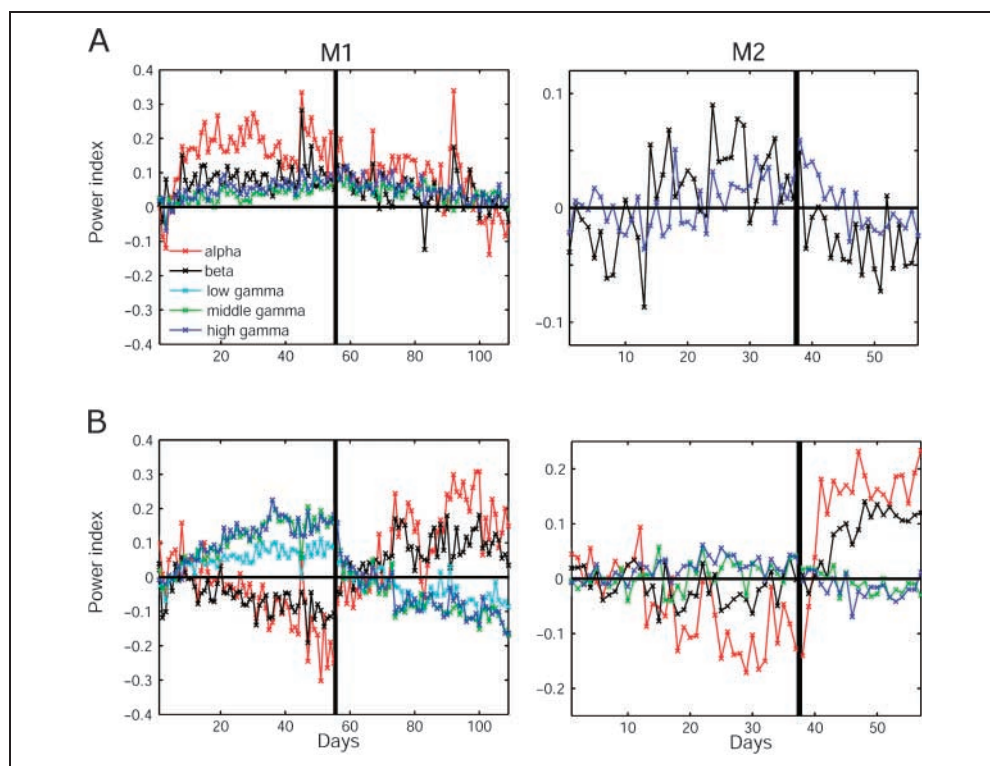


Table 1. Pearson Correlation Values (r) and p Values of the Correlations between Days of the Learning Phase before (Learning) and after the Reversal (Reversal) for Intervals 1 and 2, Respectively, and the Power Index Values for the Different Frequency Bands for Monkeys M1 and M2

	M1		M2	
	r	p	r	p
<i>Interval 1</i>				
Learning				
alpha	.387	.003	.038	.82
beta	.359	.007	.577	0
low gamma	.153	.262	-.003	.985
middle gamma	.727	0	.011	.947
high gamma	.703	0	.449	.005
Reversal				
alpha	-.565	0	.181	.444
beta	-.434	.001	-.432	.057
low gamma	-.079	.566	-.38	.097
middle gamma	-.648	0	-.094	.69
high gamma	-.747	0	-.794	0
<i>Interval 2</i>				
Learning				
alpha	-.808	0	-.71	0
beta	-.629	0	-.213	.203
low gamma	.725	0	.178	.291
middle gamma	.853	0	.376	.021
high gamma	.851	0	.578	0
Reversal				
alpha	.704	0	.636	.002
beta	.512	0	.815	0
low gamma	-.676	0	-.441	.051
middle gamma	-.824	0	-.677	.001
high gamma	-.832	0	-.307	.187

Significant values ($p < .01$) are indicated in bold.

observed when computing the “induced” power by subtracting the daily average LFP for that stimulus from the LFP of each presentation before computing the power (data not shown). Overall, the time–frequency power analysis substantiated the dissociation between the early and the late response phases seen in the LFP amplitudes. In the early interval, the power increased in high- and low-frequency bands during the first part of the training and changed slowly after reversal, whereas in the late interval

the high- and the low-frequency bands changed rapidly, and in opposite ways after reversal of the stimulus–reward pairing.

To assess whether the mere act of licking the juice or the reward itself without visual stimulation can cause an LFP signal on the visual cortical electrodes, we switched off the monitor in front of M2 and ran the experiment. The monkey did not have to fixate, and could not see the stimuli, but still obtained reward randomly. We compared the LFP between rewarded and unrewarded periods. We also measured the pressure inside the tube of the juice delivery system and confirmed that the monkey was still sucking the tube during reward delivery as in the real training. We found that without visual stimulation there was no difference in the LFP response between the rewarded and unrewarded periods, except for a small solenoid artifact from opening the switch. This confirms the visual nature of the LFP difference between rewarded and unrewarded stimuli during the conditioning.

To determine whether these training effects transfer to different task contexts, we recorded LFPs for different orientations and signal-to-noise ratios outside the conditioning procedure with a simple fixation task. In this *sensitivity test* (see Methods), a variety of stimuli, including gratings and scenes, were presented for 300 msec, followed by 400 msec fixation, and after it a reward in 50% of the presentations. *Sensitivity tests* were conducted before the training phase of the experiment, and before and after the reversal [after the 28th, 31st, 37th, 47th, and 53rd days of the first part of the training, and after the 45th and 51st days of the second part of the training (after the reversal) for M1; after the 21st and 34th days of the first part of the training, and after the 20th day of the second part of the training (after reversal) for M2]. Figure 8 shows the LFPs averaged across electrodes and the 20%, 40%, and 80% signal-to-noise gratings before the training, during the first part of the training, and after the reversal. In both animals, the overall LFP amplitude increased for the rewarded compared to the unrewarded orientation during the training and this difference decreased after the reversal. To quantify the changes in power in the *sensitivity tests*, we used the same power index as for the analysis of the data obtained during the conditioning. For this analysis, we concentrated on the gamma band because this may reflect activity in the local cortical network (Belitski et al., 2008). The power indices were computed for the mean power values of the middle and high gamma band (60–150 Hz) during the stimulus presentation time (1–300 msec) before the training, during the first part of the training, and after the reversal for the 10%, 20%, 40%, and the 80% signal-to-noise gratings (Figure 9). The power indices before conditioning, during the first part of training, before the reversal, and after the reversal were compared across monkeys and signal-to-noise ratios in repeated measures ANOVA. The effect of testing period was significant (main effect: $F = 10.89$, $p < .05$). The power index in the 60–150 Hz band

track neural activity across daily sessions in our chronic recording preparation, we were able to show dissociable neural effects of stimulus–reward pairing, each having a distinct time course. The difference between the two effects became apparent when we reversed the stimulus–reward pairing. After the reversal of the stimulus–reward pairing, the later part of the LFP response switched quickly, whereas the earlier part of the response changed more slowly. In this early interval, the response to the previously rewarded stimulus stayed high compared to that of the previously unrewarded stimulus for a much greater number of postreversal sessions than was found for the later interval. In addition to this marked dissociation in time course, the two effects also had a different signature in the frequency domain: The power at both low and high frequencies increased with training in the early response interval, whereas the late interval effect consisted of an increase in power in high frequencies (gamma) and a decrease in low frequencies.

The early and late response changes were present in both animals. However, there were also differences in the LFPs between the two animals. The latency of the first response phase was shorter in M2 than in M1, which might be related to the larger and more visible stimulus in the training of M2. The late interval response phases and the initial stimulus unselective response peaks (at about 50 msec) were of opposite polarity in the two animals. It should be noted that previous studies have also reported interindividual differences in visual cortical LFPs in monkeys (Anderson, Mruzek, Kawasaki, & Sheinberg, 2008; Taylor, Mandon, Freiwald, & Kreiter, 2005) and humans (Yoshor, Ghose, Bosking, Sun, & Maunsell, 2007). Possible causes of these interanimal variations in the shape of the LFP include differences in recording location (e.g., layer; Anderson et al., 2008), that is, foveal versus peripheral stimulus presentations or different cortical layer, in situ electrical properties of electrodes and, although less likely, differences in the location of the reference wires. It should be stressed that despite these quantitative differences between the two monkeys in the LFP, the recording location, and the stimulus parameters, both animals showed clear evidence for distinct early and late response intervals and dissociations of how these response intervals changed with stimulus–reward pairing and with reversal of these pairings. This consistent pattern of results cannot be explained as a function of recording artifacts. Indeed, qualitatively identical effects of conditioning were observed in the two animals.

Effects of stimulus–reward associations on neural responses have been observed in several monkey cortical areas, including inferior temporal area TE and perirhinal cortex (Mogami & Tanaka, 2006), lateral intraparietal area (e.g., Sugrue, Corrado, & Newsome, 2004; Coe, Tomihara, Matsuzawa, & Hikosaka, 2002; Platt & Glimcher, 1999), and prefrontal cortex (e.g., Leon & Shadlen, 1999). These are high-level cortices, and thus far, no reward-related effects have been reported in monkey early visual cortical

areas. Shuler and Bear (2006) documented responses that predicted reward timing in primary visual cortex of rats, but as yet, it is unclear to what degree the functional properties of rat and monkey visual cortex are comparable, especially regarding extra-retinal response modulations.

Few studies have investigated the effect of long-term stimulus–reward pairing on sensorial responses as in the present study. Sasaki and Gemba (1982) documented increases in intracortical evoked potentials during the learning of a motor response cued by a light stimulus (i.e., with increasing stimulus–reward correlation). However, the marked increases in evoked potentials on visual cortical electrodes might have reflected enhanced fixation of the stimulus in the free-viewing conditions or enhanced attention during the course of training. Salazar, Kayser, and Konig (2004) observed a larger LFP amplitude and gamma power for the rewarded compared to the orthogonal unrewarded orientation in visual areas 17, 18, and 21. This result in cat visual cortex is similar to our result in the monkey, except that the animals in the study of Salazar et al. (2004) were required to emit a behavioral response in order to get reward (i.e., operant instead of classical conditioning). Blake et al. (2006) found an increase in spiking activity in auditory cortex of awake monkeys after classical conditioning, but this was specific to the frequency of the rewarded tone in only one of the two trained animals. Selective spiking activity to rewarded tones was found in both monkeys after these were trained in an operant conditioning task. Whether the less consistent, stimulus-selective effect in the Blake et al. study compared to our study is due to differences in species (old-world vs. new-world monkeys), modality (visual vs. auditory cortex), neural activity measure (LFPs vs. spiking activity), or extent of training is unclear.

During the course of the conditioning, the LFP amplitude increased not only for the rewarded but also, although less so, for the unrewarded and background noise stimulus. Thus, one can distinguish an orientation-specific component and an orientation-unspecific component of the LFP change. The latter unspecific component might be due to reward-induced increases in responses of weakly orientation-tuned neurons.

Our study is, to our knowledge, the first one to show two dissociable forms of visual neural response changes with long-term stimulus–reward pairing. The fact that the time course of the effect of stimulus–reward pairing differed between the early and late response intervals indicates that different mechanisms must underlie these changes in neural responses. The fast reversal of the late interval LFP change and its correlation with the behavior (eye movement) suggest that this late LFP change is related to enhanced attention during presentations of the rewarded stimulus. Studies of spatially selective attention in V4 reported an increase in power in the gamma range and a decrease for lower frequencies with attention (Fries, Womelsdorf, Oostenveld, & Desimone, 2008; Fries, Reynolds, Rorie, & Desimone, 2001), which is similar

to what we observed here in the late response interval after reversing the stimulus–reward pairing. In the classical conditioning protocol used in the present study, the rewarded stimulus might increase attention in the sense of alerting the animal that reward is imminent and this alerting factor might underlie the changes in the late interval response. Such a mechanism requires first an identification of the stimulus as being the rewarded one, explaining the relatively long latency of this effect. Note that this alerting attention mechanism differs from the stimulus-selective attention mechanisms studied by Fries et al. and others.

We cannot compare directly the LFP responses during the conditioning and sensitivity tests because the stimulus context differs between the two paradigms (e.g., during the training, the stimuli were presented in long sequences of presentations of noise). Nonetheless, we observed changes in the late part of the LFP (around 250–300 msec poststimulus onset) during the sensitivity test. These followed the same trend as observed in the second interval during the first part of the conditioning: in M1, a significantly larger late interval LFP for the 20% signal-to-noise ratio unrewarded compared to the rewarded stimuli (Kruskal–Wallis test, $p < .0001$; data not shown), and in M2, a nonsignificant tendency for a larger late LFP for the 20% signal-to-noise ratio rewarded stimuli (data not shown; also see Figure 8). Importantly, the late LFP effects in the sensitivity test did not reverse after the reversal of the stimulus–reward pairing, unlike in the conditioning task. Thus, the late interval changes observed in conditioning task were certainly not that prominent in the sensitivity test, which agrees with the attention interpretation of this late interval LFP effect.

For Int1, the response difference between orientations ameliorated slowly after reversal, and this change was largely due to response increases in the newly rewarded orientation. The slower time course of the Int1 compared to the Int2 response changes rules out that the former are related to attention. It is tempting to speculate that the slower Int1 response changes are related to the long-lasting behavioral effects that have been described in perceptual learning (Ghose et al., 2002; Schoups et al., 2001; Vogels & Orban, 1985), and thus, may reflect a sensorial enhancement of the stimulus representation. We do not know whether the monkeys showed an improvement in the sensorial thresholds for the rewarded stimulus during the course of learning as we did not measure perceptual thresholds for the grating stimuli before and after the conditioning. Such behavioral testing is not trivial because one wants to avoid associations of stimuli and reward during the testing. However, a recent human psychophysical study using the same stimuli, fluid reward, and conditioning procedure (Seitz, Kim, & Watanabe, 2009) showed an improvement in detection of the rewarded stimulus, even when the latter could not be attended to during conditioning. It is well possible that the Int1 increase in LFP that we documented in the present study underlies the improved discrimination of the rewarded

stimulus found by Seitz et al. (2009) in human subjects after classical conditioning.

In the present study, we measured LFPs and not spiking activity. LFPs reflect mainly synaptic potentials of a large population of neurons, whereas spiking activity corresponds to the output of a much smaller population of neurons and, in the case of single-cell recording, the output of a single neuron. One cannot exclude a priori that part of the LFP signals reflect activity picked up by the reference wire instead of activity changes in area V4. Thus, the question arises whether the effect of stimulus–reward pairing that we observed here for the LFPs would also present for spiking activity. Recent reports that related spiking activity and LFPs in different frequency bands in visual cortex have found significant signal (i.e., across stimuli) correlations between multiunit spiking activity and LFPs in the gamma band (Belitski et al., 2008; Maier et al., 2008). This correlation between spiking activity and LFP appears to be particularly strong in the high gamma band above 100 Hz (Gieselmann & Thiele, 2008; Viswanathan & Freeman, 2007; also see Ray, Hsiao, Crone, Franaszczuk, & Niebur, 2008 for similar results in awake monkey somatosensory cortex), perhaps because of the intrusion of low-frequency components of the action potentials into the LFP at these high gamma frequencies. We found in both animals significant increases and decreases in the >100 Hz gamma band power with stimulus–reward pairing before and after the reversal, respectively (Table 1), which might suggest that similar effects are also present in spiking activity. However, this should be tested directly recording spiking activity.

In conclusion, we have shown that simple stimulus–reward pairings outside the context of a task are sufficient to strengthen stimulus representations in visual cortex. This selective enhancement, perhaps by means of feedback from late visual cortical or extravisual cortical areas and/or by means of neuromodulatory factors (Bao, Chan, & Merzenich, 2001), can serve to promote learning specifically for stimuli that are paired with reward (Seitz et al., 2009; Seitz & Dinse, 2007).

Acknowledgments

Funding from NIH (R21 EY017737), NSF (BCS-0549036), GSKE, GOA (2005/18), IUAP (P6/29), EF (LCCC EF/05/014), NeuroProbes (IST-2004-027017), Crea (07/027), and HFSP (RGP 18/2004) is gratefully acknowledged. We thank Dr. Matthew Fellows for his indispensable help during the array implantation surgery, and Marc De Paep, Wouter Depuydt, Piet Kayenbergh, Gerrit Meulemans, Inez Puttemans, Kirsten Vanderheyden, Stijn Verstraeten for their technical assistance, and Olivier Joly for his help with the analyses.

Reprint requests should be sent to Rufin Vogels, Laboratorium voor Neuro- en Psychofysiologie, K.U. Leuven Medical School, 3000 Leuven, Belgium, or via e-mail: Rufin.Vogels@med.kuleuven.be.

REFERENCES

- Ahissar, M., & Hochstein, S. (1993). Attentional control of early perceptual learning. *Proceedings of the National Academy of Sciences, U.S.A.*, *90*, 5718–5722.

- Anderson, B., Mruzec, R. E., Kawasaki, K., & Sheinberg, D. (2008). Effects of familiarity on neural activity in monkey inferior temporal lobe. *Cerebral Cortex*, *18*, 2540–2552.
- Bao, S., Chan, V. T., & Merzenich, M. M. (2001). Cortical remodelling induced by activity of ventral tegmental dopamine neurons. *Nature*, *412*, 79–83.
- Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M. A., Logothetis, N. K., et al. (2008). Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *Journal of Neuroscience*, *28*, 5696–5709.
- Blake, D. T., Heiser, M. A., Caywood, M., & Merzenich, M. M. (2006). Experience-dependent adult cortical plasticity requires cognitive association between sensation and reward. *Neuron*, *52*, 371–381.
- Coe, B., Tomihara, K., Matsuzawa, M., & Hikosaka, O. (2002). Visual and anticipatory bias in three cortical eye fields of the monkey during an adaptive decision-making task. *Journal of Neuroscience*, *22*, 5081–5090.
- Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: Experience and attention in primary visual cortex. *Nature Neuroscience*, *4*, 519–525.
- Desimone, R., Schein, S. J., Moran, J., & Ungerleider, L. G. (1985). Contour, color and shape analysis beyond the striate cortex. *Vision Research*, *25*, 441–452.
- Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*, 1035–1045.
- Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, *291*, 1560–1563.
- Fries, P., Womelsdorf, T., Oostenveld, R., & Desimone, R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *Journal of Neuroscience*, *28*, 4823–4835.
- Ghose, G. M., Yang, T., & Maunsell, J. H. (2002). Physiological correlates of perceptual learning in monkey V1 and V2. *Journal of Neurophysiology*, *87*, 1867–1888.
- Gieselmann, M. A., & Thiele, A. (2008). Comparison of spatial integration and surround suppression characteristics in spiking activity and the local field potential in macaque V1. *European Journal of Neuroscience*, *28*, 447–459.
- Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585–612.
- Ito, M., Westheimer, G., & Gilbert, C. D. (1998). Attention and perceptual learning modulate contextual influences on visual perception. *Neuron*, *20*, 1191–1197.
- Leon, M. I., & 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.
- Li, W., Piech, V., & Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience*, *7*, 651–657.
- Logothetis, N. K. (2003). The underpinnings of the BOLD functional magnetic resonance imaging signal. *Journal of Neuroscience*, *23*, 3963–3971.
- Maier, A., Wilke, M., Aura, C., Zhu, C., Ye, F. Q., & Leopold, D. A. (2008). Divergence of fMRI and neural signals in V1 during perceptual suppression in the awake monkey. *Nature Neuroscience*, *11*, 1193–1200.
- Mogami, T., & Tanaka, K. (2006). Reward association affects neuronal responses to visual stimuli in macaque TE and perirhinal cortices. *Journal of Neuroscience*, *26*, 6761–6770.
- Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology*, *82*, 2490–2502.
- Pasupathy, A., & Connor, C. E. (2001). Shape representation in area V4: Position-specific tuning for boundary conformation. *Journal of Neurophysiology*, *86*, 2505–2519.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*, 233–238.
- Polley, D. B., Steinberg, E. E., & Merzenich, M. M. (2006). Perceptual learning directs auditory cortical map reorganization through top-down influences. *Journal of Neuroscience*, *26*, 4970–4982.
- Raiguel, S., Vogels, R., Mysore, S. G., & Orban, G. A. (2006). Learning to see the difference specifically alters the most informative V4 neurons. *Journal of Neuroscience*, *26*, 6589–6602.
- Rainer, G., Lee, H., & Logothetis, N. K. (2004). The effect of learning on the function of monkey extrastriate visual cortex. *PLoS Biology*, *2*, E44.
- Ray, S., Hsiao, S. S., Crone, N. E., Franaszczuk, P. J., & Niebur, E. (2008). Effect of stimulus intensity on the spike-local field potential relationship in the secondary somatosensory cortex. *Journal of Neuroscience*, *28*, 7334–7343.
- Roelfsema, P. R., & van Ooyen, O. A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Computation*, *17*, 2176–2214.
- Salazar, R. F., Kayser, C., & Konig, P. (2004). Effects of training on neuronal activity and interactions in primary and higher visual cortices in the alert cat. *Journal of Neuroscience*, *24*, 1627–1636.
- Sasaki, K., & Gemba, H. (1982). Development and change of cortical field potentials during learning processes of visually initiated hand movements in the monkey. *Experimental Brain Research*, *48*, 429–437.
- Schein, S. J., & Desimone, R. (1990). Spectral properties of V4 neurons in the macaque. *Journal of Neuroscience*, *10*, 3369–3389.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, *412*, 549–553.
- Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Sciences*, *9*, 329–334.
- Seitz, A. R., & Dinse, H. R. (2007). A common framework for perceptual learning. *Current Opinion in Neurobiology*, *17*, 148–153.
- Seitz, A. R., Kim, D., & Watanabe, T. (2009). Rewards evoke learning of unconsciously processed visual stimuli in adult humans. *Neuron*, *61*, 700–707.
- Seitz, A. R., & Watanabe, T. (2003). Psychophysics: Is subliminal learning really passive? *Nature*, *422*, 36.
- Shuler, M. G., & Bear, M. F. (2006). Reward timing in the primary visual cortex. *Science*, *311*, 1606–1609.
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, *304*, 1782–1787.
- Tallon-Baudry, C., & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*, *3*, 151–162.
- Taylor, K., Mandon, S., Freiwald, W. A., & Kreiter, A. K. (2005). Coherent oscillatory activity in monkey area v4 predicts successful allocation of attention. *Cerebral Cortex*, *15*, 1424–1437.

- Viswanathan, A., & Freeman, R. D. (2007). Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience*, *10*, 1308–1312.
- Vogels, R., & Orban, G. A. (1985). The effect of practice on the oblique effect in line orientation judgments. *Vision Research*, *25*, 1679–1687.
- Watanabe, T., Nanez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, *413*, 844–848.
- Yang, T., & Maunsell, J. H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *Journal of Neuroscience*, *24*, 1617–1626.
- Yoshor, D., Ghose, G. M., Bosking, W. H., Sun, P., & Maunsell, J. H. (2007). Spatial attention does not strongly modulate neuronal responses in early human visual cortex. *Journal of Neuroscience*, *27*, 13205–13209.
- Yotsumoto, Y., & Watanabe, T. (2008). Defining a link between perceptual learning and attention. *PLoS Biology*, *6*, e221.
- Zeki, S. M. (1978). Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex. *Journal of Physiology*, *277*, 273–290.