

Visual Categorization and the Primate Prefrontal Cortex: Neurophysiology and Behavior

DAVID J. FREEDMAN,^{1,2,5} MAXIMILIAN RIESENHUBER,^{3–5} TOMASO POGGIO,^{3–5} AND EARL K. MILLER^{1,2,5}
¹Center for Learning and Memory, ²The Institute of Physical and Chemical Research–Massachusetts Institute of Technology Neuroscience Research Center, ³Center for Biological and Computational Learning and ⁴McGovern Institute for Brain Research, ⁵Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Received 18 January 2002; accepted in final form 26 March 2002

Freedman, David J., Maximilian Riesenhuber, Tomaso Poggio, and Earl K. Miller. Visual categorization and the primate prefrontal cortex: neurophysiology and behavior. *J Neurophysiol* 88: 929–941, 2002; 10.1152/jn.00040.2002. The ability to group stimuli into meaningful categories is a fundamental cognitive process. To explore its neuronal basis, we trained monkeys to categorize computer-generated stimuli as “cats” and “dogs.” A morphing system was used to systematically vary stimulus shape and precisely define a category boundary. Psychophysical testing and analysis of eye movements suggest that the monkeys categorized the stimuli by attending to multiple stimulus features. Neuronal activity in the lateral prefrontal cortex reflected the category of visual stimuli and changed with learning when a monkey was retrained with the same stimuli assigned to new categories. Further, many neurons showed activity that appeared to reflect the monkey’s decision about whether two stimuli were from the same category or not. These results suggest that the lateral prefrontal cortex is an important part of the neuronal circuitry underlying category learning and category-based behaviors.

INTRODUCTION

Our perception of the environment is not a faithful registration of its physical attributes. Instead, we carve the world into meaningful groupings or categories. This process of abstracting and storing the commonalities among like-themed individuals is fundamental to cognitive processing because it imparts knowledge. For example, knowing that a new gadget is a “camera” instantly and effortlessly provides a great deal of information about its relevant parts and functions and spares us from having to learn anew each time we encounter a new individual. The ability to categorize stimuli is a cornerstone of complex behavior. Categories are evident in all sensory modalities and range from relatively simple (e.g., color perception) to the most abstract human concepts.

Because perceptual categories often group together very different-looking things, their representation must involve something beyond the sort of neuronal tuning that typifies encoding of physical appearance: gradual changes in neuronal activity as features gradually change (e.g., shape, orientation, direction). In fact, evidence that a human or animal has stored a category is that behavior does not track smoothly with

changes in physical appearance: categories have sharp boundaries (not gradual transitions) between them and members of the same category are treated as equivalent even though their physical appearances may vary widely. A simple example is crickets sharply dividing (at 16 kHz) a continuum of pure tones into “mate” versus “bat” (a predator) (Wytttenbach et al. 1996). Other examples include humans’ perception of the phonemes “b” versus “p” (Lieberman et al. 1967) and the facial expressions of emotion (Beale and Keil 1995).

The elaborate behavioral repertoire of advanced animals naturally depends on more elaborate categorization abilities. Their mental lexicon includes categories that are characterized along multiple dimensions and are often difficult to precisely define. In addition, advanced animals have an enormous capacity to learn and adapt. Monkeys, for example, have been taught categories such as animal versus nonanimal (Roberts and Mazmanian 1988), food versus nonfood (Fabre-Thorpe et al. 1998), tree versus nontree, fish versus nonfish (Vogels 1999), and ordinal numbers (Orlov et al. 2000). Such categories could be processed in brain areas involved in object recognition such as the inferior temporal cortex (ITC) (Desimone et al. 1984; Gross 1973; Logothetis and Sheinberg 1996; Tanaka 1996) as well as those involved in orchestrating voluntary, visually guided behaviors, such as the prefrontal cortex (PFC) (Fuster 1997; Goldman-Rakic 1987; Miller 2000; Miller and Cohen 2001). The PFC and ITC are directly connected (Ungerleider et al. 1989; Webster et al. 1994) and both contain neurons that often exhibit highly specific responses to complex stimuli such as trees, fishes, faces, brushes, etc. (Bruce et al. 1981; Desimone et al. 1984; Gross et al. 1972; Miller et al. 1996; Perrett et al. 1982; Scialidhe et al. 1999; Tanaka et al. 1991) and are influenced by experience (Booth and Rolls 1998; Kobatake et al. 1998; Logothetis et al. 1995; Miyashita 1988; Rainer and Miller, 2000). Whether or not their activity reflects stimulus categories has not been clear. These neurons have not been tested for the diagnostic characteristics of categories, (e.g., sharp boundaries and within-category generalization); their specificity might reflect similarities and differences in physical appearance of the stimuli, not necessarily their category membership.

Address for reprint requests: E. K. Miller, Bldg. E25, Room 236, Massachusetts Institute of Technology, Cambridge, MA 02139 (E-mail: ekm@ai.mit.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

To evaluate the role of the PFC in visual categorization, we trained monkeys to categorize computer-generated stimuli into two categories, cats and dogs. A novel three-dimensional (3-D) morphing system was used to create a large set of parametric blends of six prototype images (3 species of cats and 3 breeds of dogs) (Beymer and Poggio 1996; Shelton 2000). This allowed us to establish a sharp category boundary between stimuli that were physically similar yet include in the same category stimuli that were visually dissimilar. A brief report of these results appeared previously (Freedman et al. 2001).

METHODS

Subjects

Two female adult rhesus monkeys (*Macacca mulatta*) weighing 6.0 and 7.5 kg were used in this study. Using previously described methods (Miller et al. 1993), they were implanted with a head bolt to immobilize the head during recording and with recording chambers. Eye movements were monitored and stored using an infrared eye-tracking system (Iscan, Cambridge, MA). All surgeries were performed under sterile conditions while the animals were anesthetized with isoflurane. The animals received postoperative antibiotics and analgesics and were handled in accord with National Institutes of Health guidelines and the recommendations of the Massachusetts Institute of Technology Animal Care and Use Committee.

Recording techniques

Electrode penetration sites were determined using magnetic resonance imaging scans obtained prior to surgery. The recording chambers were positioned stereotactically over the lateral prefrontal cortex such that the principal sulcus and ventrolateral prefrontal cortex were readily accessible. Neuronal activity was isolated using arrays of four to eight independently moveable tungsten microelectrodes (FHC Instruments, Bowdoinham, ME). The electrodes were advanced using custom-made screw-driven mini-microdrives (Nichols et al. 1998) mounted on a plastic grid (Crist Instruments, Damascus, MD). Neuronal activity was amplified, filtered, and stored for off-line sorting into individual neuron records (Plexon Systems, Dallas, TX). This allowed us to isolate an average of nearly two neurons per electrode. We did not prescreen neurons for task-related activity such as visual responsiveness or stimulus selectivity. Rather, we randomly selected neurons for study by advancing 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 prefrontal activity.

Stimuli

A large continuous set of images was generated from three prototype cats and three prototype dogs (Fig. 1) with a novel algorithm (Shelton 2000). It found corresponding points between one of the prototypes and the others and then computed their differences as vectors. Morphs were created by linear combinations of these vectors added to that prototype. For more information see <http://www.ai.mit.edu/people/cshelton/corr/>. By morphing different amounts of the prototypes, we could generate thousands of unique images, continuously vary shape, and precisely define one or more arbitrary category boundaries. For most of the experiments, the images were divided into two groups, cats and dogs, with the boundary at an equal blend of cat and dog. Thus category membership was defined by whichever category contributed more (>50%) to a given morph. As a result, stimuli that were close to but on opposite sides of the boundary were visually similar, while stimuli that belonged to the same category could be visually dissimilar [e.g., the "housecat" (C1) and "cheetah" (C2)

prototypes]. The stimuli differed along multiple features and were smoothly morphed, i.e., without sudden appearance or disappearance of any feature. They were 4.2° in diameter, had identical color, shading, orientation and scale, and were presented at the center of gaze.

We confirmed that the morphs did indeed vary smoothly by using an image correlation analysis. This analysis was used merely to ensure that the morphing system functioned as designed and generated stimuli that had no a priori discontinuities that the monkeys could exploit to solve the task. A two-dimensional (2-D) correlation coefficient was calculated for neighboring images at six levels of blends of cat and dog (cat:dog: 100:0, 80:20, 60:40, 40:60, 20:80, 0:100) along each of the nine between-category morph lines. The correlation was calculated by computing the 2-D correlation coefficient separately for each color plane and then averaging across planes. The correlation coefficient between neighbors remained constant and high (~0.9) across the entire morph space. The coefficients between stimuli directly across the cat/dog boundary did not differ from the coefficients calculated between adjacent morphs within the same category (1-way ANOVA, $P = 0.44$).

Behavioral tasks

The monkeys performed a delayed match-to-category task that required them to judge whether two successive stimuli were from the same category (Fig. 2). The trial began when the monkey grasped a metal bar and fixated a small (0.3°) white spot at the center of a CRT screen. They were required to maintain gaze within a $\pm 2^\circ$ square window around the fixation spot for the entire trial. After the initial 500 ms of fixation, a sample image was presented at the center of the screen for 600 ms, followed by a 1,000-ms delay. Then a choice image appeared. If the sample and choice stimuli were from the same category (a category match), the monkeys were required to release the lever before the stimulus disappeared 600 ms after its onset to receive a juice reward. If the choice image was from a different category (a category nonmatch), there was an additional brief delay (600 ms) followed by another image that was always a match and thus required a lever release. As a result, a category judgment was only required for the first choice image. The second delay and match image were used so that a behavioral response would be required on every trial. This ensured that the monkeys were always paying attention. Because a decision was only required for the first choice image and the forthcoming behavioral response was predictable from the second delay onwards, that delay and subsequent match image will not be considered further. Note that with this design, the behavioral response (lever release) is not uniquely associated with a category (it was used to signal "match," not cat or dog) and, further, the monkeys could not predict whether the first choice stimulus would require a response. Thus any differential activity to the sample categories could not be related to the behavioral response. A 2,000- to 3,000-ms inter-trial interval followed correct trials. An error was defined as a lever release to a nonmatch or failure to release to a match; breaks of fixation were not counted among the error rates in behavioral analyses. An additional 3,000-ms "time-out" was added to the inter-trial interval following an error. Monkeys typically performed >700 correct trials per day.

The monkeys were gradually trained to categorize the images as cats and dogs by beginning with a delayed match-to-sample task in which the prototypes were used as samples, the match was always identical to it, and the nonmatches were a prototype from the other category. We then gradually included more and more morphs as samples and chose images at increasing distances from the prototypes. In parallel, matches were chosen from an increasingly greater distance of morph space around the sample while respecting the category boundary. Nonmatches were always from the other category.

During the course of training, >1,000 sample stimuli were used from all over the morph space. This prevented monkeys from solving

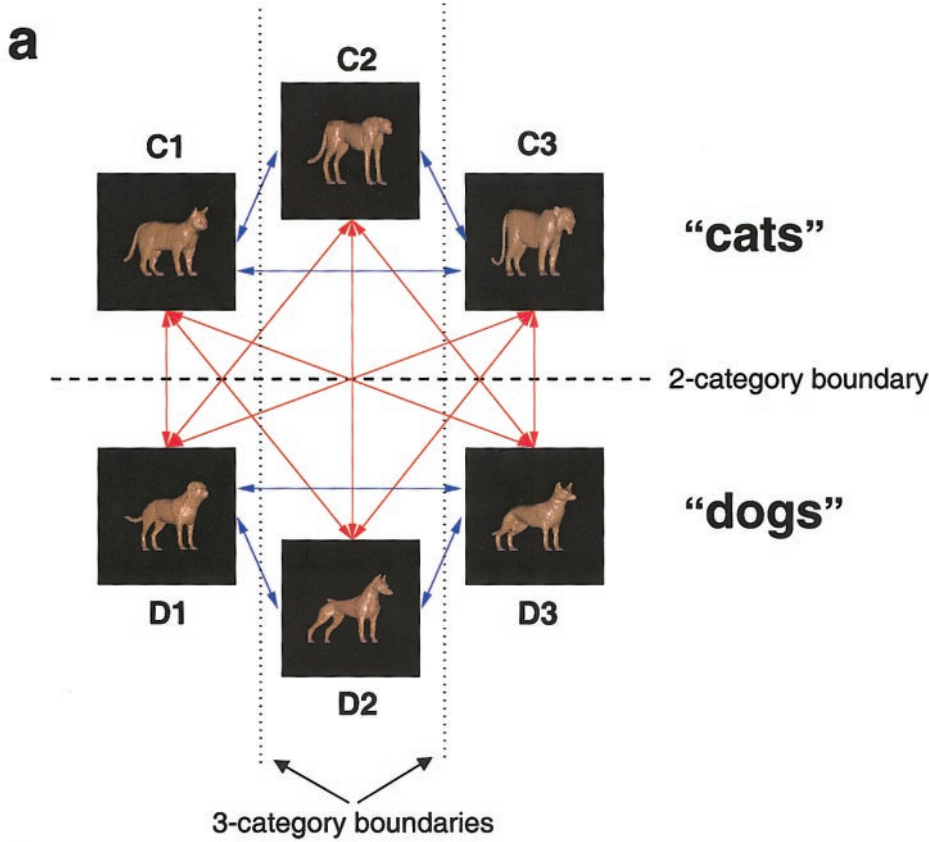



FIG. 1. Organization of stimulus set. A: the 6 prototype images and 15 morph lines. The sample stimulus set was composed of 54 unique images: 6 prototypes (as shown), 4 images evenly placed (20, 40, 60, 80%) along the 9 lines (in red) connecting each cat to each dog prototype, and 2 images (at 40 and 60%) along each of the 6 lines (in blue) between prototypes of the same category (with respect to the 2-class boundary). B: an example of the morphs generated between the C1 and D1 prototypes.



the task by simply remembering specific stimulus-response contingencies. Neurophysiological recording, however, requires that a limited number of stimuli be used so that each can be repeated multiple

times and neuronal variability can be assessed. Thus for recording experiments, we limited the samples to 54 images. This included the six prototype images and morphs from equally spaced intervals across

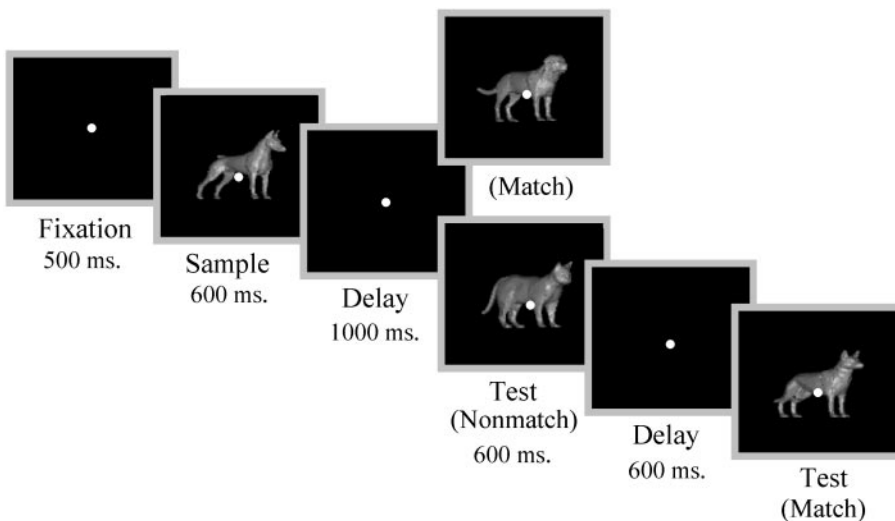


FIG. 2. Task design and behavioral performance. The trial began with central fixation (500 ms) after which a sample stimulus appeared at the center of gaze for 600 ms. This was followed by a 1-s delay and then by a choice (Test) stimulus (600 ms). If the category of the choice matched that of the sample, monkeys had to release a lever to the choice stimulus within 600 ms of its presentation to obtain a fruit-juice reward. If the choice was a nonmatch, there was another delay interval (600 ms) followed by a presentation of a match, which required a lever release for a reward. There were an equal number of match and nonmatch trials and they were randomly interleaved.

each of the nine morph lines that connected each cat prototype to each dog prototype (Fig. 1A). There were six levels of blends of cat and dog (cat:dog, 100:0, 80:20, 60:40, 40:60, 20:80, 0:100) along the nine morph lines that crossed the two-category boundary (the red lines in Fig. 1A) and two levels along the six within-category morph lines (60:40, 40:60; the blue lines in Fig. 1A). To prevent monkeys from learning to memorize specific stimulus-response contingencies during the recording experiments, the choice stimuli were 100 randomly generated morphs from each category that were randomly paired with sample stimuli of the appropriate category. To ensure that category judgment errors were due to confusion over the sample category, the choice stimuli unambiguously belonged to a given category: they were always chosen to be at a distance of $\geq 20\%$ from the boundary.

The monkeys' categorization abilities were further examined with separate psychophysical tests employing an additional 14 morphs that were equally and tightly spaced (6.67% intra-stimulus distance) along each of the morph lines that crossed a category boundary. This allowed for a more precise description of the monkeys' ability to categorize stimuli near the category boundary. This task was identical in all timing and behavioral events except that the monkeys were randomly rewarded on trials in which the sample stimulus was very close to the category boundary (<10% difference). As monkeys were not shown morphs of closer than 10% distance from the category boundary during training, feedback during psychophysical testing was withheld on trials where such stimuli were presented as samples to discourage learning and changes in performance during those sessions.

To test the effects of learning on neuronal activity, we trained a monkey to re-categorize the cat and dog images into three new categories. Two new category boundaries were defined that were orthogonal to the original two-category boundary (Fig. 1A). This resulted in three new classes that each contained morphs centered around one cat prototype and one dog prototype. The same 54 sample stimuli were used for neurophysiological recording under the two and three-category schemes. As in the two-category experiment, the choice stimulus set consisted of 100 randomly generated morphs from each category that had a maximum component of 20% from each of the other two categories.

Data analysis

Neuronal activity level was calculated in four time epochs: baseline, sample presentation, first delay, and first choice stimulus presentation. Baseline neuronal activity was averaged over the 500 ms of fixation preceding sample presentation. Sample period activity was averaged over an 800-ms epoch beginning 100 ms after sample onset to account for the latency of PFC neuronal responses and included the first 300 ms following sample offset to include any activity related to that event. Delay activity was assessed over an 800-ms epoch beginning 300 ms after sample offset and ending 100 ms after first choice stimulus onset. Activity to that choice stimulus was averaged over an epoch that began 100 ms after its onset and ended 2 SD before the monkeys' average reaction time during each recording session to exclude any effects related to the execution of the behavioral response.

Category information in neuronal activity was assessed using several methods. We computed an index of category tuning by calculating each neuron's average firing rate difference to pairs of sample morphs from the same category (within-category difference, WCD) and its average difference to samples from different categories (between-category difference, BCD) using images from the morph lines that crossed the category boundary. The WCD was defined by computing the absolute difference between the 100 and 80% morphs and between the 80 and 60% morphs for both categories and averaging these values. The BCD was computed by averaging the across-boundary differences between the 60% cats and 60% dogs. As a result, the distance between stimuli in morph space was identical

(20%) for the BCD and WCD comparisons. A standard index was computed for each neuron by dividing the difference between their BCD and WCD values by their sum. This index can have values ranging from -1 to 1 . Positive values indicate a larger difference between categories, whereas negative values reflect larger differences within the categories than between categories. BCD and WCD values were computed for neurons recorded during the three-category task in a similar fashion by determining differences in activity to samples that differed by 20% along the morph lines that crossed the three-category boundaries (Fig. 1). To ensure that the previously learned two-category scheme did not contribute to the values obtained when calculating category effects in the three category scheme, we excluded from this analysis the morph lines that crossed both the two- and three-category boundaries (e.g., the morph line connecting cat prototype 1 and dog prototype 2).

In addition to computing an index, we also compared between and within category differences by using a receiver-operating characteristics (ROC) analysis (Green and Swets 1966; Tolhurst et al. 1983; Vogels and Orban 1990). The ROC analysis measures the degree of overlap between two distributions of values. It has several advantages. First, it makes no assumptions about the two distributions, *A* and *B* and thus returns an unbiased estimation of overlap. Second, it can be interpreted as the performance of an ideal observer in a two-way forced choice task; values of 0.5 indicate 50% correct classification (guessing) while values of 0 or 1 indicate error-free classification. Third, it is independent of neuronal firing rate and number of observations. While the category index described above explicitly tests for sharp tuning across the category boundary, the ROC value gives a general measure of the degree of category selectivity.

To determine the time course of category information in neuronal activity, we computed the ROC area within a time bin of 200 ms that was slid in 10-ms steps. We began 500 ms prior to sample stimulus onset and ended 100 ms following the first choice stimulus onset. This was computed for all neurons that were "category selective" (according to a 2-tailed *t*-test comparing the average response to cats and dogs, evaluated at $P < 0.01$) during the sample and/or delay epochs.

The latency for neuronal activation (irrespective of category information) was determined by compiling the average histogram of firing rate values for all responsive neurons (i.e., neurons that showed significantly different activity during the sample and/or delay periods compared with baseline, evaluated by 2-tailed *t*-test at $P < 0.01$.) This average histogram was smoothed with a 30-ms Gaussian window, and the latency was defined as the point of maximum inflection (determined by computing the 2nd derivative at all points along the histogram) of this curve following sample onset.

Because neurons have a wide range of firing rates, these firing rates were normalized when computing histograms of average effect size across the population. For each neuron, the mean firing rates at each of the six steps from the cat to dog prototypes were computed. Then, the range of firing rates for each neuron were rescaled according to the minimum and maximum values across those six groups such that each neuron's minimum and maximum rate was 0.0 and 1.0, respectively. This allowed each neuron's range of modulation to contribute equally to the population average. Similar results were obtained by conducting these analyses using raw firing rates. Category information was not limited to a specific range of activity; as our single neuron examples will illustrate, it was evident in neurons exhibiting both low and high firing rates.

RESULTS

Behavioral data

CATEGORY JUDGMENTS. The monkeys were very accurate at the two-category judgments. During the recording sessions, performance was high ($\sim 90\%$ correct), even when the samples were close to the category boundary; the monkeys classified

dog-like cats (60:40 cat:dog) correctly ~90% of the time, and misclassified them as dogs only 10% of the time and vice-versa (Fig. 3). The results of psychophysical tests with more closely spaced morphs are shown in Fig. 4A. Even very near the boundary, when stimuli were very similar to (only 3% different from) the other category (i.e., a 53.3% cat or dog), the monkeys still performed significantly above chance (~65%, chance = 50%). Thus even with closely spaced morph images, the sudden change in behavior characteristic of category representations were evident in behavior.

Figure 4B shows the performance of *monkey A* after it had been trained to re-categorize the same images under the three-category scheme. Performance here was somewhat lower than during the two-category task. This is presumably because there were two boundaries in the three-category task and thus a higher percentage of stimuli were close to the border (the data in the figure are collapsed across the boundaries). Still, the sharp drop-off in performance indicative of a category representation was evident; monkeys continued to perform above chance for morphs that were only ~3% different from the boundary. The greater difficulty of the three-category task was also evident in the monkey's behavioral reaction times. They were significantly longer during the three-category task (average = 307 ms) than the two-category task (264 ms, t -test at $P < 0.01$).

STIMULUS FEATURES USED FOR CATEGORIZATION. To explore which features monkeys tended to focus on when categorizing the images into cats and dogs, we conducted further psychophysical testing. In one set of experiments, we removed the requirement to maintain central fixation (by removing the fixation point) and allowed the monkeys to freely gaze at the images. Given the close link between attention and gaze during unconstrained viewing, the assumption was that monkeys would spend more time gazing at the features that they were using to define the categories.

It seemed that the monkeys were not focusing on a single feature to categorize the images. Even though the sample presentation was brief (600 ms), they typically made several

saccades while viewing the stimulus. One monkey made an average of 3.45 saccades and the other monkey averaged 2.25 saccades during sample presentation (defined as the number eye movements exceeding $50^\circ/\text{s}$, equivalent to 0.5° of movement in adjacent 10-ms time bins). Interestingly, the two monkeys seemed to use different combinations of features to categorize the images. One monkey tended to look toward the tail of the sample images; its gaze was on average 1.46° to the left and 0.60° below the center of the screen. The other monkey tended to direct its gaze toward the head region; on average, its gaze was 0.57° to the right and 1.16° above central fixation. Figure 5A shows representative traces from one trial for each monkey. The gaze patterns for the two monkeys were significantly different from one another (along both the horizontal and vertical axes, t -tests, $P < 0.01$).

We also tested the monkeys' ability to categorize the images after removing the "heads" or "tails" of the morph stimuli and then interleaving them with nondegraded samples. The assumption here was that if the monkey was relying on a single feature on the head or tail of the image, its removal should cause a decrease in categorization performance to chance. This was not the case. As shown in Fig. 5B, we found that the monkeys' performance remained high (~80% correct) when either the head or tail was absent. This pattern of results suggests that each monkey attended to a unique combination of features and that neither of the monkeys used a single stimulus feature to categorize the images.

Neuronal data

BASIC PROPERTIES. A total of 395 lateral prefrontal cortex neurons were recorded from three hemispheres of two monkeys during performance of the two-category DMC task (130 from *monkey A*, 265 from *monkey B*, Fig. 6). Visual responsiveness was evaluated by comparing the activity in the sample and delay intervals to baseline activity using two-tailed t -tests (evaluated at $P < 0.01$). Based on this criterion, 259/395 (66%, 113 from *monkey A*, 146 from *monkey B*) of neurons were activated during one or more task intervals. The onset of

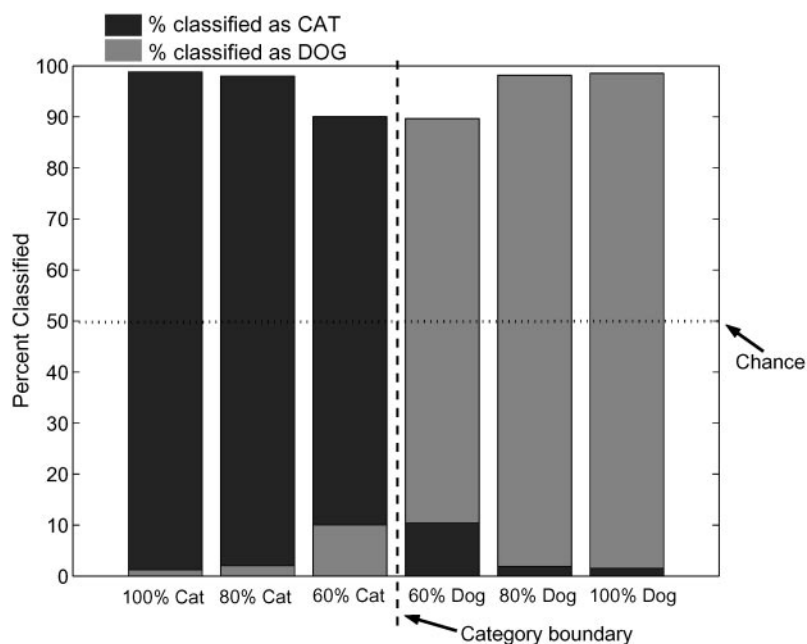


FIG. 3. Average performance of both monkeys during neurophysiological recordings for the 2-category task. Dark grey bars indicate the percent of samples classified as cat and light grey bars the percent classified as dog.

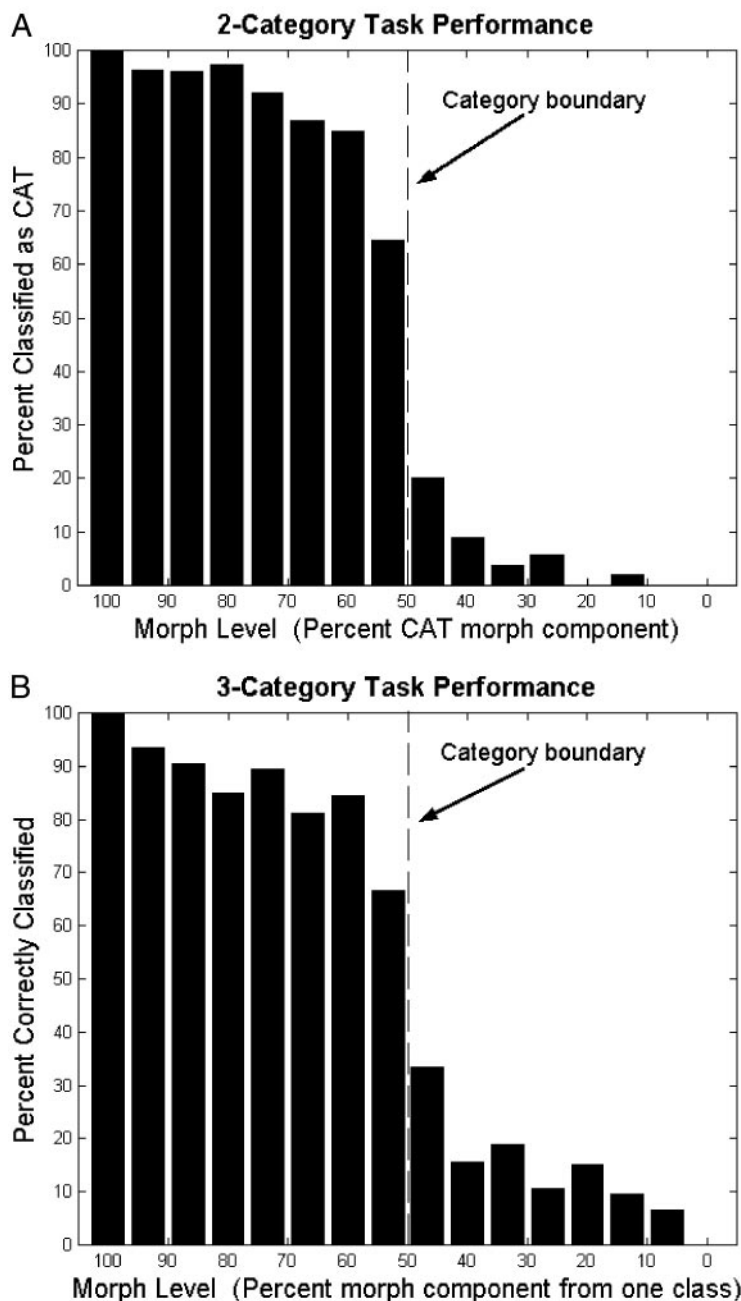


FIG. 4. Psychophysical performance for the 2- and 3-category tasks. For both the 2-category (A) and 3-category (B) tasks, the monkeys' error rates did not increase linearly as stimuli approached the category boundary but changed more sharply at the category boundary.

neuronal responsiveness across the population of responsive PFC neurons occurred at ~100 ms following sample stimulus onset (METHODS).

An initial assessment of neuronal category selectivity was made with a *t*-test of the activity to all cat stimuli versus all dog stimuli (evaluated at $P < 0.01$). This revealed that nearly a quarter of all (randomly selected) neurons (96/395, 24%, 60 and 36 in *monkeys A* and *B*, respectively) showed a significant difference in their overall activity to cats versus dogs in the sample and/or delay intervals (74 sample, 51 delay). Many neurons (78/395, or 20%; 67 sample, 32 delay) also showed a significant effect of the individual samples (i.e., were stimulus-selective) according to a one-way ANOVA (with the 54 sample stimuli as the factor; evaluated at $P < 0.01$). A majority of these stimulus-selective neurons also showed an overall effect

of category (56/78, or 72%; 46 sample, 21 delay, *t*-test, $P < 0.01$). Similar numbers of category selective neurons preferred cats (39/74 sample, 27/51 delay) as dogs (35/74 sample, 24/51 delay). In both monkeys, there was a greater incidence of category-selective neurons in the sample than the delay interval (*monkey A*: 48 sample, 31 delay; *monkey B*: 26 sample, 20 delay).

The activity of many neurons showed a sharp differentiation between the two categories that mirrored the monkeys' behavior. That is, they showed relatively large differences in activity to samples from different categories and relatively similar activity to samples from the same category. Two examples of single neurons are shown in Figs. 7, A and B. They show each neuron's average activity to all samples at different blends of cats and dogs. Both seem to encode the category of stimuli. Note that their activity was significantly different to dog-like

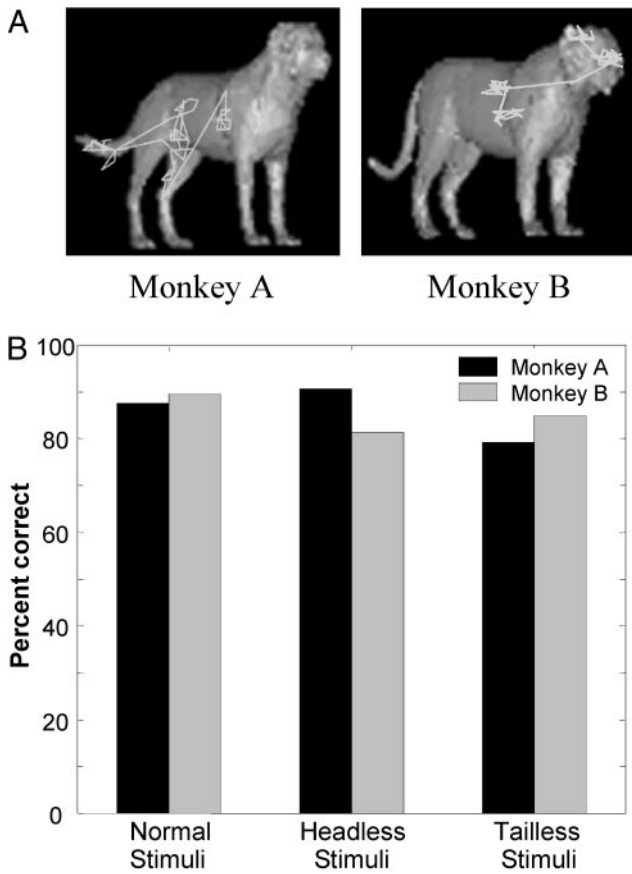


FIG. 5. Eye movements and degraded stimuli during behavioral testing. *A*: an example of monkey *A*'s (left) and monkey *B*'s (right) eye movements during the sample period of a single trial superimposed on the sample stimulus shown on that trial. Monkeys were allowed to gaze freely at the stimuli during behavioral testing as fixation was not required. During neurophysiological recordings, monkeys were always required to maintain fixation within a $\pm 2^\circ$ window. *B*: behavioral performance to degraded stimuli.

(60%) cats and cat-like (60%) dogs (t -test, $P < 0.001$), but there was no difference in activity between these stimuli and their respective prototypes ($P > 0.1$).

These effects were also evident in the average activity across the population of all stimulus-selective neurons. For this analysis, we chose neurons that were stimulus selective not category selective per se (ANOVA with the individual samples as a factor, $P < 0.01$, $n = 55$ for the sample interval, 29 for the

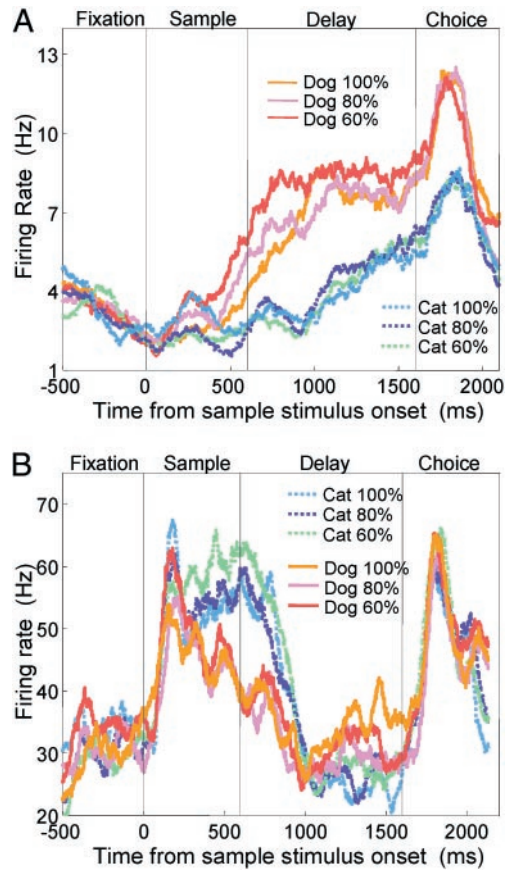


FIG. 7. Single neuron examples (2-category task). *A*: the average activity of a single neuron that showed greater activity to dogs during the memory delay. The histogram traces represent the neuron's average activity to stimuli at each of the 6 morph levels. *B*: the average activity of a single neuron that showed greater activity to cats during the late sample and early delay period.

delay, excluding neurons with mean firing rates < 2 Hz, as they can produce spurious results when normalized). Figure 8 shows the mean normalized firing rates for the six levels of morphs. Each neuron's preferred category was determined by the category of the single sample stimulus (of 54) that evoked the maximal firing rate, computed separately for the sample and delay intervals. By determining the preferred category by a single stimulus instead of the average across all category members, we ensured that this test was not biased toward finding a category effect. During both time epochs, there was

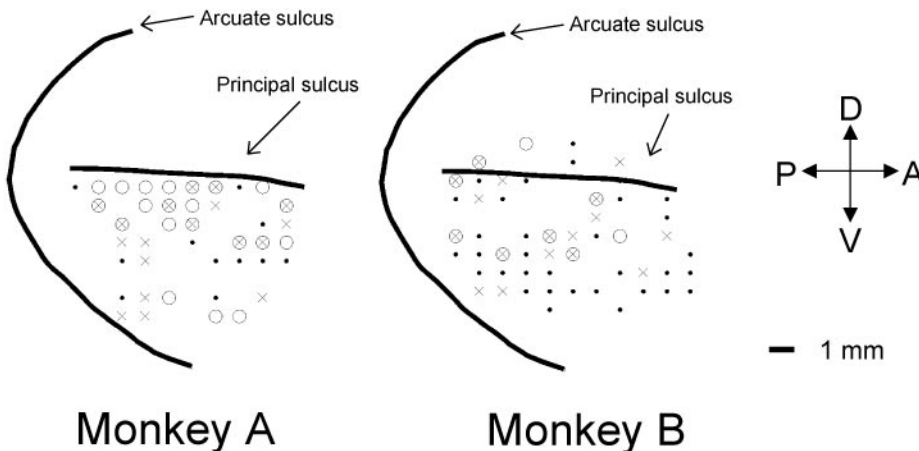


FIG. 6. Anatomical location of recording sites and category selective neurons in both monkeys. A, anterior; P, posterior; D, dorsal; V, ventral. X and O, recording sites at which cat- or dog-selective neurons were found, respectively. ●, each location at which no category selective neurons were encountered. There was no obvious topography to task-related neurons.

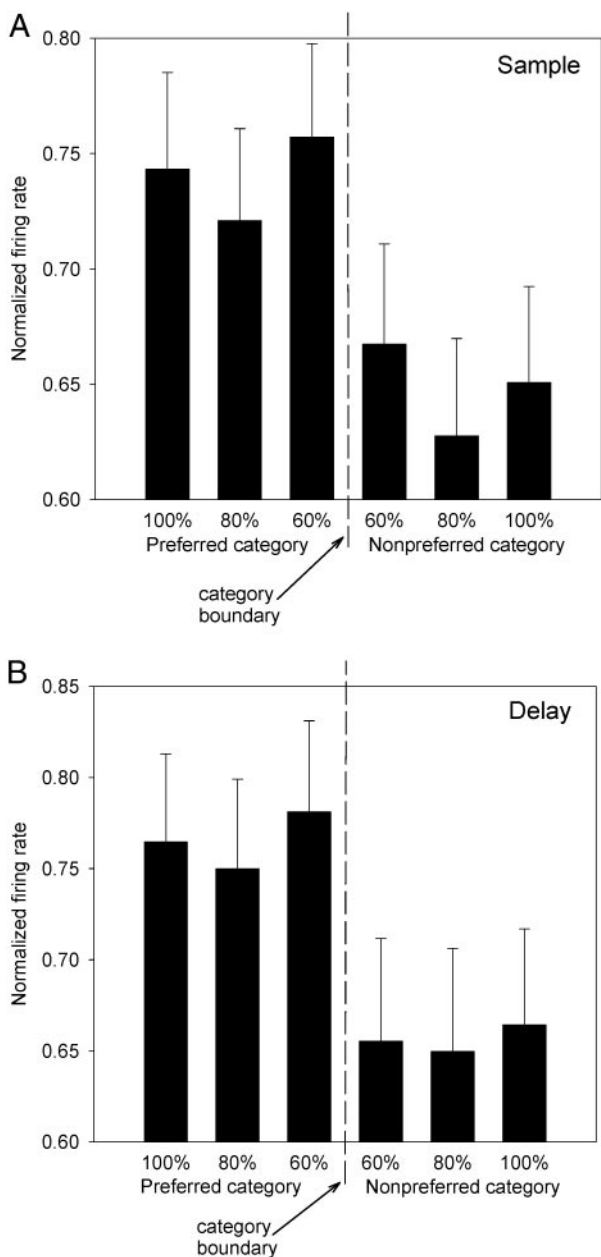


FIG. 8. Average neuronal response to preferred and nonpreferred categories during the sample (A) and delay (B). Each bar represents the population's average normalized response to stimuli at each of the 6 morph levels. The error bars represent SE.

a significant difference between the categories ($P < 0.01$) but no differences between the different morph levels within each category ($P > 0.6$, 2-way ANOVA with category and distance from the category boundary as factors).

QUANTIFICATION OF CATEGORY EFFECT. To quantify the effect of category membership on the neuronal population, we computed a category index that reflected each neuron's average difference in activity to samples across the category boundary versus its difference to samples that were from the same category (see METHODS). Positive values indicate greater differences across the category boundary than within each category and negative index values indicate the opposite.

We examined all stimulus-selective neurons irrespective of

whether they were category selective per se ($n = 78$, 67 sample, 32 delay). The distributions of category index values for the sample and delay periods are shown in Fig. 9. During both epochs, mean category index values were significantly greater than zero, i.e., the distribution was shifted toward category tuning (sample: 0.09, delay: 0.16, 1-tailed t -test, $P < 0.001$). For the subset of stimulus selective neurons that were category selective ($n = 46$ sample, 21 delay), the category indices were significantly greater (more shifted toward category tuning) during the delay than the sample interval (index = 0.12 sample, 0.21 delay, 2-tailed t -test, $P = 0.04$). Similar comparisons were also made by computing ROC values, which reflect how well an ideal observer would do at categorization using each neuron's firing rate (see METHODS). Across the population of stimulus selective neurons, the average ROC value was 0.59 (range: 0.50–0.75) in the sample interval and 0.59 in the delay (range: 0.50–0.82).

These analyses demonstrate that a significant degree of category information was evident even across the entire population of stimulus-selective neurons. The average index or ROC values obtained were somewhat modest because activity was averaged across an entire trial epoch and across all stimulus-selective neurons. As will be shown next, the strength of category signals varied widely with individual neurons and with time; individual neurons could convey very strong category signals at particular points in the trial.

TEMPORAL CHARACTERISTICS OF CATEGORY INFORMATION. To examine the temporal dynamics of the representation of cate-




FIG. 9. Distribution of 2-category index values across the population of 67 and 32 stimulus selective neurons during the sample (A) and delay (B) epochs, respectively. Positive values indicate larger differences in neuronal firing to samples across the category boundary than within a category. Negative values indicate larger differences within category than between categories.

gory information in PFC activity, we used a sliding ROC analysis (see METHODS). For this analysis, we included neurons whose average activity in the sample and/or delay intervals was significantly category-selective (t -test on activity to all cats vs. all dogs, evaluated at $P < 0.01$, $n = 96$ neurons).

Figure 10A shows the ROC values for each neuron in 10-ms time steps. The ROC values are sorted by their magnitude separately for each time bin to better illustrate the number of neurons exhibiting ROC values >0.5 (chance) at each moment in time. This revealed that, in general, more neurons conveyed category signals late in the sample epoch than during the delay interval but that the strongest category signals occurred in the late delay and early choice presentation epoch. Figure 10 indicates that there was a greater number of neurons with moderate or small ROC values for the time bins during the sample epoch (i.e., there are more “foothills” leading up the “peaks”) but that the highest ROC values occurred during the late delay/choice presentation (the “peaks” are highest then).

EFFECTS OF LEARNING ON CATEGORY REPRESENTATIONS. As our monkeys had no prior experience with cats and dogs, it seemed likely that the category information in the PFC was acquired through learning. To test the effects of learning on category representations, we retrained one monkey with the samples reassigned to three new categories (see Fig. 1 and METHODS). We then recorded from 103 neurons at similar depths and locations as those recorded during the two-category task. The incidence of neuronal responsiveness and stimulus selectivity during the three-category task was similar to that during the two-category task: $\sim 63\%$ (65/103) of neurons were visually responsive (t -test vs. baseline, as in the preceding text, $P < 0.01$) and $\sim 23\%$ (24/103, 14 sample, 14 delay) were stimulus selective (ANOVA with stimulus as factor, $P < 0.01$).

An example of a neuron recorded during the three-category task is shown in Fig. 11. It showed a significant effect of category during the delay period when the data were sorted

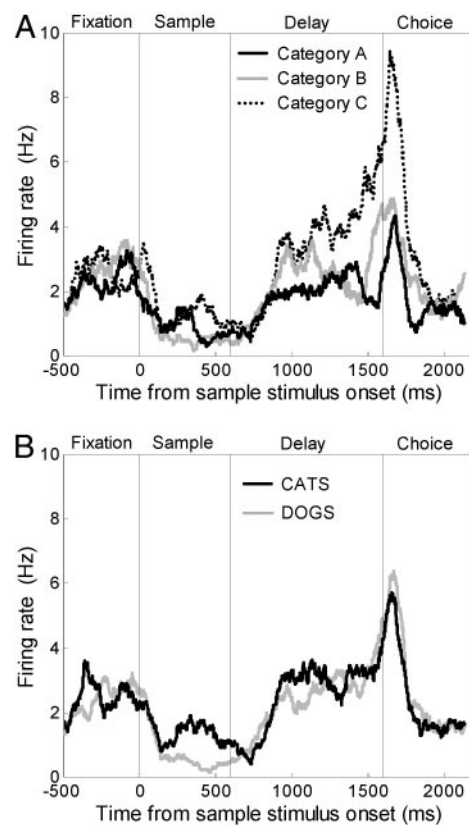


FIG. 11. An example of a single neuron recorded during the 3-category task that showed selectivity for the newly learned 3 categories (A) but not the old (now irrelevant) 2 categories (B).

according to the new, currently relevant, three-category scheme (ANOVA, $P < 0.001$); it distinguished one of the categories from the other two (Fig. 11A). By contrast, when the data were sorted according to the old (now irrelevant) cats and dogs category scheme (Fig. 11B), there was no significant difference (ANOVA, $P = 0.74$).

To test for these effects in this population of neurons, we first examined all those that were stimulus-selective ($n = 24/103$, 14 sample, 14 delay, ANOVA, $P < 0.01$). When the category index was computed using the old (now irrelevant) cat and dog categories, there was no evidence of category effects; the two-category index was not significantly greater than zero for the sample interval (2-category index = 0.01, 1-tailed t -test, $P = 0.5$) nor the delay (2-category index = -0.10 , 1-tailed t -test, $P = 0.9$). However, when the category index was computed using the new (relevant) three-category boundaries, a significant category effect was observed in the delay (3-category index = 0.16, 1-tailed t -test, $P = 0.008$). As we found for the two-category task, three-category tuning was stronger during the delay than the sample interval (2-tailed t -test, $P = 0.04$). In fact, we did not detect significant category tuning across the population of stimulus selective neurons during the sample interval (3-category index = -0.01 , $P = 0.5$), although it was detected when we computed the index for all neurons recorded during the three-category task ($n = 103$, see following text).

The same pattern of effects was observed across the entire population of neurons. Figure 12 shows the distribution of the category indices for all 103 (randomly sampled) cells recorded

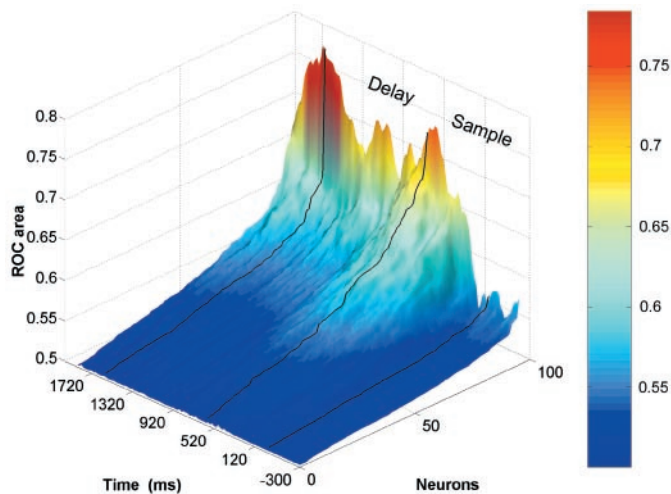


FIG. 10. Time course of category selectivity. Category selectivity across the population of 96 category selective neurons was computed using a sliding receiver-operating characteristics (ROC) analysis (see METHODS). The ROC values for all 96 neurons were sorted from minimum to maximum for each time bin independently. Higher ROC values indicating stronger category tuning. The black lines correspond (from right to left) to sample onset, offset, and choice-stimulus onset. Time is aligned to the end of the 200-ms sliding window (i.e., the values at time = 0 indicates the ROC values during the -200 to 0 time epoch relative to sample onset).

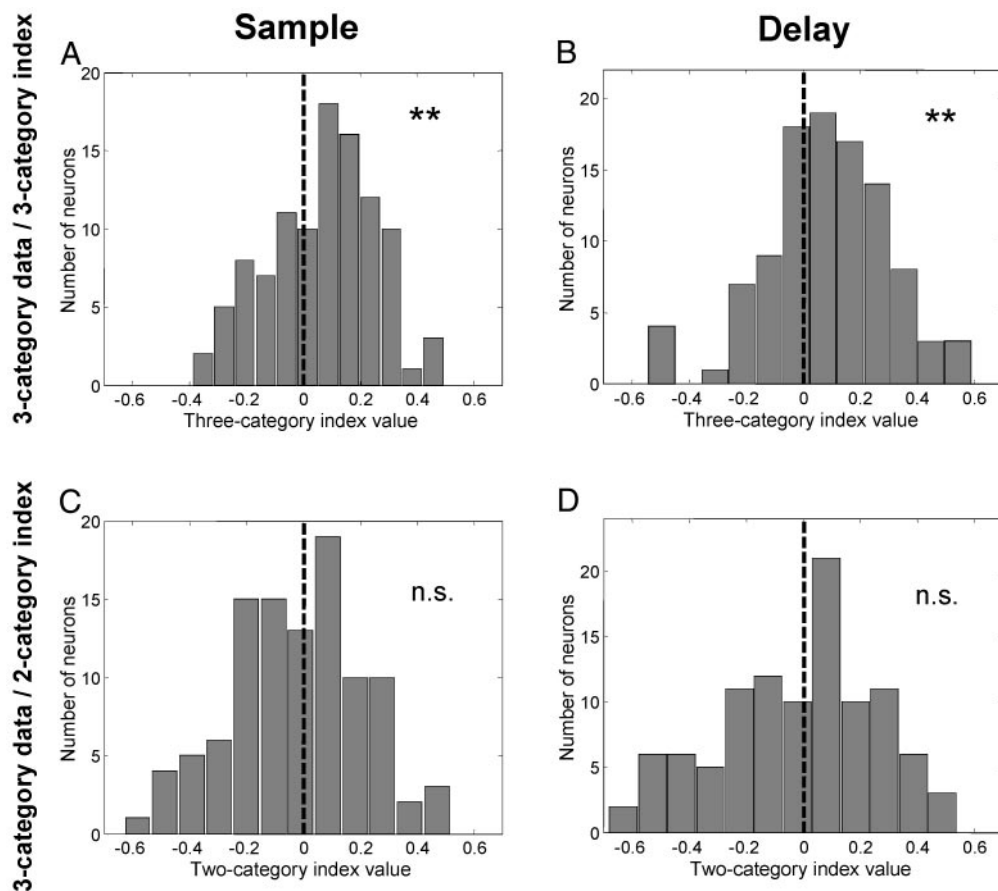


FIG. 12. Distribution of 3-category (A: sample, B: delay) and 2-category (C: sample, D: delay) index values across the entire population of 103 neurons recorded during the 3-category task. The index can range from -1 to 1 . Positive values indicate larger differences in neuronal firing between categories than within categories. Negative values indicate larger differences within categories than between categories. **, significantly positive values of the category index (1-tailed t -test, $P < 0.001$).

during the three-category task. The indices computed using the three-category scheme revealed significant category information (i.e., the distribution was shifted to the right) for both the sample interval (Fig. 12A, 3-category index = 0.065, 1-tailed t -test, $P = 0.0007$) and for the delay (Fig. 12B, 3-category index = 0.08, 1-tailed t -test, $P = 0.0005$). By contrast, when the indices were computed using the two-category scheme, there were no significant category effects during the sample (Fig. 12C, 2-category index = -0.02 , 1-tailed t -test, $P = 0.83$) or the delay interval (Fig. 12D, 2-category index = -0.03 , 1-tailed t -test, $P = 0.82$). Thus while information about the three-category scheme was evident in the population of PFC neurons, we could no longer detect information about the previously learned, now-irrelevant, cat and dog categories.

CATEGORY MATCH/NONMATCH EFFECTS. When the choice stimulus was presented, the monkeys needed to categorize it and then decide whether or not its category matched that of the sample. Both signals were present in neuronal responses to the choice stimulus. We evaluated activity in this interval with a two-way ANOVA (factor 1: choice stimulus category, factor 2: match vs. nonmatch, evaluated at $P < 0.01$). Just more than 9% (37/395) of the entire population of PFC neurons reflected the category of the choice stimulus while 11% (43/395) reflected its match/nonmatch status. More than two-thirds of the latter neurons (29/43) showed an effect of matching/nonmatching that was similar regardless of whether the choice stimulus was a cat or dog (main effect of match/nonmatch, no interaction with choice stimulus category). An example of a neuron that exhibited greater activity to matches is shown in Fig. 13A, and

an example of a neuron showing greater activity to nonmatches is shown in Fig. 13B. This activity could have encoded the monkeys' decisions about the match/nonmatch status of stimuli and/or the motor aspects of the task (the lever release to matches). The remaining third of these neurons (14/43) showed an interaction between the match/nonmatch status and the category of the choice stimulus ($P < 0.01$). In other words, they showed match/nonmatch effects that were limited, or much stronger, to one of the categories. An example of a "cat match" neuron is shown in Fig. 13C. For match/nonmatch selective neurons, a similar number preferred matches (22/43 or 51%) as nonmatches (21/43 or 49%).

ANALYSIS OF ERROR TRIALS. For insight into neuronal correlates of the monkey's errors, we compared category effects and match/nonmatch effects on correctly performed trials versus those in which monkeys made errors in category judgments. For these analyses, we included neurons that showed significant effects on correct trials. Figure 14 shows the results of these comparisons. Category information was evident during the sample interval on both correct and error trials; the average activity to the preferred versus nonpreferred category was significantly different for both types of trials (t -test, $P < 0.001$, Fig. 14A). But category information seemed to be lost in the delay. A significant difference between the average activity to the two categories was evident on correct trials ($P < 0.001$) but not on error trials ($P = 0.79$, Fig. 14B). Match/nonmatch effects also depended on whether the trial was correctly performed or not. For these analyses, the choice stimulus status (match or nonmatch) that elicited the greater activity on correct

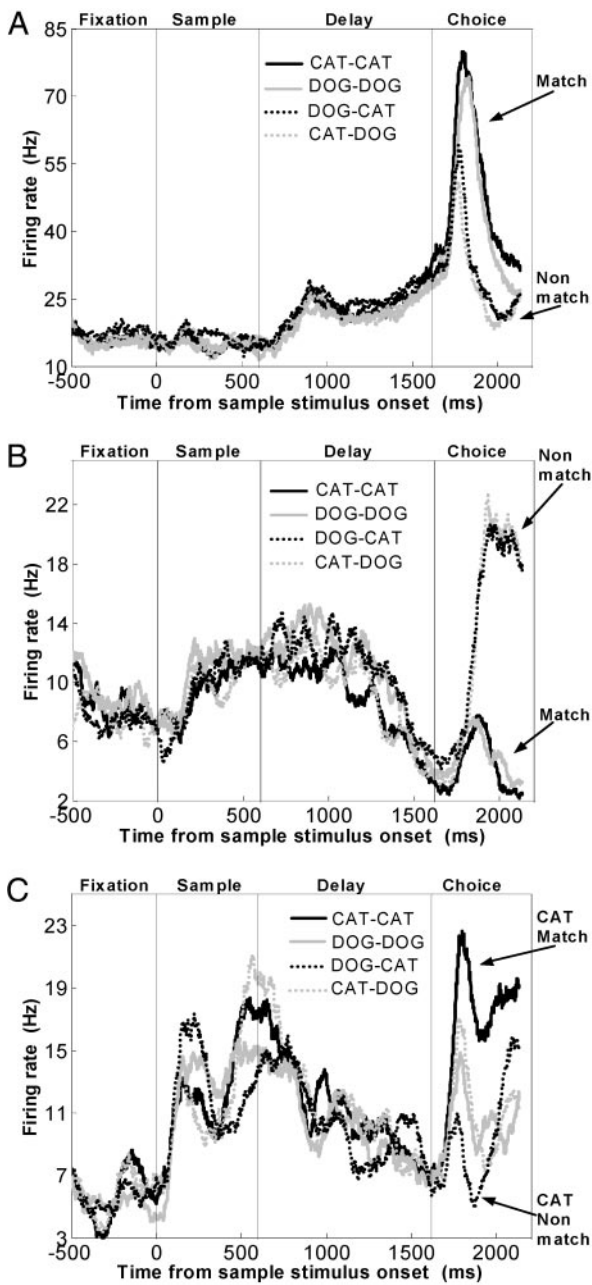


FIG. 13. Single neuron examples of match/nonmatch effects. Neuronal activity is grouped according to the category of the choice stimulus and whether or not the category of the choice stimulus matched the category of the sample. *A*: an example of a neuron that showed enhanced activity to category matches. *B*: enhanced activity to category nonmatches. *C*: a neuron that showed enhanced activity to a cat choice stimulus when it was a match and a suppressed response to a cat nonmatch. It did not differentiate between match and nonmatch trials when the choice stimulus was a dog.

trials was termed the “preferred condition.” For all neurons that showed pure match versus nonmatch effects ($n = 25$, i.e., match vs. nonmatch factor: $P < 0.01$, choice-category and interaction factors: $P > 0.01$), there was a significant difference ($P < 0.001$) in average activity to the preferred and nonpreferred conditions on correct trials. On error trials, however, the pattern reversed; there was an increase in activity to nonpreferred over preferred conditions that reached significance at $P < 0.05$ (Fig. 14C). This is presumably because the

monkeys mistakenly responded to nonmatches as if they were matches.

DISCUSSION

We report that neurons in the PFC, a brain region central to many visual behaviors, exhibited properties that mirrored the behavioral characteristics of perceptual categories. They made sharper distinctions between stimuli from different categories than between stimuli from the same category, irrespective of

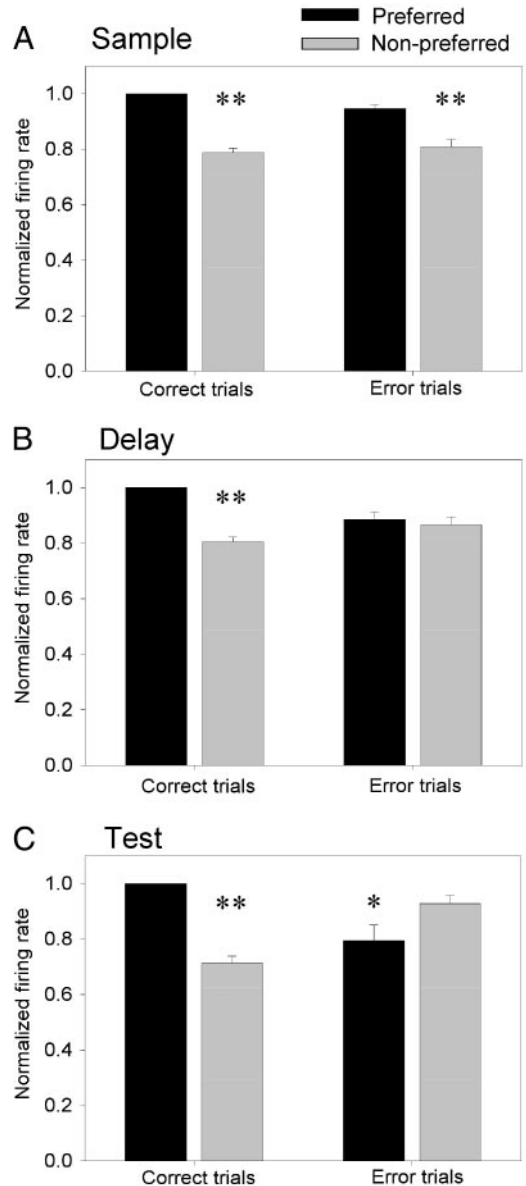


FIG. 14. Comparison of neuronal selectivity on correct and error trials. **, significance (evaluated by a *t*-test) at $P < 0.001$. *, significance at $P < 0.05$. The average response of all category selective neurons during the sample ($n = 74$; *A*) and delay ($n = 51$; *B*) phases is shown for correct and incorrect trials. Each neuron’s preferred category was determined by the category that evoked greater average firing rates during correct trials. *C*: average activity to choice stimuli for all match/nonmatch selective neurons ($n = 25$ neurons with only a significant effect of match/nonmatch) for correct and incorrect trials. Each neuron’s preferred response (match or nonmatch) was determined according to the group that elicited greater average firing rates during correctly executed trials.

their relative physical similarity. This explicit encoding of category membership in the activity of single neurons did not have to be the case. In principle, categories might have only been reflected on the ensemble level, as an emergent property of neurons encoding different defining features. Our results illustrate instead that familiar categories are reflected on the single-neuron level, much as physical attributes of stimuli are. This ability to carve category membership into the tuning of single neurons might allow for the quick and effortless classification of familiar objects. We also observed neuronal correlates of category match/nonmatch effects, suggesting a role for the PFC in these judgments and/or in issuing resulting motor commands. Finally, the observation that neuronal correlates of categories and category judgments waned or changed on error trials suggests that PFC activity was directly related to task performance.

The presence of category information in the PFC makes sense given its position at the apex of the perception-action cycle (Fuster 1990). Categories are defined by their functional relevance. Therefore they might be strongly represented in a brain area that mediates the functions needed to transform perceptions into voluntary actions, functions such as the integration of temporally separated events (Fuster et al. 2000), the acquisition and representation of behavior-guiding rules (Asaad et al. 1998; Wallis et al. 2001; White and Wise 1999) and visuomotor decisions (Kim and Shadlen 1999). The relative specialization of PFC in guiding behavior is reflected in the fact that its damage or reversible inactivation in monkeys cause deficits in performance of tasks demanding attention, working memory and response inhibition (Dias et al. 1996; Funahashi et al. 1993; Goldman and Rosvold 1970; Goldman et al. 1971; Gross and Weiskrantz 1962; Mishkin 1957; Mishkin and Manning 1978; Mishkin et al. 1969; Passingham 1975) but usually spares more purely perceptual functions such as object recognition, visual long-term memory, and “high level” visual analysis of form.

But an important contribution must also come from brain areas that mediate these visual functions, such as the ITC. Its damage causes deficits in visual discrimination, recognition, and learning (Blum et al. 1950; Kluver and Bucy 1938, 1939; Mishkin 1954, 1966; Mishkin and Pribram 1954) and category-specific agnosias (e.g., for faces) in humans (Damasio et al. 1982). Since the seminal work of Gross and coworkers, who reported a small population of “face cells,” numerous studies have shown that ITC neurons show selectivity for objects that cannot be explained by sensitivity to low-level features, such as orientation or color (Desimone et al. 1984; Gross et al. 1972; Perret et al. 1992; Kobatake and Tanaka 1994; Tanaka et al. 1991). There has even been some recent evidence that suggests that these neurons play a direct role in categorization. Vogels (1999) recorded from the ITC in monkeys trained to categorize stimuli as tree versus nontree or fish versus nonfish and found that many neurons were selectively activated by the trained class (photographs of trees or fish) but not by distracter objects (photos of household objects or scenes containing neither trees nor fish). Kreiman et al. (2000) recorded from medial temporal lobe neurons in epileptic human patients while they classified stimuli into nine categories (e.g., faces, cars, food) and found neurons that selectively responded to stimuli from one of the categories. However, it has not been clear whether ITC neuronal selectivity encodes the category membership of stimuli,

their physical appearance or some combination of these two factors. With a large, amorphous set of stimuli (such as trees or food), the category boundaries are unknown and the sharp transitions that are diagnostic of categories cannot be evaluated independently of stimulus similarity. Hence, neuronal selectivity for, say, trees could reflect the fact that trees look more like one another than other stimuli. Our results indicate that PFC neurons can convey information about the category of stimuli largely irrespective of their physical appearance.

The relative roles of the PFC and ITC in perceptual categorization remain to be determined. A recent theory of object recognition suggests that category tuning in the PFC could arise from converging inputs from ITC neurons that are stimulus, but not category, tuned (Riesenhuber and Poggio 2000). In this model, category-tuned neurons perform a weighted sum of the inputs from neurons broadly tuned for individuals followed by a thresholding operation. This suggests a greater role for the PFC in the explicit representation of categories. Another possibility is that category information is “loaded” into the PFC from long-term storage in the ITC. A recent study by Tomita et al. (1999) suggested that recall of long-term visual memories involved top-down signals from the PFC that activate representations stored in the ITC. Similar mechanisms might mediate the retrieval of category information stored in the ITC.

In sum, our results have provided insight into how perceptual categories and category-related behaviors are encoded in the PFC, a brain area that receives the outputs of sensory cortex and helps mediate voluntary action. How and whether category membership is encoded in sensory systems and the respective roles of the PFC and visual areas like the ITC in representing and storing category information remains to be determined.

We thank C. Shelton for the morphing software and K. Anderson, D. Applewhite, W. Asaad, M. Machon, M. Mehta, A. Nieder, A. Pasupathy, J. Wallis, and M. Wicherski for valuable comments, help, and discussions.

This work was supported by a National Institute of Mental Health grant, a National Science Foundation-Knowledge and Distributed Intelligence grant, RIKEN-MIT Neuroscience Research Center, a McDonnell-Pew Fellowship (M. Riesenhuber), the Whitaker Chair (T. Poggio), and the Class of 1956 Chair (E. K. Miller).

REFERENCES

- ASAAD WF, RAINER G, AND MILLER EK. Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21: 1399–1407, 1998.
- BEALE JM AND KEIL FC. Categorical effects in the perception of faces. *Cognition* 57: 217–239, 1995.
- BEYMER D AND POGGIO T. Image representations for visual learning. *Science* 272: 1905–1909, 1996.
- BLUM JS, CHOW KL, AND PRIBRAM KH. A behavioral analysis of the organization of the parieto-temporo-preoccipital cortex. *J Comp Neurol* 93: 53–100, 1950.
- BOOTH MC AND ROLLS ET. View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cereb Cortex* 8: 510–523, 1998.
- BRUCE C, DESIMONE R, AND GROSS CG. Visual properties of neurons in a polysensory area in superior temporal sulcus in the macaque. *J Neurophysiol* 46: 369–384, 1981.
- DAMASIO AR, DAMASIO H, AND VAN HOESEN GW. Prosopagnosia: anatomic basis and behavioral mechanisms. *Neurology* 32: 331–341, 1982.
- DESIMONE R, ALBRIGHT TD, GROSS CG, AND BRUCE C. Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4: 2051–2062, 1984.
- DIAS R, ROBBINS TW, AND ROBERTS AC. Dissociation in prefrontal cortex of affective and attentional shifts. *Nature* 380: 69–72, 1996.

- FABRE-THORPE M, RICHARD G, AND THORPE SJ. Rapid categorization of natural images by rhesus monkeys. *Neuroreport* 9: 303–308, 1998.
- FREEDMAN DJ, RIESENHUBER M, POGGIO T, AND MILLER EK. Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291: 312–316, 2001.
- FUNAHASHI S, BRUCE CJ, AND GOLDMAN-RAKIC. Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic “scotomas.” *J Neurosci* 13: 1479–1497, 1993.
- FUSTER JM. Prefrontal cortex and the bridging of temporal gaps in the perception-action cycle. *Ann NY Acad Sci* 608: 318–329, 1990.
- FUSTER JM. Animal neuropsychology. In: *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*. Philadelphia, PA: Lippincott-Raven, 1997, p. 66–101.
- FUSTER JM, BODNER M, AND KROGER JK. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature* 405: 347–351, 2000.
- GOLDMAN-RAKIC PS. Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In: *Handbook of Physiology: The Nervous System. Higher Functions of the Brain*. Bethesda, MD: Am. Physiol. Soc., 1987, sect. 1, vol. V, p. 373–417.
- GOLDMAN PS AND ROSVOLD HE. Localization of function within the dorsolateral prefrontal cortex of the rhesus monkey. *Exp Neurol* 27: 291–304, 1970.
- GOLDMAN PS, ROSVOLD HE, VEST B, AND GALKIN TW. Analysis of the delayed-alternation deficit produced by dorsolateral prefrontal lesions in the rhesus monkey. *J Comp Physiol Psychol* 77: 212–220, 1971.
- GREEN DM AND SWETS JA. *Signal Detection Theory and Psychophysics*. New York: Wiley, 1966.
- GROSS CG. Visual functions of inferotemporal cortex. In: *Handbook of Sensory Physiology*, edited by Autrum H, Jung R, Lowenstein W, McKay D, and Teuber H-L. Berlin: Springer-Verlag, 1973, vol. VII/3B, p. 451–482.
- GROSS CG, ROCHA-MIRANDA CE, AND BENDER DB. Visual properties of neurons in inferotemporal cortex of the macaque. *J Neurophysiol* 35: 96–111, 1972.
- GROSS CG AND WEISKRANTZ L. Evidence for dissociation of impairment on auditory discrimination and delayed response following lateral frontal lesions in monkeys. *Exp Neurol* 5: 453–476, 1962.
- KIM JN AND SHADLEN MN. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2: 176–185, 1999.
- KLÜVER H AND BUCY L. An analysis of certain effects of bilateral temporal lobectomy in the rhesus monkey, with special reference to “psychic blindness.” *J Psychol* 5: 33–54, 1938.
- KLÜVER H AND BUCY L. Preliminary analysis of functions of the temporal lobes in monkeys. *Arch Neurol Psychiatry* 42: 979–1000, 1939.
- KOBATAKE E AND TANAKA K. Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *J Neurophysiol* 71: 856–867, 1994.
- KOBATAKE E, WANG G, AND TANAKA K. Effects of shape discrimination training on the selectivity of inferotemporal cells in adult monkeys. *J Neurophysiol* 80: 324–330, 1998.
- KREIMAN G, KOCH C, AND FRIED I. Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci* 3: 946–953, 2000.
- LIEBERMAN AM, COOPER FS, SHANKWEILER DP, AND STUDDERT-KENNEDY M. Perception of the speech code. *Psychol Rev* 74: 431–461, 1967.
- LOGOTHETIS NK, PAULS J, AND POGGIO T. Shape representation in the inferior temporal cortex of monkeys. *Curr Biol* 5: 552–563, 1995.
- LOGOTHETIS NK AND SHEINBERG. Visual object recognition. *Annu Rev Neurosci* 19: 577–621, 1996.
- MILLER EK. The prefrontal cortex and cognitive control. *Nat Rev Neurosci* 1: 59–65, 2000.
- MILLER EK AND COHEN JD. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24: 167–202, 2001.
- MILLER EK, ERICKSON CA, AND DESIMONE R. Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13: 1460–1478, 1993.
- MILLER EK, ERICKSON CA, AND DESIMONE R. Neuronal mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16: 5154–5167, 1996.
- MISHKIN M. Visual discrimination performance following partial ablations of the temporal lobe. II. Ventral surface vs. hippocampus. *J Comp Physiol Psychol* 47: 187–193, 1954.
- MISHKIN M. Effects of small frontal lesions on delayed alternation in monkeys. *J Neurophysiol* 20: 615–622, 1957.
- MISHKIN M. Visual mechanisms beyond the striate cortex. In: *Frontiers in Physiological Psychology*, edited by Russell RW. New York: Academic, 1966, p. 93–119.
- MISHKIN M AND MANNING FJ. Non-spatial memory after selective prefrontal lesions in monkeys. *Brain Res* 143: 313–323, 1978.
- MISHKIN M AND PRIBRAM K. Visual discrimination performance following partial ablations of the temporal lobe. I. Ventral vs. lateral. *J Comp Physiol Psychol* 47: 14–20, 1954.
- MISHKIN M, VEST B, WAXLER M, AND ROSVOLD HE. A re-examination of the effects of frontal lesions on object alternation in monkeys. *Neuropsychologia* 7: 357–364, 1969.
- MIYASHITA Y. Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature* 335: 817–820, 1988.
- NICHOLS AM, RUFFNER TW, SOMMER MA, AND WURTZ RH. A Screw microdrive for adjustable chronic unit recording in monkeys. *J Neurosci Methods* 81: 185–188, 1998.
- ORLOV T, YAKOVLEV V, HOCHSTEIN S, AND ZOHARY E. Macaque monkeys categorize images by their ordinal number. *Nature* 404: 77–80, 2000.
- PASSINGHAM RE. Delayed matching after selective prefrontal lesions in monkeys (*Macaca mulatta*). *Brain Res* 92: 89–102, 1975.
- PERRETTI DI, HIETANEN JK, ORAM MW, AND BENSON PJ. Organization and functions of cells responsive to faces in the temporal cortex. *Philos Trans Roy Soc Lond B Biol Sci* 335: 23–30, 1992.
- PERRETTI DI, ROLLS ET, AND CAAN W. Visual neurons responsive to faces in the monkey temporal cortex. *Exp Brain Res* 47: 329–342, 1982.
- RAINER G AND MILLER EK. Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron* 27: 179–189, 2000.
- RIESENHUBER M AND POGGIO T. Models of object recognition. *Nature Neurosci* 3 Suppl: 1199–1204, 2000.
- ROBERTS WA AND MAZMANIAN DS. Concept at different levels of abstraction by pigeons, monkeys and people. *J Exp Psychol Anim Behav Process* 14: 247–260, 1988.
- SCALADHE SP, WILSON FA, AND GOLDMAN-RAKIC PS. Face selective neurons during passive viewing and working memory performance of rhesus monkeys: evidence for intrinsic specialization of neuronal coding. *Cereb Cortex* 9: 459–475, 1999.
- SHELTON C. Morphable surface models. *Int J Comp Vis* 38: 75–91, 2000.
- TANAKA K. Inferotemporal cortex and object vision. *Annu Rev Neurosci* 19: 109–139, 1996.
- TANAKA K, SAITO H, FUKADA Y, AND MORIYA M. Coding visual images of objects in the inferotemporal cortex of monkeys. *J Neurophysiol* 66: 170–189, 1991.
- TOLHURST DJ, MOVSHON JA, AND DEAN AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23: 775–785, 1983.
- TOMITA H, OHBAYASHI M, NAKAHARA K, HASEGAWA I, AND MIYASHITA Y. Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* 401: 699–703, 1999.
- UNGERLEIDER LG, GAFFAN D, AND PELAK VS. Projections from inferior temporal cortex to prefrontal cortex via the uncinate fascicle in rhesus monkeys. *Exp Brain Res* 76: 473–484, 1989.
- VOGELS R. Categorization of complex visual images by rhesus monkeys. *Eur J Neurosci* 11: 1223–1238, 1999.
- VOGELS R AND ORBAN GA. How well do response changes of striate neurons signal differences in orientation: a study in the discriminating monkey. *J Neurosci* 10: 3542–3558, 1990.
- WALLIS JD, ANDERSON KC, AND MILLER EK. Single neurons in prefrontal cortex encode abstract rules. *Nature* 411: 953–956, 2001.
- WEBSTER MJ, BACHEVALIER J, AND UNGERLEIDER LG. Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. *Cereb Cortex* 4: 470–483, 1994.
- WHITE IM AND WISE SP. Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res* 126: 315–335, 1999.
- WYTTENBACH RA, MAY ML, AND HOY RR. Categorical perception of sound frequency by crickets. *Science* 273: 1542–1544, 1996.