Effects of Visual Experience on the Representation of Objects in the Prefrontal Cortex

Gregor Rainer*† and Earl K. Miller
Center for Learning and Memory
RIKEN-MIT Neuroscience Center and
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Summary

The perception and recognition of objects are improved by experience. Here, we show that monkeys' ability to recognize degraded objects was improved by several days of practice with these objects. This improvement was reflected in the activity of neurons in the prefrontal (PF) cortex, a brain region critical for a wide range of visual behaviors. Familiar objects activated fewer neurons than did novel objects, but these neurons were more narrowly tuned, and the object representation was more resistant to the effects of degradation, after experience. These results demonstrate a neural correlate of visual learning in the PF cortex of adult monkeys.

Introduction

Learned behaviors typically benefit from practice. Motor skills become more accurate, efficient, and automatic and can be executed more quickly. Visual processing likewise benefits from experience. Familiarity with a set or class of objects can lead to an increase in our accuracy in discriminating and recognizing them, particularly in situations when the objects are occluded or otherwise difficult to discriminate. However, while much is understood about how sensory visual attributes are represented in neural activity, relatively little is known about how experience improves their processing.

One robust effect of visual experience is the waning of neuronal responsiveness as an initially novel stimulus becomes familiar. This seems to be a widespread phenomenon, evident in several cortical visual areas, including the inferior temporal (IT) cortex (a brain region important for object recognition) (Gross, 1973; Tanaka, 1996; Logothetis, 1998; Miyashita and Hayashi, 2000) and the prefrontal (PF) cortex (which plays a central role in visual cognition) (Goldman-Rakic, 1995; Miller, 1999). It has been suggested that this reflects a process that winnows the neural ensemble activated by a new stimulus, leaving behind fewer neurons that optimally represent the now familiar object (Li et al., 1993; Desimone, 1996; Wiggs and Martin, 1998). Support for this comes from observations that some IT neurons show an apparent increase in responsiveness to familiar stimuli, at least with weeks or months of experience (Kobatake et al., 1998). Establishing a link between these changes and

improvements in object recognition has remained elusive, however. These effects have not been shown to increase the amount of object information conveyed by neural activity and they have not been directly tied to increased ability of monkeys to recognize objects.

To provide evidence for such a link, we examined the lateral PF cortex, a brain region that receives inputs from the IT cortex (Barbas and Mesulam, 1985; Pandya and Yeterian, 1990; Petrides and Pandya, 1999) and plays a central role in behaviors that involve discriminating, remembering, and making decisions about visual stimuli (Fuster, 1997). Damage to the PF cortex disrupts a variety of visually guided behaviors (Petrides, 1994), and, like IT neurons, many lateral PF neurons are selectively activated by complex objects (Miller et al., 1996). Frontal cortex has been particularly implicated in the processing of novel information and is critical for learning a wide variety of behaviors (Chen and Wise, 1995; Knight et al., 1995; Asaad et al., 1998; Parker et al., 1998). Thus, it is a good candidate for studying the effects of experience.

We used a modified version of a delayed matching-to-sample (DMS) task that required monkeys to discriminate and remember each of a set of five objects (see Experimental Procedures). Monkeys were briefly shown a sample object, then after a short delay, a test object. If the test object matched the sample, monkeys were required to release a bar to receive a juice reward. Monkeys can typically perform this task very well, making correct judgments on 90% or more of the trials, leaving little room for improvement with practice. We made the task more difficult by parametrically degrading the sample objects by interpolation with noise patterns (Figure 1). Monkeys had to recognize which object was present during the sample period and release a lever if that object appeared during the test period. Similar parametric behavioral paradigms have been previously employed in monkeys for the study of coherent motion (Newsome et al., 1989; Kim and Shadlen, 1999), tactile (Romo et al., 1999), and object (Amit et al., 1997; B. Jagadeesh and R. Desimone, 1997, Soc. Neurosci., abstract) stimuli. To assess the effects of experience, we compared behavior and neural activity during the performance of the task with novel and familiar objects (see Experimental Procedures).

Results

Experience Improved Psychophysical Performance

We found that experience resulted in improved psychophysical performance. The monkeys' performance varied systematically as a function of stimulus level for both novel and familiar objects (Figure 2). When nondegraded objects were used as samples (100% stimulus level), monkeys performed near ceiling with both familiar and novel objects (95% and 92% correct, respectively, p = 0.14, t test). Performance was at chance level at the 0% stimulus level because there was no object information present in the sample, and the monkeys were guessing. At intermediate stimulus levels, however, monkeys performed better with familiar objects than with novel objects (e.g., at the 65% stimulus level, performance was

*To whom correspondence should be addressed (e-mail: gregor. rainer@tuebingen.mpg.de).
†Present address: Max-Planck Institute for Biological Cybernetics, Spemannstrasse 38, D-72076 Tübingen, Germany.
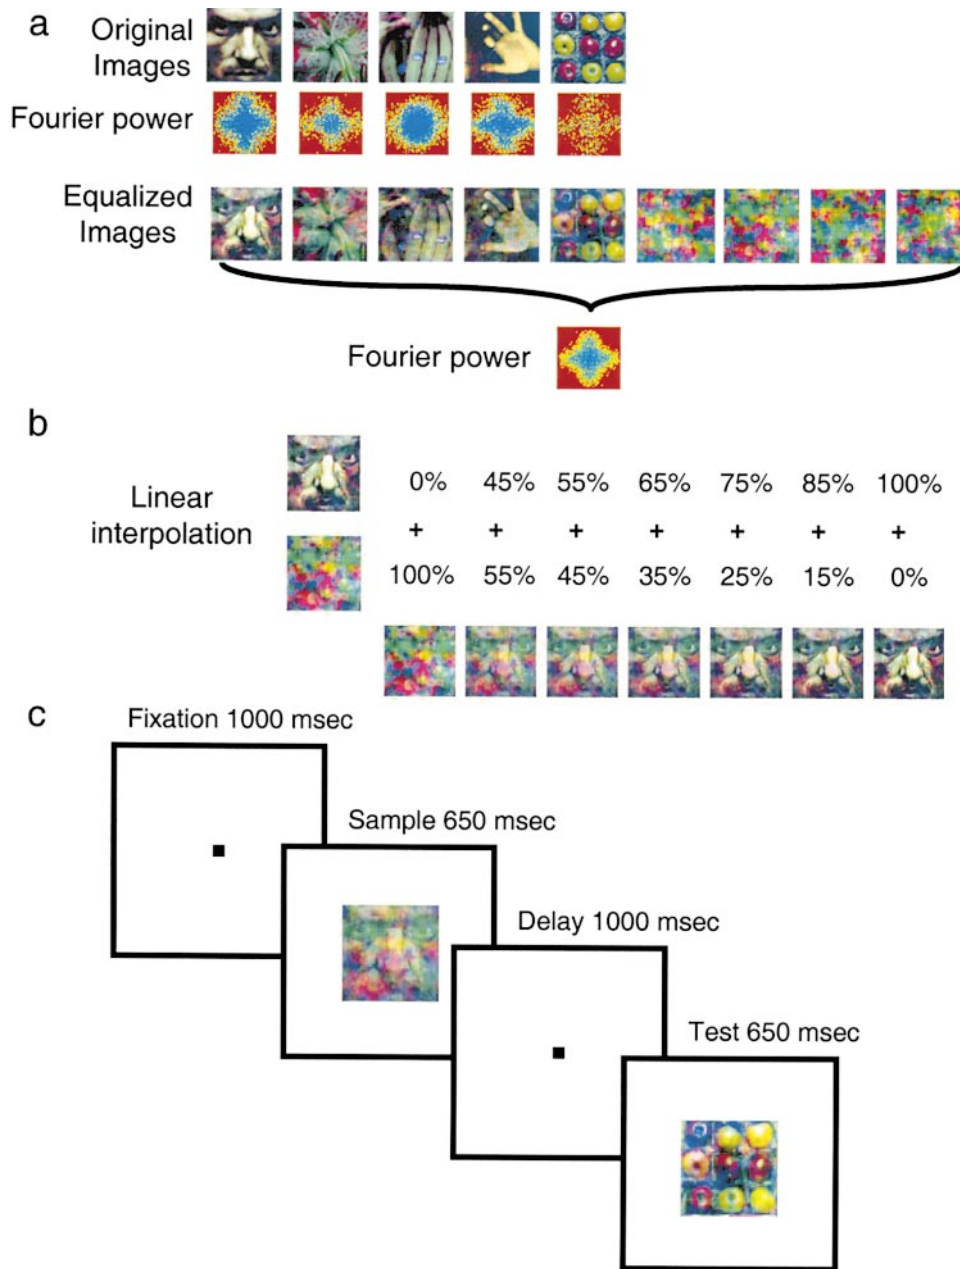
Figure 1. Task and Stimuli

(a) Natural objects and their Fourier power spectra before and after equalization for spatial frequency. Also shown are four examples of random noise patterns matched for spatial frequency and luminance.

(b) Linear (pixel-by-pixel) interpolation that was used to generate intermediate images between the natural images and the noise patterns.

(c) Example of a behavioral trial. A sample object (one of the five objects at one of seven stimulus levels) was presented for 650 ms. After a 1000 ms delay, either a match or a nonmatch object (one of the five objects, always at 100% stimulus) was presented for 650 ms.

82% correct with familiar and 56% correct with novel objects, $p \ll 0.0001$, t test). We attribute this improvement in performance at intermediate stimulus levels to visual experience with the objects.

Neural Activity: Effects of Familiarity

We recorded a total of 324 neurons from the cortices of two monkeys (164 neurons were recorded with novel objects, and 160 neurons with familiar objects; see Figure 3a). We first asked whether there were any general differences in activity depending on whether an object

was novel or familiar. To do this, we examined neural responses to nondegraded (100% stimulus level) objects. The most obvious effect was that familiar objects elicited less overall activity than did novel objects. This is shown in Figure 3b, which illustrates the average activity of all neurons to all novel objects and the corresponding average activity to all familiar objects. During the sample interval and first part of the memory delay, novel objects elicited significantly more activity than did familiar objects ($p = 0.002$, t test). No such difference was apparent in the last part of the delay ($p > 0.1$, t test),
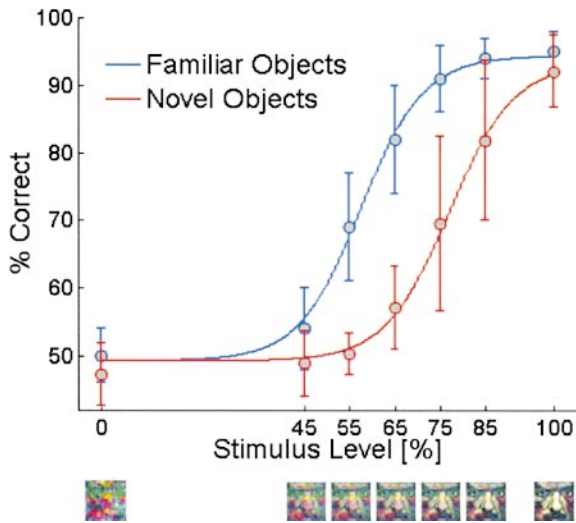
Figure 2. Behavioral Performance

Psychometric functions pooled across two monkeys (familiar objects: 14 sessions, novel objects: 11 sessions) with the corresponding standard deviations.
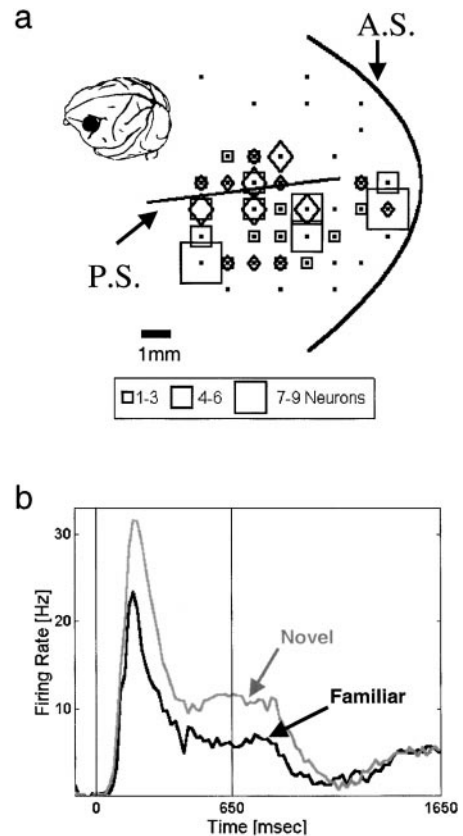

Figure 3. Recording Sites and Average Activity

(a) Electrode penetration sites of object-selective neurons (t tests at 100% stimulus level, $p < 0.05$) recorded in the familiar object experiment (n = 42, diamonds) and the novel object experiment (n = 90, squares). The small black dots indicate penetration sites where one or more non object-selective neurons were successfully isolated. The inset shows the general location of the recording sites on the surface of the brain. Abbreviations: P.S., principal sulcus; A.S., arcuate sulcus.

(b) Average activity plots for the novel and familiar object experiments. Each curve represents the neural activity to nondegraded objects (100% stimulus), averaged across all objects and all neurons recorded in each experiment (novel objects: n = 160, familiar objects: n = 164).

which has been associated with "prospective" coding of a forthcoming stimulus or action (Quintana and Fuster, 1992; Rainer et al., 1999). Average activity during the late delay was, however, significantly above baseline for both familiar and novel objects ($p < 0.01$, t tests).

We next assessed how many neurons showed activity that was dependent on the identity of the sample object. A one-way ANOVA was applied to the activity elicited by nondegraded objects (100% stimulus level; see Experimental Procedures), averaged across the sample and delay intervals. Given the greater overall level of activity elicited by novel objects, it was not surprising to find that a significantly greater proportion of neurons showed selectivity for novel objects (90 of 160, or 56%) than for familiar objects (42 of 164, or 26%, $p \ll 0.0001$, χ^2 test). This was also true late in the delay (novel objects: 78 of 160, or 49%, familiar objects: 37 of 160, or 23%, $p \ll 0.0001$, χ^2 test). It has been hypothesized that as a stimulus becomes familiar, neurons coding features not essential for recognizing it reduce their responses, leaving behind a smaller number of more selective neurons that optimally represent the familiar stimulus. To test whether this was the case for the PF object representation, we computed the depth of selectivity (S) for neurons that showed significant effects on the ANOVA. The depth of selectivity provides a measure of how narrowly tuned a neuron is to a given stimulus set (see Experimental Procedures). We found that neural activity to familiar objects was modestly, but significantly, more narrowly tuned for familiar than for novel objects ($S_{familiar} = 0.30$, $S_{novel} = 0.24$, $p = 0.046$, t test). Thus, familiarity of a set of objects resulted in (1) a reduction in average PF activity and (2) a decrease in the number of neurons involved in representing the objects but (3) narrower tuning to these familiar objects.

Neural Activity: Effects of Degradation

Similar to the behavioral results, familiarity had strong effects on neural activity when the objects were degraded. For example, a single neuron's activity in response to novel objects is shown in Figures 4a–4d. At the 100% stimulus level, this neuron showed greater activity in response to a preferred object than to a nonpreferred object. Differences in activity between the objects were less evident for degraded stimuli, however. To quantify the amount of object information conveyed by the activity, we used standard receiver operating–characteristic (ROC) methods (see Experimental Procedures). Activity was evaluated across time using two methods, a "sliding" and a "cumulative" method. The former provides a measure of the "moment-to-moment" account of object information, whereas the latter provides a measure of the total amount of information conveyed up to a given point in a trial. The data for the example neuron are shown in Figures 4e and 4f. At the 100% stimulus level, the object selectivity values reached around 0.95, indicating that this neuron could distinguish between the two objects on 95% of trials. However, the object selectivity fell off rapidly toward lower stimulus levels. This indicates that the ability of this neuron to convey object information dropped off
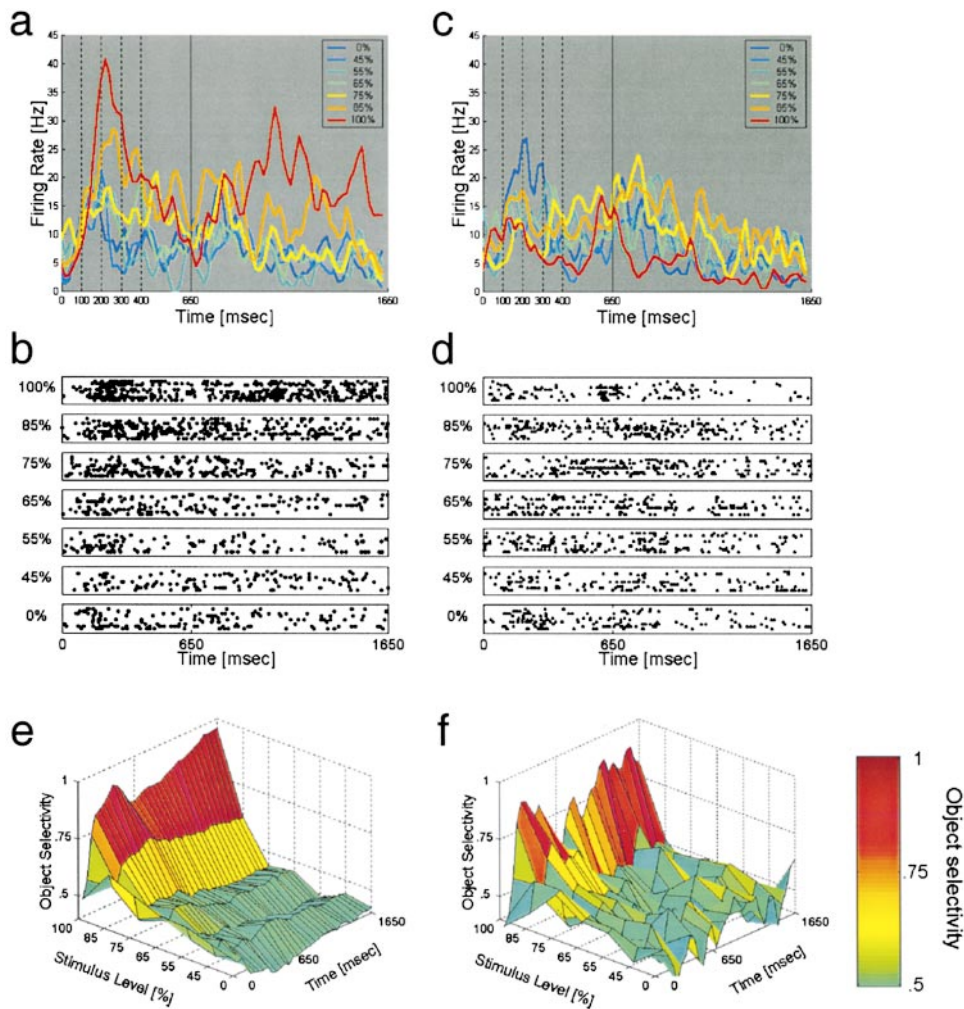
Figure 4. Response of Single Neuron during the Novel Object Experiment

The panels in the top row show the peristimulus time histogram (PSTH) as a function of stimulus level for this neuron's preferred (a) and nonpreferred (c) objects. The sample presentation (0–650 ms) and delay periods (650–1650 ms) are shown. The corresponding rasters are shown below each histogram in (b) and (d). (e) and (f) provide a measure of the object selectivity for this neuron by directly comparing the activity to the preferred and nonpreferred objects at each stimulus level across time (see Experimental Procedures). Values near 0.5 indicate similar neural activity to the two objects, and values near 1 indicate very different neural activity and thus strong object selectivity. (e) was constructed using the cumulative, and (f) using the sliding bin, method.

sharply as the object was degraded. Below a stimulus level of 85%, its object selectivity values were near chance level (0.50), indicating that it was communicating no reliable object information.

A neuron that showed selectivity for familiar objects is shown in Figure 5. Activity differences between a preferred and a nonpreferred object were evident not only at the 100% stimulus level but also when the objects were degraded (Figures 5a–5d). At the 100% stimulus level, this neuron's object selectivity values were high (Figures 5e and 5f), but unlike the neuron shown in Figure 4, they remained high at lower stimulus levels. This indicates that for this neuron, object selectivity was robust with respect to degradation. Even at stimulus levels of 65%–55%, the values were at 0.75 or above. Figure 6 shows object selectivity values for four neurons studied with novel objects (left column) and four neurons studied with familiar objects (right column). All eight neurons

were highly selective when objects were presented at the 100% stimulus level. However, note that for novel objects there was a steep drop in selectivity values for stimulus levels below 100%, while for familiar objects the values formed a broad plateau. This means that selectivity collapsed when novel objects were degraded but remained high when familiar objects were degraded.

To quantify this effect across the entire population of neurons, we compared each neuron's optimal performance at discriminating 100% objects to its optimal performance at discriminating 65% objects (see Experimental Procedures). The 65% level was chosen because familiarity was most beneficial to the monkeys' behavior at that level (Figure 2). For both novel objects (Figures 7a and 7c) and familiar objects (Figures 7b and 7d), most of the data points fall below the diagonal, indicating that PF neurons conveyed more information about nondegraded objects (100%) than about degraded objects
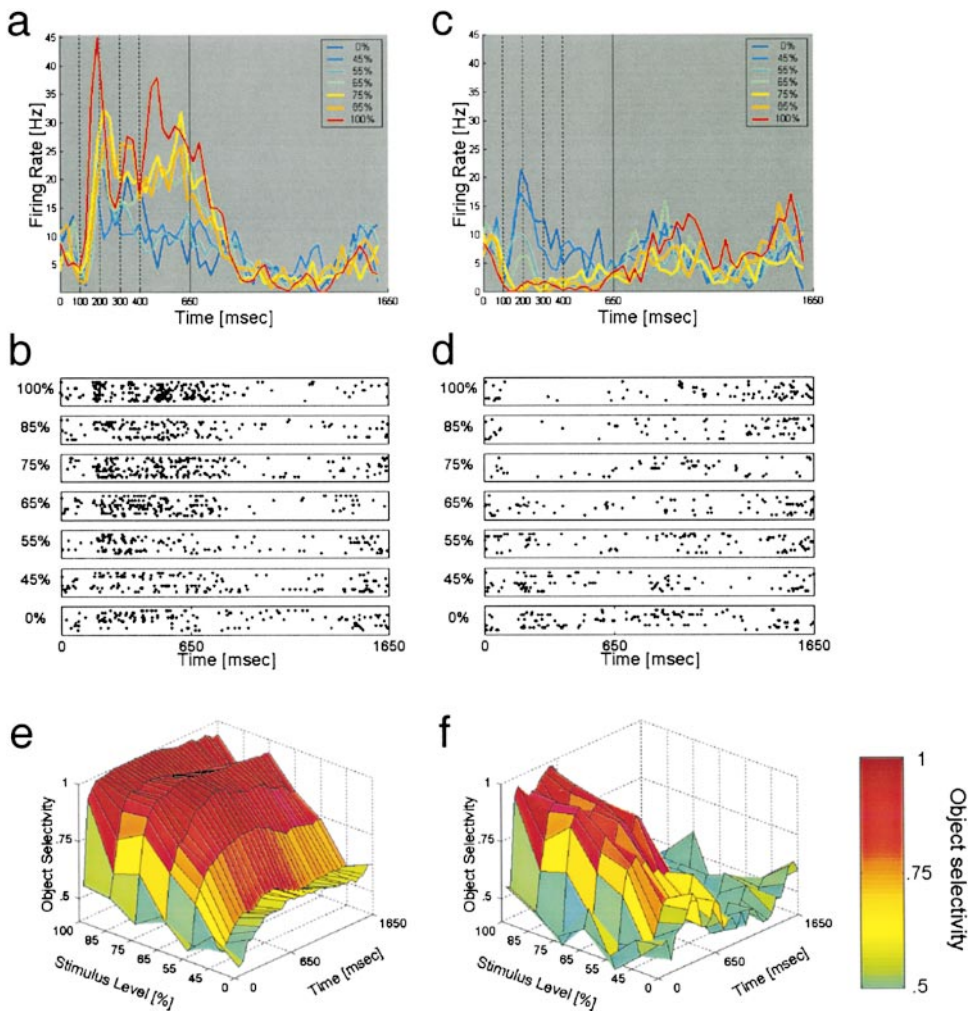
Figure 5. Response of Single Neuron during the Familiar Object Experiment

The panels in the top row show the peristimulus time histogram (PSTH) as a function of stimulus level for this neuron's preferred (a) and nonpreferred (c) objects, as well as the corresponding rasters (b and d) and the object selectivity computed using cumulative (e) and sliding (f) methods. Conventions are identical to Figure 4.

(65%). For novel objects, many neurons were highly selective (exceeded an arbitrary object selectivity value of 0.85) at the 100% stimulus level (sliding: 63 of 160, or 39%, cumulative: 67 of 160, or 42%). At the 65% stimulus level, however, only a small fraction of them maintained their selectivity, i.e., exceeded a corresponding criterion value of 0.75 (sliding: 9 of 63, or 14%, cumulative: 4 of 67, or 6%; Figures 7a and 7c). We used a criterion of 0.75 (rather than 0.85) because behavioral performance was below ceiling at the 65% stimulus level (Figure 2), and we expected this to be reflected in lower neural performance. In contrast to the selectivity for novel objects, fewer neurons were highly selective for familiar objects, but these were better able to discriminate the degraded objects. Of the neurons that were highly selective to familiar 100% objects (sliding: 40 of 164, or 24%, cumulative: 39 of 164, or 24%), the majority retained their selectivity when objects were degraded to the 65% stimulus level (sliding: 31 of 40, or 78%, cumulative: 25 of 39, or 64%). The proportions of neurons exhibiting selectivity for degraded objects

depending on familiarity were highly significantly different ($p < 0.0001$, χ^2 test), and this was also true if only the subset of neurons that were recorded at identical penetration sites and at similar depths was considered (see Experimental Procedures). The fact that we found very few selective neurons in the novel object experiment at 65% stimulus was not merely a consequence of a general lack of selectivity of the neurons studied with novel objects. Indeed, at the 100% stimulus level, more neurons were highly selective for novel than for familiar objects; this selectivity was just not robust with respect to degradation.

The above analysis focuses on neurons that were highly selective at 100% stimulus. The same pattern of results was apparent when we considered the entire population of neurons, i.e., regardless of their selectivity at 100% stimulus. At the 65% stimulus level, fewer neurons had selectivity values of at least 0.75 for novel objects (sliding: 5 of 160, or 3%, cumulative: 7 of 160, or 4%) than for familiar objects (sliding: 30 of 164, or 18%, cumulative: 22 of 164, or 13%, $p < 0.001$, χ^2-test).
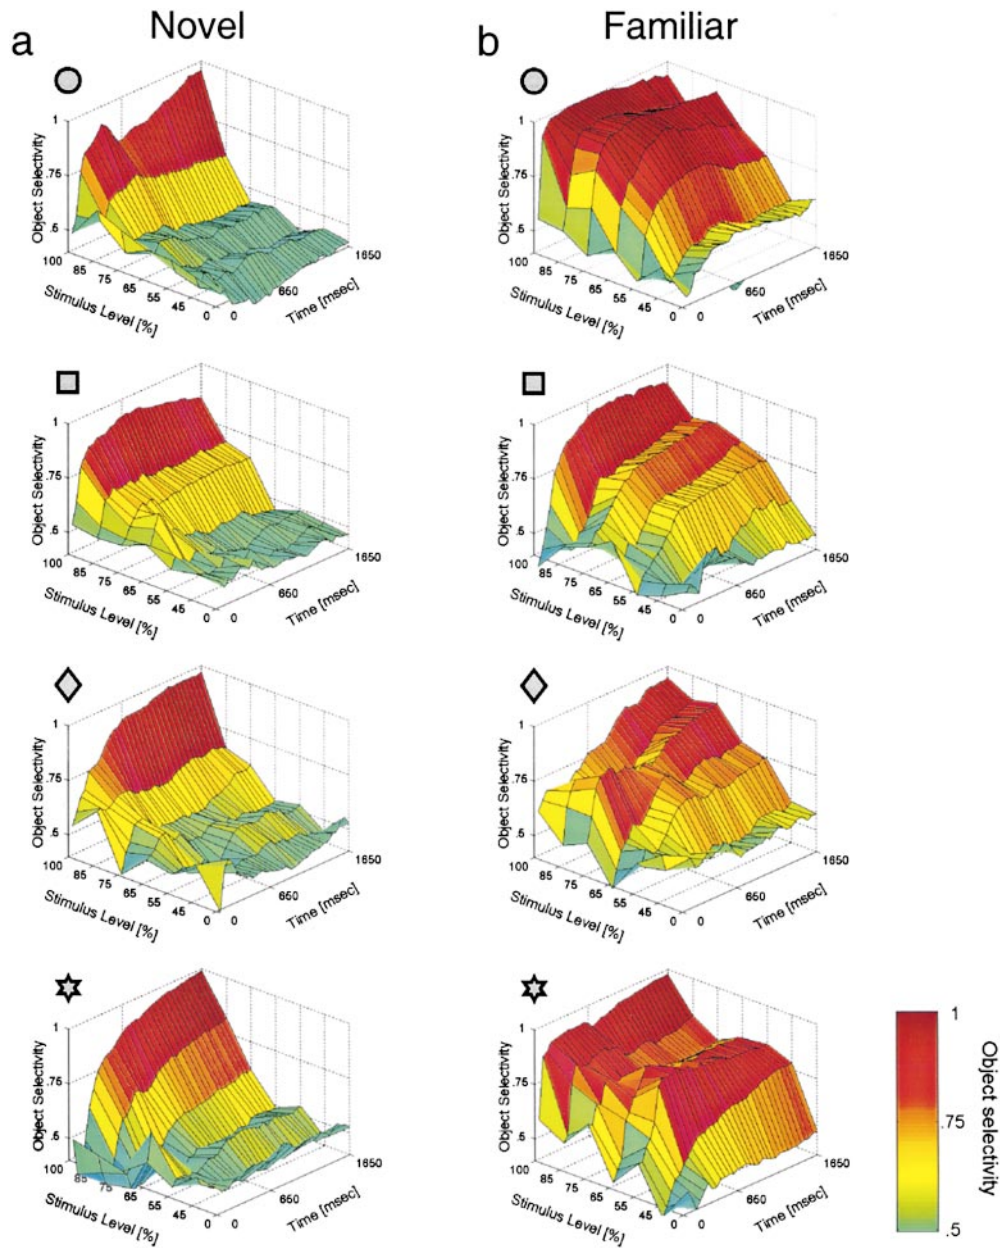
Figure 6. Object Selectivity Plots for Single Neurons Recorded in the Familiar and Novel Object Experiments

The object selectivity for four example neurons recorded in the familiar object experiment (a) and the novel object experiment (b) is shown. The top row corresponds to the examples from Figures 4 and 5. The symbols allow the identification of these example neurons in the population scatterplot shown in Figure 7. Object selectivity was robust to degradation for familiar but not for novel objects.

Familiarity resulted in relatively little improvement in neural performance at higher stimulus levels (100% and 85%), presumably because, like the monkeys' behavior, performance was already high, and there was little room for further improvement. Also, at very low stimulus levels (45% and 0%), object discrimination was very difficult (or impossible), and there was little change in either neural or behavioral performance with familiarity. The same was true for neurons; the largest improvements in neural performance with familiarity were evident at the same stimulus levels for which there was also the largest improvement in behavioral performance. This is illustrated in Figure 8, which shows that the average improvement in the object selectivity values of the highly selective neurons was correlated with the behavioral improvement at each stimulus level (correlation coefficient: $R = 0.93$, $R = 0.95$ for neurons recorded at "common" locations; see Experimental Procedures). The average ROC values for familiar objects were significantly greater than those for novel objects at the three intermediate stimulus levels (55%, 65%, and 75% stimulus, $p < 0.001$, t test). Monkeys also showed a significant improvement in behavior at these levels ($p < 0.001$, t test; Figure 2).
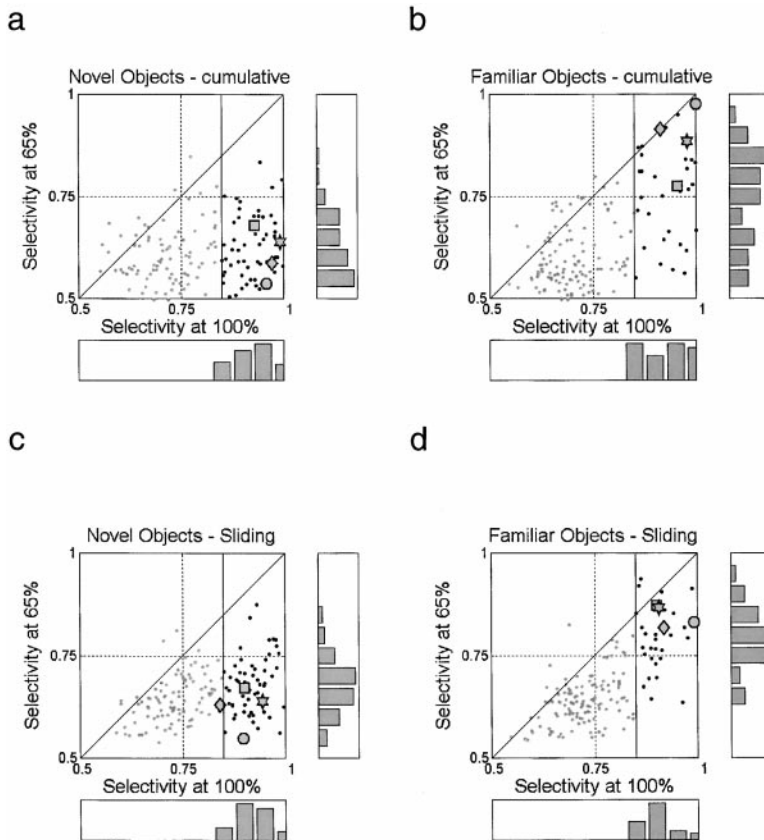
Figure 7. Distribution of Object Selectivity across the Population

These four panels compare the object selectivity at 100% stimulus to 65% stimulus for novel and familiar objects. (a) and (b) were generated using the cumulative, (c) and (d) using the sliding window, technique. Each data point corresponds to a single neuron; the symbols show the locations of the example neurons from Figure 6. For each neuron, the preferred and nonpreferred objects were determined at 100% stimulus. Object selectivity values correspond to the maximum object selectivity observed at any time point. Highly selective neurons (at 100% stimulus) are shown in bold (familiar objects: n = 39, novel objects: n = 67) and lie to the right of a vertical line drawn at 0.85. Data points near the diagonal indicate similar object selectivity at 100% and 65% stimulus, whereas data points near the x axis indicate that object selectivity was not robust to degradation. To the right and below the scatterplots, histograms of the object selectivity for the highly selective neurons are shown.

Discussion

We report a neural correlate of visual learning in the PF cortex. Experience improved monkeys' ability to discriminate a set of objects that were degraded with noise patterns. It also improved the ability of PF neurons to communicate information about object identity. Familiar objects activated a smaller population of PF neurons than did novel objects, but neural activity was more narrowly tuned to familiar objects, and their representation was more robust with respect to stimulus degradation.
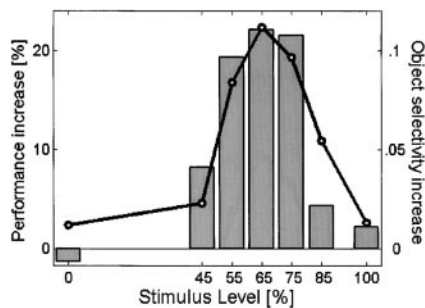
Single PF neurons communicated object information

Figure 8. Increases in Psychophysical and Neural Performance

The relative increases between the novel and the familiar experiments are shown. The bar graph shows changes in behavioral performance (left axis), and the line graph shows the changes in object selectivity (right axis), as a function of degradation.

with high fidelity. The most object-selective neurons recorded in the present study performed at levels comparable to the psychophysical performance of the monkeys. This differs somewhat from results reported for extrastriate area MT (middle temporal) during direction discrimination, in which the average neuron performed similarly to the monkey (Newsome et al., 1989; Britten et al., 1992). This is not surprising since for MT neurons, the stimulus was adjusted to be optimal for the neuron under study (in terms of receptive field location and preferred direction), whereas we did not and could not do this in the present study because learning was specific to the studied objects.

Effects of experience on neural activity have been extensively studied in the somatosensory (Buonomano and Merzenich, 1998) and auditory (Weinberger, 1995) systems. In vision, effects of experience have been reported in the MT (Zohary et al., 1994), as well as in the IT, cortex, where familiarity with objects can lead to both decreases and increases in activity at short (Rolls et al., 1989; Li et al., 1993), as well as long, timescales (Fahy et al., 1993; Kobatake et al., 1998). These and other studies have suggested that IT cortex is modifiable by experience (Sakai and Miyashita, 1991; Logothetis et al., 1995). However, changes in activity in the IT cortex have not been shown to improve processing capacity, nor have they been directly linked to improved behavioral performance. Indeed, one study has reported a lack of an effect of training on IT neurons (Vogels and Orban, 1994). We report here that neurons in the lateral PF cortex can indeed improve their ability to communicate object information with experience and that this

improvement is correlated with improved behavioral performance.

It is possible that the effects reported in this study are at least in part due to plasticity within the PF cortex. Indeed, experience has been correlated with the differential activation of the human frontal lobe (Ranganath and Paller, 1999). Training on visuomotor tasks has been shown to modulate neural activity in the supplementary eye field (Chen and Wise, 1995) and the lateral PF cortex (Asaad et al., 1998), as well as to induce color selectivity in neurons in the frontal eye field (FEF) (Bichot et al., 1996). Similar to these results in the FEF, activity differences as a function of familiarity were apparent early after the onset of the visual response in our study. This suggests that the effect of familiarity was a result of experience-dependent plasticity rather than attentional modulations since attention does not affect the initial, phasic visual response (Rainer et al., 1998; Schall and Thompson, 1999). A given object-selective PF neuron may receive inputs from many feature-sensitive neurons located in extrastriate and temporal cortex. We suggest that through experience, selective strengthening of connections occurs between PF neurons and visual system neurons that best discriminate a given set of objects. The end result may then be a small population of highly selective neurons that robustly represent the familiar objects.

The capacity of humans to improve their psychophysical performance has been studied using a variety of tasks (Goldstone, 1998; Wallis and Bulthoff, 1999), including discrimination of orientation (Karni and Sagi, 1993; Fahle et al., 1995; Dosher and Lu, 1999; Sigman and Gilbert, 2000), brightness (Ito et al., 1998; Ito and Gilbert, 1999), motion (Ball and Sekuler, 1982, 1987), and complex objects (Gold et al., 1999; Furmanski and Engel, 2000). It has been suggested that synaptic plasticity at very early stages of visual processing—possibly even in primary visual cortex—might be responsible for some of these effects (Gilbert, 1996; Sagi and Tanne, 1994) because learning is often specific to the trained location and stimuli. Recent evidence suggests that visual learning may also retune intermediate areas by modifying connections from low-level sensory to high-level decision structures (Dosher and Lu, 1999). Our results are compatible with both of these views; we do not know where in the visual processing hierarchy learning effects first appear. This question will have to be addressed by further experiments. Similar to the present results in monkeys, human performance on recognition tasks in the presence of noise can increase dramatically with practice as a result of enhanced signal strength (Gold et al., 1999). The increase in robustness with respect to degradation that we observed at intermediate stimulus levels is consistent with this and suggests a possible neural substrate for this kind of learning in humans as well.

In sum, our results show that visual experience with objects has strong effects on both behavioral performance and neural activity in the PF cortex. It has been hypothesized that experience winnows the neural ensemble representing a given object, leaving behind a smaller group of neurons that optimally represent the learned object (Li et al., 1993). The present study supports this view and further describes an experience-dependent increase in the robustness of the representation for degraded inputs. This makes processing less susceptible to noise and may reflect a general form of learning in the adult brain that enables it to optimize interactions with the environment.

Experimental Procedures

Recording Technique

Two adult rhesus monkeys, *macaca mulatta* (monkey A: female, 8 kg, monkey B: male, 11 kg), participated in the experiments. Eye coil, head restraint, and recording chamber were implanted under aseptic conditions while the animals were anesthetized with isoflurane. The animals received postoperative antibiotics and analgesics and were always handled in accordance with National Institutes of Health guidelines and the recommendations of the Massachusetts Institute of Technology Animal Care and Use Committee. During recording sessions, monkeys were seated in primate chairs within sound-attenuating enclosures (Crist Instruments, Damascus, MD). Their heads were restrained, and a juice spout was placed at their mouths for automated delivery of reward (apple juice). We employed a grid system (Crist Instruments, Damascus, MD) with custom made modifications that enabled us to use eight single tungsten electrodes (catalog number UEWLGDSMNN1E, 3 μm tip diameter, FHC, Bowdoin, ME) simultaneously. Penetrations were made perpendicular to the surface of the skull, and the minimum separation between sites was 1 mm. Recording sites were localized using magnetic resonance imaging and consisted of the lateral PF cortex around and ventral to the principal sulcus. Neural waveforms were stored on disk and sorted into clusters offline (DataWave Instruments, Longmont, CO). We did not prescreen neurons for task-related responses. Instead, we advanced each electrode until the activity of one or more neurons was well isolated and then began data collection. This procedure was used to ensure an unbiased estimate of PF activity. Monkeys completed an average of 865 trials during each of the 25 recording sessions, resulting in about 25 repetitions for each object at each of seven stimulus levels (0%–100%; see Figure 1). Each session lasted $\sim$3 hr.

Recording Locations

Due to the number of conditions required and limitations on the number of trials a monkey can work on a given day, it was in general not possible to complete both the familiar and the novel object experiment during a single recording session. The two experiments are thus based on different neuronal samples. This was necessary because the familiarity effect took several days to appear, and thus comparisons could not be made for the same neurons and objects. However, care was taken to record neurons in the same locations (Figure 3a) and depths in both experiments. Further, the use of multiple electrodes allowed rapid acquisition of data, and thus experiments with novel and familiar objects were completed within 2 weeks. To ensure that our results are not a consequence of sampling different portions of PF cortex in the novel and familiar object experiments, we identified ten "common" locations, where we successfully isolated selective neurons in both experiments (see Figure 3a). We found that most (33 of 42, 79%) of the neurons studied with familiar objects and 36 of 90 (40%) of neurons studied with novel objects were in fact recorded at such "common" locations. There was no significant difference in recording depths at each of the ten recording sites (p $>$ 0.1, t tests). We repeated the major analyses for this subset of neurons and found that results were virtually identical to those obtained when the entire population was considered.

Behavioral Task

The behavioral paradigm was a modified version of DMS. Each trial began when the monkey grabbed a metal lever and initiated fixation on a small spot of light at the center of the screen. After 1000 ms of fixation, a sample object was presented for 650 ms. This sample object could be any of the five objects at any of the seven stimulus levels (0%, 45%, 55%, 65%, 75%, 85%, or 100%). After a 1000 ms delay, a test object was presented for 650 ms. The test object was always at 100% stimulus and could be any one of the five objects. If on a given trial, the sample object was identical to or was a degraded version of the test object, monkeys had to release a lever to obtain a juice reward. In the case of the 0% stimulus level, monkeys were randomly rewarded for releasing the lever on half of the

trials because the sample image contained none of the objects, and monkeys had to simply guess. If the test object did not match the sample object, monkeys had to hold the lever for the entire 650 ms of test object presentation, and a second delay followed that was always terminated by a correct match requiring a lever release. Monkeys did not need to retain object information in memory during this second delay; it was included only to ensure a behavioral response on every trial and was not included in any of the analyses. Match and nonmatch trials occurred equally often, and monkeys had no way of predicting whether a particular trial would be a match or a nonmatch. Prior to the participation in the present study, both monkeys had extensive experience with a standard DMS task using undegraded stimuli over a period of about 2 years.

Stimuli

To generate a stimulus set for one recording session, we started by selecting five images of natural objects at 24-bit color depth, adjusted to 50 by 50 pixels in size. First, the mean intensity of each object was adjusted for each RGB channel to 100/255. Then we computed the Fourier power spectrum (FPS) for each object using the fast Fourier transform (FFT) and calculated an average FPS by averaging the five individual ones from each of the objects. This average FPS had the $1/f^{\alpha}$ spatial frequency dependence that is characteristic of natural images (Field, 1987; Tadmor and Tolhurst, 1993; Brady, 1997). The average FPS was used together with the unchanged Fourier phase to generate five new objects using inverse FFT. This procedure was applied to each RGB channel independently. The objects created in this way were thus equalized for intensity and Fourier power. Next, random noise patterns were obtained by inverse FFT of random phase distributions (values ranging from $-\pi$ to $+\pi$), together with the average FPS from the five original objects. This resulted in random noise patterns that were matched to the objects in terms of luminance and spatial frequency. Linear pixel-by-pixel interpolation was performed by linear combination of each of the objects with each of four random noise patterns at seven stimulus levels, determined by pilot psychophysics. Isoluminance of all frames within $\pm10\%$ was confirmed using a Photometer. Average luminance of the entire 4° by 4° stimulus was 14 cd/m². Stimuli were displayed on a 17 inch computer monitor after appropriate gamma correction. Images were always presented at the center of gaze. Central fixation within a $\pm1.25°$ window was required at all times during a trial.

Novelty/Familiarity

During the novel object experiment, a new set of objects was presented to the animal each day. Thus, each set of objects presented in each session of the novel object experiment was completely unfamiliar to the monkeys. In preparation for the familiar object experiment, we allowed monkeys to perform the task with one set of objects for about five consecutive behavioral sessions. We found that there was a continuous improvement in performance over 5 days, but after that behavior was stable. Then, the (now familiar) set of objects was used for several recording sessions. Recording sessions with novel and familiar objects were conducted in an intermixed fashion, such that a session with familiar objects would often be followed by a session with novel objects and vice versa. The behavioral data summarized in Figure 2 were collected during these recording sessions. To counterbalance for novelty, we used novel noise patterns for every behavioral and recording session in the familiar object experiment and the same noise patterns for all sessions during the novel object experiment. Noise patterns had to be novel every day in the familiar object experiment to prevent monkeys from cheating by learning details about particular noise patterns. Familiar noise patterns were employed in the novel object experiment to equate the total amount of novelty between the two experiments.

Data Analysis

Data were analyzed using custom written MATLAB (MathWorks, Natick, MA) programs and the SPSS statistical software package. Background activity was assessed by averaging activity over the 400 ms of fixation prior to presentation of the sample object. Incidence of object selectivity at the 100% stimulus level was assessed with a one-way ANOVA (evaluated at $p < 0.05$) using object identity as a factor for the entire sample period and delay period (from 100 to 1650 ms after sample onset). In addition, we also performed a similar ANOVA (evaluated at $p < 0.05$) for the "prospective" period (from 1250 to 1650 ms after sample onset). To quantify object selectivity, we employed ROC analyses (Green and Swets, 1966; Tolhurst et al., 1983; Vogels and Orban, 1990). Essentially, a ROC measures the degree of overlap between two response distributions. Given two distributions of neural activity A and B, we start by plotting for each possible firing rate the proportion of distribution A that exceeded this criterion versus the proportion of distribution B that exceeded it. Calculation of the area under this ROC curve yields a single number for that comparison. This analysis has several advantageous properties. First, it provides an assumption-free estimate of the degree of overlap between A and B; values near 0.5 indicate large overlap between A and B, whereas values near 0 or 1 indicate small overlap. Second, it can be conveniently interpreted as the performance of an ideal observer in a two-way forced choice task. Third, it is independent of the firing rate of the neuron under study and can thus be used to compare the activity of neurons with widely different baseline and dynamic firing rates.

In the context of the present experiment, we computed ROC area values comparing each of the five objects to every other object (10 comparisons) at each stimulus level (7 comparisons), resulting in a total of 70 comparisons. Note that at the 0% stimulus level, we essentially compared one set of responses to noise patterns to another set of responses to the identical patterns. We performed these calculations using all trials (correct and incorrect) and correct trials only. Results were virtually identical, and data presented here are based on all trials because at 0% stimulus, half of the trials are arbitrarily assigned as incorrect, and excluding these from the analysis makes comparison across stimulus levels more difficult since they would be based on an unequal number of trials.

Because we were also interested in the time course of the object selectivity, we computed ROC area values at 32 different time points. We employed both a cumulative and a sliding bin method. For the sliding bin, we moved a bin of fixed size (200 ms) through the trial by shifting it each time by 50 ms. For the cumulative bin method, estimation started with a window of 50 ms duration and began 100 ms after sample presentation to adjust for the response latencies of PF neurons. Successive time points were obtained by adding the data from an additional 50 ms until the entire sample period and delay period (1650 ms) were considered together to obtain the last data point. For each neuron, we thus computed 2240 ROC area coefficients. Then, we determined the preferred and nonpreferred objects by selecting the comparison, which yielded the maximal ROC area at any time point (at 100% stimulus). We used the preferred and nonpreferred objects because we wanted to quantify how much object information it was communicating. To generate the surface plots of Figures 4–6, we plotted the ROC area values of this comparison—the object selectivity—as a function of time and stimulus level. The values in Figure 7 are based on the maxima (in time) of this optimal comparison for 100% and 65% stimulus levels (not necessarily the same time slice). Note that preferred and nonpreferred objects were those selected as yielding maximum object selectivity values for each neuron. Thus, overall each neuron responded more vigorously to the preferred than to the nonpreferred objects. However, for some time bins (especially for the sliding bin method), object selectivity could fall below the value of 0.5, indicating that during these intervals there were actually more spikes in response to the nonpreferred object.

Depth of Selectivity

The object selectivity defined above quantifies firing rate differences between the most and least preferred objects. In addition, we measured the depth of selectivity (S), which takes all five objects into account (Rainer et al., 1998). This analysis was performed at the 100% stimulus level for the entire sample period and delay period, treated together (from 100 to 1650 ms after sample onset). The depth of selectivity is defined as:

$$S = \frac{n - \frac{\sum R_i}{R_{max}}}{n - 1}$$

where n = number of objects, R_i = firing rate to ith object, and R_{max} = max{R_i}. S can vary from zero to 1, where a value of zero means identical responses to all objects, and a value of 1 means activation by one object and silence by all others.

References

Amit, D.J., Fusi, S., and Yakovlev, V. (1997). Paradigmatic working memory (attractor) cell in IT cortex. Neural Comput. 9, 1071–1092.

Asaad, W.F., Rainer, G., and Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. Neuron 21, 1399–1407.

Ball, K., and Sekuler, R. (1982). A specific and enduring improvement in visual motion discrimination. Science 218, 697–698.

Ball, K., and Sekuler, R. (1987). Direction-specific improvement in motion discrimination. Vision Res. 27, 953–965.

Barbas, H., and Mesulam, M.M. (1985). Cortical afferent input to the principalis region of the rhesus monkey. Neuroscience 15, 619–637.

Bichot, N.P., Schall, J.D., and Thompson, K.G. (1996). Visual feature selectivity in frontal eye fields induced by experience in mature macaques. Nature 381, 697–699.

Brady, N. (1997). Spatial scale interactions and image statistics. Perception 26, 1089–1100.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. J. Neurosci. 12, 4745–4765.

Buonomano, D.V., and Merzenich, M.M. (1998). Cortical plasticity: from synapses to maps. Annu. Rev. Neurosci. 21, 149–186.

Chen, L.L., and Wise, S.P. (1995). Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. J. Neurophysiol. 73, 1101–1121.

Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. Proc. Natl. Acad. Sci. USA 93, 13494–13499.

Dosher, B.A., and Lu, Z.L. (1999). Mechanisms of perceptual learning. Vision Res. 39, 3197–3221.

Fahle, M., Edelman, S., and Poggio, T. (1995). Fast perceptual learning in hyperacuity. Vision Res. 35, 3003–3013.

Fahy, F.L., Riches, I.P., and Brown, M.W. (1993). Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. Exp. Brain Res. 96, 457–472.

Field, D.J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. J.Opt. Soc. Am. A 4, 2379–2394.

Furmanski, C.S., and Engel, S.A. (2000). Perceptual learning in object recognition: object specificity and size invariance. Vision Res. 40, 473–484.

Fuster, J.M. (1997). The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe (Lippincott Williams & Wilkins Publishers).

Gilbert, C.D. (1996). Plasticity in visual perception and physiology. Curr. Opin. Neurobiol. 6, 269–274.

Gold, J., Bennet, P.J., and Sekuler, A.B. (1999). Signal but not noise changes with perceptual learning. Nature 402, 176–178.

Goldman-Rakic, P.S. (1995). Cellular basis of working memory. Neuron 14, 477–485.

Goldstone, R.L. (1998). Perceptual learning. Annu. Rev. Psychol. 49, 585–612.

Green, D.M., and Swets, J.A. (1966). Signal Detection Theory and Psychophysics (New York: J. Wiley and Sons).

Gross, C.G. (1973). Visual functions of inferotemporal cortex. In Handbook of Sensory Physiology, Volume Seven, R. Jung, ed. (Berlin: Springer-Verlag), pp. 451–482.

Ito, M., and Gilbert, C.D. (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. Neuron 22, 593–604.

Ito, M., Westheimer, G., and Gilbert, C.D. (1998). Attention and perceptual learning modulate contextual influences on visual perception. Neuron 20, 1191–1197.

Karni, A., and Sagi, D. (1993). The time course of learning a visual skill. Nature 365, 250–252.

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. 2, 176–185.

Knight, R.T., Grabowecky, M.F., and Scabini, D. (1995). Role of human prefrontal cortex in attention control. Adv. Neurol. 66, 21–26.

Kobatake, E., Wang, G., and Tanaka, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. J. Neurophysiol. 80, 324–330.

Li, L., Miller, E.K., and Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. J. Neurophysiol. 69, 1918–1929.

Logothetis, N.K. (1998). Object vision and visual awareness. Curr. Opin. Neurobiol. 8, 536–544.

Logothetis, N.K., Pauls, J., and Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. Curr. Biol. 5, 552–563.

Miller, E.K. (1999). The prefrontal cortex: complex neural properties for complex behavior. Neuron 22, 15–17.

Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J. Neurosci. 16, 5154–5167.

Miyashita, Y., and Hayashi, T. (2000). Neural representation of visual objects: encoding and top-down activation. Curr. Opin. Neurobiol. 10, 187–194.

Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. Nature 341, 52–54.

Pandya, D.N., and Yeterian, E.H. (1990). Prefrontal cortex in relation to other cortical areas in rhesus monkey: architecture and connections. Prog. Brain Res. 85, 63–94.

Parker, A., Wilding, E., and Akerman, C. (1998). The Von Restorff effect in visual object recognition memory in humans and monkeys. The role of frontal/perirhinal interaction. J. Cogn. Neurosci. 10, 691–703.

Petrides, M. (1994). Frontal lobes and working memory: evidence from investigations of the effects of cortical excisions in nonhuman primates. In Handbook of Psychology, F. Boller and J. Grafman, eds. (New York: Elsevier), pp. 59–82.

Petrides, M., and Pandya, D.N. (1999). Dorsolateral prefrontal cortex: comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. Eur. J. Neurosci. 11, 1011–1036.

Quintana, J., and Fuster, J.M. (1992). Mnemonic and predictive functions of cortical neurons in a memory task. Neuroreport 3, 721–724.

Rainer, G., Asaad, W.F., and Miller, E.K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. Nature 393, 577–579.

Rainer, G., Rao, S.C., and Miller, E.K. (1999). Prospective coding for objects in primate prefrontal cortex. J. Neurosci. 19, 5493–5505.

Ranganath, C., and Paller, K.A. (1999). Frontal brain potentials during recognition are modulated by requirements to retrieve perceptual detail. Neuron 22, 605–613.

Rolls, E.T., Baylis, G.C., Hasselmo, M.E., and Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the

cortex in the superior temporal sulcus of the monkey. Exp. Brain Res. *76*, 153–164.

Romo, R., Brody, C.D., Hernandez, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470–473.

Sagi, D., and Tanne, D. (1994). Perceptual learning: learning to see. Curr. Opin. Neurobiol. *4*, 195–199.

Sakai, K., and Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. Nature *354*, 152–155.

Schall, J.D., and Thompson, K.G. (1999). Neural selection and control of visually guided eye movements. Annu. Rev. Neurosci. *22*, 241–259.

Sigman, M., and Gilbert, C.D. (2000). Learning to find a shape. Nat. Neurosci. *3*, 264–269.

Tadmor, Y., and Tolhurst, D.J. (1993). Both the phase and the amplitude spectrum may determine the appearance of natural images. Vision Res. *33*, 141–145.

Tanaka, K. (1996). Inferotemporal cortex and object vision. Annu. Rev. Neurosci. *19*, 109–139.

Tolhurst, D.J., Movshon, J.A., and Dean, A.F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. Vision Res. *23*, 775–785.

Vogels, R., and Orban, G.A. (1990). How well do response changes of striate neurons signal differences in orientation: a study in the discriminating monkey. J. Neurosci. *10*, 3543–3558.

Vogels, R., and Orban, G.A. (1994). Does practice in orientation discrimination lead to changes in the response properties of macaque inferior temporal neurons? Eur. J. Neurosci. *6*, 1680–1690.

Wallis, G., and Bulthoff, H. (1999). Learning to recognize objects. Trends Cogn. Sci. *3*, 22–31.

Weinberger, N.M. (1995). Dynamic regulation of receptive fields and maps in the adult sensory cortex. Annu. Rev. Neurosci. *18*, 129–158.

Wiggs, C.L., and Martin, A. (1998). Properties and mechanisms of perceptual priming. Curr. Opin. Neurobiol. *8*, 227–233.

Zohary, E., Celebrini, S., Britten, K.H., and Newsome, W.T. (1994). Neuronal plasticity that underlies improvement in perceptual performance. Science *263*, 1289–1292.