

A Comparison of Primate Prefrontal and Inferior Temporal Cortices during Visual Categorization

David J. Freedman,^{1,2,5} Maximilian Riesenhuber,^{3,4,5} Tomaso Poggio,^{3,4,5} and Earl K. Miller^{1,2,5}

¹The Picower Center for Learning and Memory, ²RIKEN-MIT 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

Previous studies have suggested that both the prefrontal cortex (PFC) and inferior temporal cortex (ITC) are involved in high-level visual processing and categorization, but their respective roles are not known. To address this, we trained monkeys to categorize a continuous set of visual stimuli into two categories, “cats” and “dogs.” The stimuli were parametrically generated using a computer graphics morphing system (Shelton, 2000) that allowed precise control over stimulus shape. After training, we recorded neural activity from the PFC and the ITC of monkeys while they performed a category-matching task. We found that the PFC and the ITC play distinct roles in category-based behaviors: the ITC seems more involved in the analysis of currently viewed shapes, whereas the PFC showed stronger category signals, memory effects, and a greater tendency to encode information in terms of its behavioral meaning.

Key words: categorization; monkey; vision; object vision; inferior temporal cortex; prefrontal cortex; learning

Introduction

Comparing and contrasting the neural properties of different brain regions can yield important insight into their respective contributions and, hence, the neural circuitry underlying a given function. Take, for example, perceptual categorization, a process fundamental for normal cognition because it gives meaning to our sensory environment. Several recent studies have reported neuronal correlates of visual categories in two interconnected cortical areas involved in visual recognition, memory, and other visual functions: the inferior temporal cortex (ITC) and the prefrontal cortex (PFC) (Vogels, 1999; Freedman et al., 2001, 2002; Nieder et al., 2002; Sigala and Logothetis, 2002). However, the respective roles of these and other brain areas in categorization remain essentially unknown. The PFC and the ITC have been studied by different investigators using different behavioral paradigms, different stimuli, etc., which are confounding factors that render comparisons between them difficult at best. So we trained two monkeys to perform a category judgment task and directly compared neuronal activity in the ITC and the PFC.

As in a previous study that focused on the PFC, we used a continuous set of cat and dog stimuli (see Fig. 1) constructed from six prototypes with a three-dimensional morphing system (Shelton, 2000). By morphing different amounts of the prototypes, we could generate thousands of unique images, continuously vary stimulus shape, and precisely define a category boundary. The category of a stimulus was defined by whichever category

contributed more (>50%) to a given morph. We used a behavioral paradigm that required monkeys to release a lever if two stimuli (separated by a 1 sec delay) were from the same category (a category match). Thus, the behavioral responses (release or hold) indicated “match” or “non-match,” respectively, and were not directly linked to category membership. This allowed us to attribute any neuronal signals related to the category of a stimulus to perceptual categorization because the behavioral responses did not differentiate between the categories.

We previously reported neurons that encoded the cat and dog categories in the lateral PFC (Freedman et al., 2001, 2002). One possibility is that category representations are encoded “upstream” from the PFC (i.e., in areas closer to sensory input, such as the ITC) and then this information is merely copied to the PFC via direct interconnections between it and the ITC (Ungerleider et al., 1989; Webster et al., 1994). ITC neurons, after all, seem to have properties appropriate for a role in visual categorization. They are selectively activated by complex visual stimuli (Gross, 1973; Bruce et al., 1981; Perrett et al., 1982; Desimone et al., 1984; Logothetis and Sheinberg, 1996; Tanaka, 1996), influenced by visual experience (Miyashita, 1988; Logothetis et al., 1995; Booth and Rolls, 1998; Kobatake et al., 1998; Baker et al., 2002), and sensitive to diagnostic features of categories (Sigala and Logothetis, 2002). Thus, neuronal category selectivity might be similar between the two areas or even stronger in the ITC than the PFC. Alternatively, the PFC may play a more active role in categorization. One model of visual recognition (Riesenhuber and Poggio, 2000) suggests that the PFC further enhances the behaviorally relevant aspects of the information that it receives from the ITC. PFC neuronal activity does emphasize behaviorally relevant information (Rainer et al., 1998), and categories are typically defined by their functionality. This predicts a stronger representation of category in the PFC than ITC. Other questions concern the mechanisms that guide category-based behaviors, such as those involved in behavioral decisions. Neuronal correlates of the

Received Jan. 28, 2002; revised March 21, 2003; accepted March 27, 2003.

This work was supported by National Institute of Mental Health Grant MH22588, a National Science Foundation–Knowledge and Distributed Intelligence grant, the RIKEN-MIT Neuroscience Research Center, and a McDonnell-Pew award in cognitive neuroscience (M.R.). Additional support was provided by the Eugene McDermott Foundation and the Whitaker Foundation (T.P.). We thank C. Shelton for the morphing software and K. Anderson, W. Asaad, M. Machon, A. Nieder, A. Pasupathy, J. Wallis, and M. Wickerski for valuable comments, help, and discussions, and Ulf Knoblich for performing the model simulations and creating Figure 7.

Correspondence should be addressed to David J. Freedman, Building E25, Room 236, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: davidf@mit.edu.

Copyright © 2003 Society for Neuroscience 0270-6474/03/235235-12\$15.00/0

match/non-match status of stimuli have been reported in the PFC and the ITC during identity judgments (Miller et al., 1996), but these areas have not been compared during category judgments. We address these questions here.

Materials and 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 recording hardware. Eye movements were monitored and stored using an infrared eye tracking system (Iscan, Cambridge, MA) at a sampling rate of 120 Hz. 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 Massachusetts Institute of Technology Committee on Animal Care.

Stimuli and task. A large continuous set of images was generated from three cat prototypes and three dog prototypes (Fig. 1) using a novel algorithm (Shelton, 2000; Freedman et al., 2001, 2002). 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/projects/cbcl/morph>. By morphing different amounts of the prototypes we could generate thousands of unique images, continuously vary shape, and precisely define a category boundary. The category of a stimulus was defined by whichever category contributed more (>50%) to a given morph. The stimuli differed along multiple feature dimensions and were smoothly morphed, i.e., without the 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 trained monkeys to perform a delayed match-to-category task (DMC) (Fig. 2). The monkeys viewed two stimuli that were separated by a brief delay. They were trained to indicate (by releasing a lever) whether the second (test) stimulus was from the same category as a previously seen (sample) stimulus. Thus, monkeys learned to group a continuous set of stimuli into discrete categories. Performance was high (~90% correct) even when the sample stimuli were close to the category boundary and therefore strongly resembled stimuli from the other category. Note that by using this task design, the monkey's responses (release or hold) indicated match or non-match, and neither was uniquely associated with either category. Furthermore, an equal proportion of match and non-match trials were presented in a pseudorandom order, and the monkeys could not predict whether a trial would be a match or non-match trial until the test stimulus appeared. Therefore, any neuronal signals related to stimulus category must be related to perceptual categorization and not the motor responses because the responses did not differentiate between the categories. Of course, because the motor response (lever release) always indicated match, any neuronal signals during the test epoch that differentiated between match and non-match could be related to the monkeys' decision-making process or preparation for motor responses, or both.

During recordings, we used 54 stimuli as samples, each belonging to one of six levels of cat/dog blends (cat/dog) (100:0, 80:20, 60:40, 40:60, 20:80, 0:100) along the nine morph lines that crossed the category boundary and two levels along the six within-category morph lines (60:

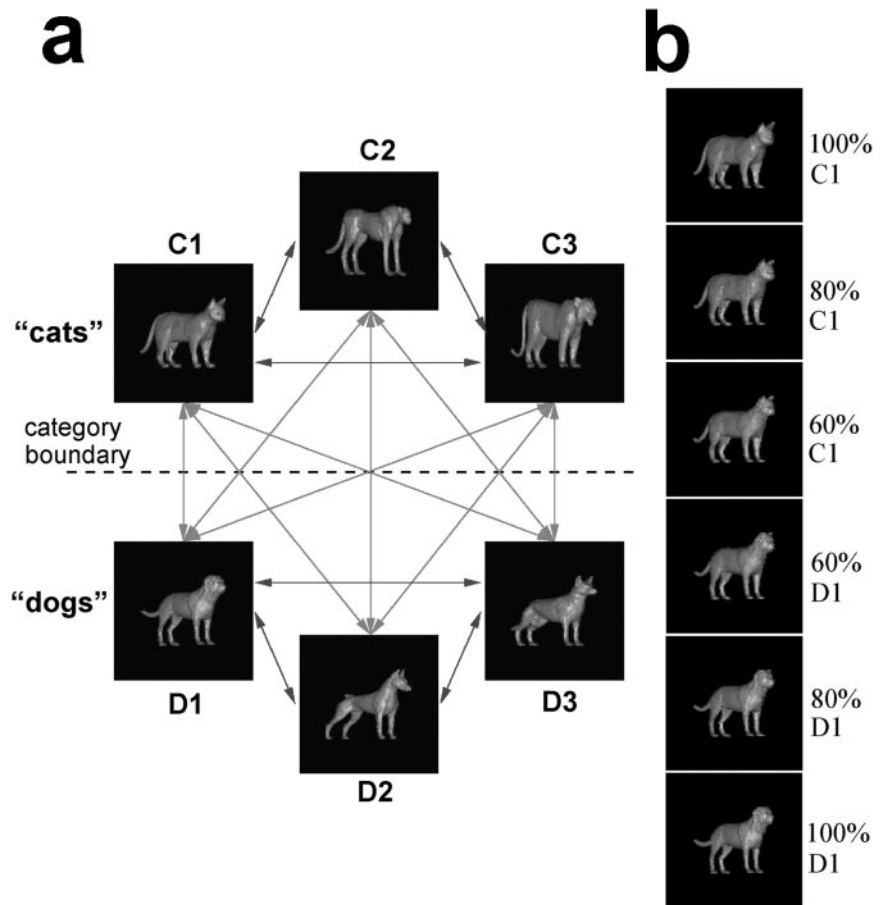


Figure 1. Organization of stimulus set. *A*, The 6 prototype images and 15 morph lines. The sample stimulus set was composed of 54 unique images: 6 prototypes (as shown), 4 images evenly placed (20, 40, 60, and 80%) along the nine lines connecting each cat to each dog prototype, and 2 images (at 40 and 60%) along each of the six lines between prototypes of the same category. *B*, An example of the morphs generated between the C1 and D1 prototypes.

40, 40:60) (Fig. 1*a*). Before recordings, monkeys were trained with thousands of randomly generated cat and dog stimuli that covered the vast majority of possible morphs using all combinations of the six prototypes. Thus, monkeys were not trained to simply memorize the 54 sample stimuli used during neuronal recordings. To prevent monkeys from memorizing specific stimulus–response contingencies during the recording experiments, the test stimuli were 100 randomly generated morphs from each category that were randomly paired with sample stimuli of the appropriate category. The set of test stimuli was frequently regenerated to further discourage monkeys from learning associations between specific sample and test images. The test stimuli unambiguously belonged to a given category: they were always chosen to be at a distance of at least 20% from the boundary.

Behavioral and recording methods. Monkeys typically performed >10 correct repetitions for each of the 54 unique sample stimuli. Eye movements were monitored and stored using an infrared eye tracking system (ISCAN, Cambridge, MA), and monkeys were required to maintain gaze within a $\pm 2^\circ$ fixation window. Breaks of fixation were not included when computing error rates.

PFC recording chambers were placed stereotaxically over the principal sulcus and anterior to the arcuate sulcus using structural magnetic resonance imaging (MRI) scans acquired before surgery. PFC recordings were obtained primarily from the areas ventral to the principal sulcus (areas 45, 46, and 12). ITC recordings were conducted between anterior–posterior = 14–20 mm and lateral = 15–21 mm (Fig. 3). ITC recording locations, as determined by stereotaxic coordinates, MRI scans, and white–gray matter transitions encountered during electrode penetrations were in areas TEa, TEm, TE2, and TE1 (Paxinos et al., 2000). The

a

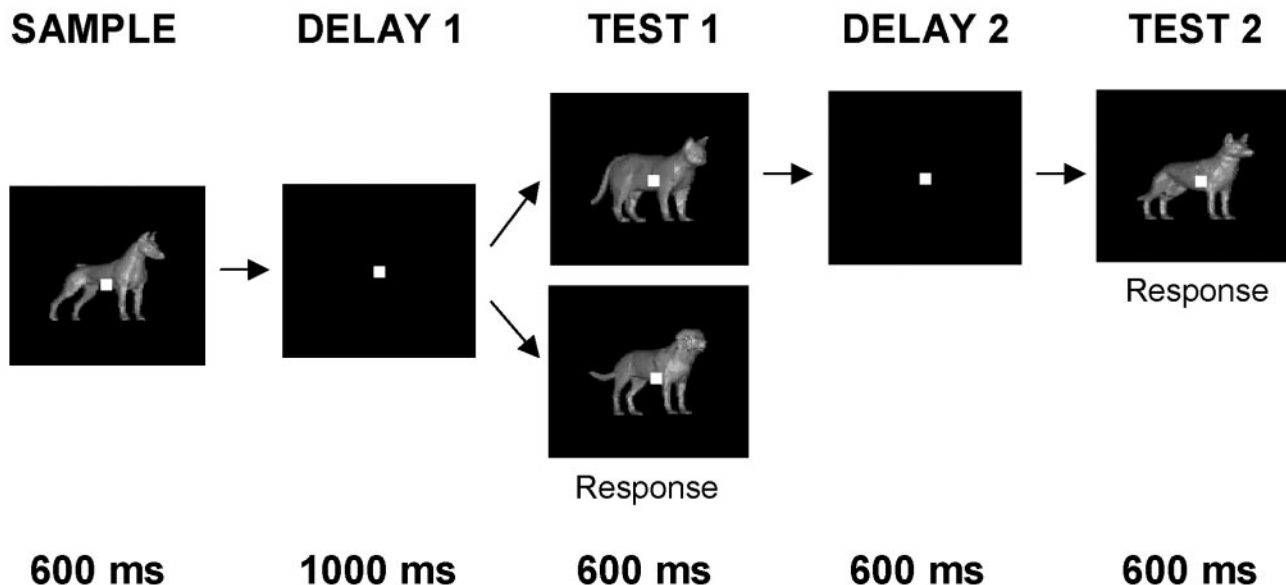


Figure 2. Behavioral task. A sample stimulus (600 msec) was followed by a 1 sec delay and a test stimulus (600 msec). If the sample and test were from the same category (a match), monkeys released a lever to receive a juice reward. If the test stimulus did not match the category of the sample, there was another brief delay that was always followed by a match. Because the monkeys' response indicated match and was not uniquely associated with either category, any differential neuronal responses to the two categories could not be related to the behavioral response. Performance of the monkeys was excellent (>90% correct) even for sample stimuli that were very close to the category boundary (e.g., the 60% morphs) and resembled stimuli from the other category (Freedman et al., 2002).




Figure 3. Location of ITC and PFC recordings. *A*, The vertical dotted white lines superimposed on this structural MRI scan (from monkey B obtained just before recordings, 14 mm anterior to the intra-aural line) show the medial and lateral bounds of the ITC recording region. The horizontal dotted lines show the dorsal-ventral range. The arrows indicate the positions of the superior temporal sulcus (STS), rhinal fissure (RF), and anterior medial temporal sulcus (AMT). ITC recordings were concentrated in the lower bank of the STS (areas TEa, TEm) and the ventral surface of the inferior temporal cortex lateral to the AMT (area TE). ITC neurons were recorded between 14 and 20 mm anterior from the intra-aural line. The placement of the recording chambers and location of ITC recordings were identical for the two monkeys. *B*, The dorsal and ventral bounds of the PFC recordings are indicated by the two dotted white lines superimposed on this MRI image (from monkey B, 20 mm anterior to the intra-aural line). The arrows indicate the superior arcuate sulcus (SAR), principal sulcus (PS), and inferior arcuate sulcus (IAR). For a detailed description of recording locations, see Materials and Methods.

locations of ITC recordings were similar to those reported in studies by several laboratories (Logothetis et al., 1995; Booth and Rolls, 1998; Kobatake et al., 1998; Baker et al., 2002). No attempt was made to prescreen neurons for task-related activity such as stimulus or category selectivity. Neuronal waveforms were amplified, digitized, and stored for off-line sorting into individual neuron records using principal components analysis clustering software (Plexon Inc., Dallas, TX).

Data analysis. For the analysis of neural activity related to the sample

stimulus, average neuronal activity was calculated in four time epochs: baseline, sample presentation, delay, and test stimulus epochs. Baseline activity was averaged over the 500 msec of fixation preceding sample presentation. Sample period activity was averaged over a 600 msec epoch beginning 100 msec after sample onset. Delay activity was assessed over an 800 msec epoch beginning 300 msec after sample offset and ending 100 msec after test stimulus onset. After the memory delay, a test stimulus appeared, and the monkey had to decide whether it matched the category of the previously presented sample stimulus. Therefore, the monkey needed to maintain information about the previously presented sample stimulus in short-term memory. To analyze signals related to the previously presented sample stimulus during the test epoch, neuronal activity was averaged over an epoch beginning at the onset of the test stimulus and ending 2 SDs before the monkeys' average reaction time (RT) during each recording session to exclude any effects related to the execution of the behavioral response (monkey A: mean RT: 242 msec, mean test epoch duration: 183 msec; monkey B: mean RT: 321 msec, mean test epoch duration: 254 msec).

During the test epoch, the monkey also needed to represent the category of the currently visible test stimulus and whether it was a categorical match to the sample stimulus. To analyze category selectivity for the currently visible test stimulus and match/non-match effects, neuronal activity was averaged over an epoch that began 100 msec after test stimulus onset (to account for the latency of the visual response in PFC and ITC) and ended 2 SDs before the monkeys' average RT during each recording session to exclude any effects related to the




Figure 4. Single neuron examples. *A*, The average activity of a category-sensitive ITC neuron to the six levels of morphs. The dotted vertical lines correspond (from left to right) to sample onset, offset, and test stimulus onset. *B*, A category-sensitive PFC neuron. The color plot below each histogram (*C*, *D*) shows the average activity of each neuron to each of the 42 stimuli along the nine between-class morph lines (Fig. 1). The prototypes (C1, C2, C3, D1, D2, D3) are represented in the outermost columns, and the category boundary is represented by the blue line in the middle. Each prototype contributes to three morph lines. A color scale indicates the activity level. For the ITC neuron, activity was averaged over the first half of the sample epoch. For the color plot of the PFC neuron, activity was averaged across the delay and test epoch. This figure also shows each the category index for each neuron and ROC values for that epoch (see Results). The PFC neuron shows sharper between-category differences and less within-category variability than the ITC neuron.

execution of the behavioral response (monkey A: mean RT: 242 msec, mean test epoch duration: 83 msec; monkey B: mean RT: 321 msec, mean test epoch duration: 154 msec).

To determine the time course of neuronal activity, we computed normalized response histograms across the populations of PFC and ITC neurons that showed a significant difference in their activity to all cats versus all dogs (t test, $p < 0.01$) during the sample and/or delay epoch. To ensure that each neuron contributed equally to the population histogram, the activity of each neuron was normalized according to its minimum and maximum value (although qualitatively similar results were observed using the raw, non-normalized, firing rates). To determine the latency of neuronal responses, we first determined the mean and SD of each sample epoch category-selective neuron's average firing rate across the 500 1 msec time bins during the fixation period. The response latency of each neuron was then defined as the time at which the mean fixation period firing rate was exceeded by 3 SDs for three consecutive time bins.

To evaluate the strength of category tuning, an index was calculated using each neuron's average difference in activity to pairs of stimuli along the morph lines that crossed the category boundary. These included pairs of stimuli that were adjacent [within-category differences (WCDs): 100 vs 80% and 80 vs 60% cat or dog; between-category differences (BCDs): 60% cat vs 60% dog] and pairs that differed by "two steps" (WCD: 100 vs 60%; BCD: 80% cat vs 60% dog and vice versa)]. Thus, the average morph distance between stimuli was identical for BCDs and WCDs. A standard contrast index was computed for each neuron by dividing the difference between BCDs and WCDs by their sum, giving values ranging from -1 to 1 . Positive values of the category index indicated larger dif-

ferences between categories (BCD) compared with within a category (WCD), whereas negative values indicate the opposite.

The time course of category selectivity was evaluated by a sliding receiver operating characteristics (ROC) analysis (Freedman et al., 2002). This was computed by calculating an ROC value for the activity of each neuron to all trials with a cat sample versus all trials with a dog sample over a 200 msec window that was stepped in 10 msec intervals over the course of the trial. Qualitatively similar results were observed with ROC windows of various widths. The latency of category selectivity was determined using similar methods as used to compute response latency, but using each neuron's sliding ROC values (rather than the average neuronal response as used for computing response latency). For each sample period category-selective neuron, we first determined the mean and SD of the sliding ROC values during the fixation period. The selectivity latency of each neuron was defined as the time at which the mean fixation period ROC value was exceeded by 3 SDs for three consecutive time epochs. The time course of match/non-match selectivity during the test epoch was determined using similar methods. For each test epoch match/non-match selective neuron, a sliding ROC analysis (window width of 100 msec) that compared all match trials versus all non-match trials was computed in consecutive 10 msec steps. The latency of each neuron was defined as the point at which ROC values exceeded mean baseline levels (computed over the 500 msec before test stimulus onset) by 3 SDs for three consecutive time epochs.

Comparison of neuronal tuning to units from a hierarchical model of object recognition. To look at the contribution of shape versus category tuning in neuronal activity, we used the model of object recognition in




Figure 5. Neuronal responses to individual sample stimuli. These six color plots show the average activity of three ITC (*a–c*) and three PFC (*d–f*) neurons to stimuli along each of the nine between-class morph lines (Fig. 1). Conventions are the same as the color plots in Figure 4. The category index of each neuron and ROC values are indicated below each color plot. The plots are arranged into three columns. The columns, from left to right, show examples of neurons with selective activity during the sample, delay, and test epochs, respectively. These plots show the activity of six different neurons and not the activity of an individual neuron from each area across time.

cortex of Riesenhuber and Poggio (1999, 2000), shown schematically in Figure 7. It consists of a hierarchy of layers containing units performing template matching and units performing a MAX operation (i.e., they select the maximum of a cell's inputs). A cascade of these two operations leads to C2 units (roughly corresponding to neurons at the mid level of the cortical processing hierarchy, areas V4 and/or the posterior ITC), which are tuned to complex features tolerant to changes in position and scale. The outputs of these units provide the inputs to the shape-tuned units (STUs), corresponding to view-tuned neurons in ITC (Logothetis et al., 1995; Booth and Rolls, 1998). Importantly, the responses of a shape-tuned model unit are completely determined by the shape of the preferred stimulus of the unit, with no explicit influence of category information.

We performed simulations using a population of 144 STUs, each tuned to a different stimulus from the cat/dog morph space (Knoblich et al., 2002). The 144 morphed animal stimuli were a subset of the stimuli used to train the monkey [i.e., chosen at random from the cat/dog morph space, excluding cats (dogs) with a dog (cat) component >40%]. This population of STUs was used to model a general stimulus representation consisting of neurons tuned to various shapes, which might then provide input to recognition task-specific neurons (such as for cat/dog categorization) in higher areas (Riesenhuber and Poggio, 2000). Each STU had a tuning width of $\sigma = 0.2$ and was connected to the 32 C2 afferents that were most strongly activated by its respective preferred stimulus, which produced neurons with realistic broadness of tuning.

We compared the tuning properties of model units with those of the ITC and the PFC neurons recorded from the two monkeys during the

cat/dog categorization task. For the comparison of model and experimental data, we restricted our analysis to the neurons that showed stimulus selectivity by an ANOVA ($p < 0.01$). Extension of the analysis to include all (responsive) neurons (relative to baseline, $p < 0.01$) added mainly untuned neurons with ROC area values close to 0.5.

Results

General properties

We recorded from a total of 968 neurons during DMC task performance, 525 PFC neurons (260 from monkey A, 265 from monkey B) and 443 ITC neurons (157 from monkey A, 286 from monkey B), from three hemispheres of two monkeys during 121 recording sessions (Fig. 3). Of these, 130 PFC and 117 ITC neurons were collected during 21 recording sessions with four to eight electrodes implanted simultaneously in both the PFC and the ITC of one monkey. Because the properties of neurons were similar between simultaneous versus single-area recording sessions and the two monkeys, we collapsed across them in reporting our results. During recordings, the monkeys' performance was >90% correct, even for sample stimuli that were close to the category boundary (e.g., the 60% morphs) and resembled stimuli from the other category (Freedman et al., 2002).

The activity of most of the neurons in both areas deviated from baseline rates during the sample presentation, memory delay, and/or test epoch (PFC: 423 of 525 or 81%; ITC: 345 of 443 or 78%; t test, $p < 0.01$; χ^2 test, PFC vs ITC; $p > 0.1$). We identified neurons that were stimulus selective (but not necessarily category selective) by using a one-way ANOVA with the 54 sample stimuli as a factor (at $p < 0.01$) on their average activity across the sample epoch and across the delay epoch. During the sample presentation, fewer PFC neurons (102 of 525 or 19%) than ITC neurons (213 of 443 or 48%) were stimulus selective (but not necessarily category selective) (χ^2 test, $p < 0.001$). By contrast, the opposite was true during the delay; the incidence of stimulus selectivity was greater in the PFC (58 of 525 or 11%) than ITC (29 of 443 or 7%; χ^2 test, $p < 0.05$). A similar pattern of results was obtained for neurons that were "category sensitive" (defined as those that showed a difference activity to all cat versus all dog samples; t test at $p < 0.01$). During the sample epoch, more ITC (119 of 443 or 27%) than PFC neurons (90 of 525 or 17%) were category sensitive (χ^2 test, $p < 0.001$), whereas the opposite was true during the delay epoch (39 of 443 or 9% of ITC neurons and 94 of 525 or 18% PFC neurons; χ^2 test, $p < 0.001$).

When the test stimulus was presented, monkeys needed to judge whether it was the same category as the sample. We found a greater incidence of ITC neurons (113 of 443 or 26%) than PFC neurons (57 of 525 or 11%) that showed an overall difference in activity to all test stimulus cats versus all dogs (t test, $p < 0.01$; χ^2 test, $p < 0.001$). Neural activity during the test epoch also reflected the previously seen sample. There was a greater incidence of neurons in the PFC ($n = 81/525$ or 15%) than ITC ($n = 48$ of 443 or 11%; χ^2 test, $p < 0.05$) that showed a difference in test epoch activity for all sample cats versus all dogs (t test, $p < 0.01$). During this epoch, a few neurons in each area (PFC: 33 of 525 or 6%; ITC: 19 of 443 or 4%) were stimulus selective for the sample (one-way ANOVA with the 54 sample stimuli as the factor; $p < 0.01$). We did not analyze stimulus selectivity for the test stimuli because they were a large number of randomly generated morphs (see Materials and Methods).

Comparison of the strength of category effects across areas

Figure 4 shows the activity of a PFC and an ITC neuron to stimuli as a function of distance from the category boundary. Both neu-

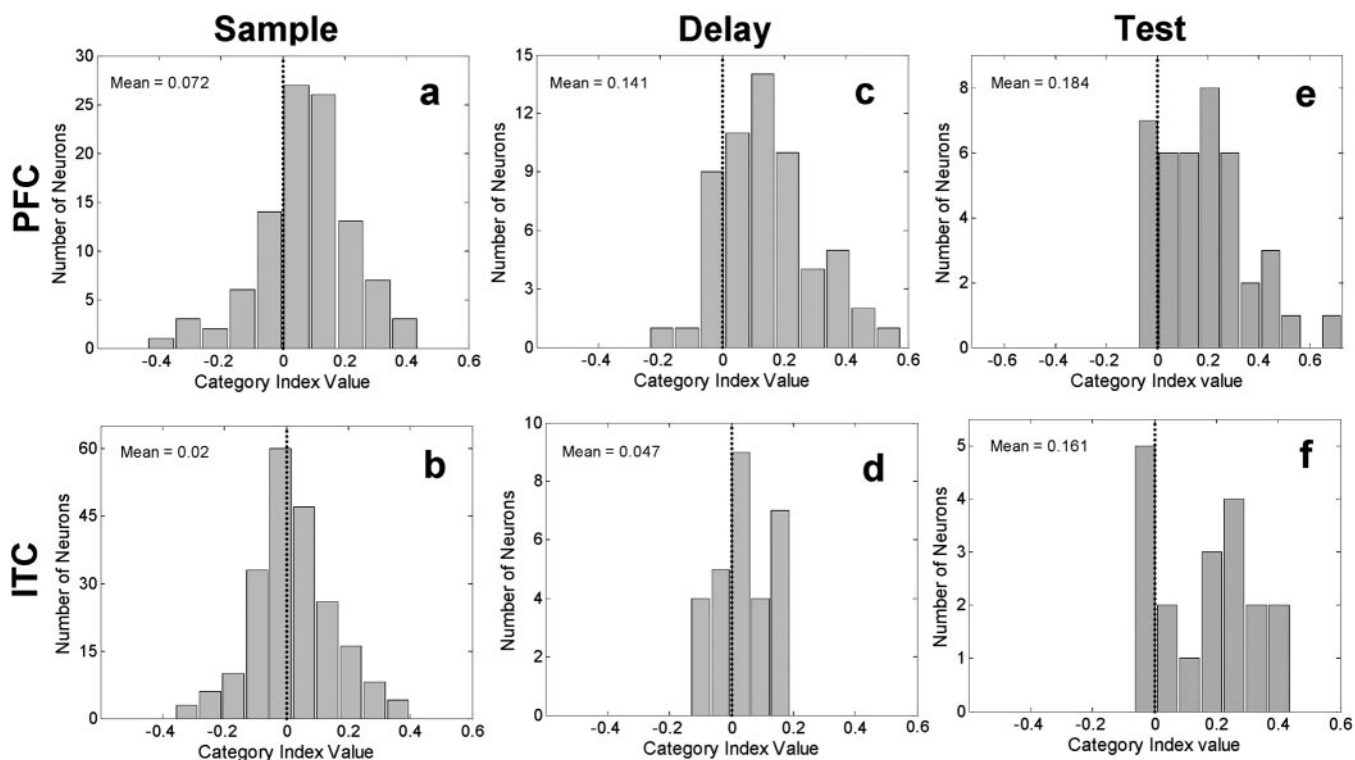


Figure 6. Distribution of PFC and ITC category indices for neurons that were stimulus selective during sample [PFC: $n = 102$; (a); ITC: $n = 213$ (b)], delay [PFC: $n = 58$ (c); ITC: $n = 29$ (d)], and test epochs [PFC: $n = 33$ (e); ITC: $n = 19$ (f)]. Indices were computed for stimulus-selective neurons separately for each epoch. Positive category index values indicate larger differences in activity to samples from different categories than from the same category. Negative values indicate larger differences to samples from the same category than between categories.

Table 1. Strength of the category-tuning index in PFC and ITC during sample, delay, and test epochs

	PFC	ITC
Number of stimulus-selective neurons (sample)	102 of 525 (19%)	213 of 443 (48%)
Category index value (mean)	0.072	0.025
<i>t</i> test on index value, PFC versus ITC	$p = 0.002$	
Number of stimulus-selective neurons (delay)	58 of 525 (11%)	29 of 443 (7%)
Category index value (mean)	0.141	0.047
<i>t</i> test on index value, PFC versus ITC	$p = 0.003$	
Number of stimulus-selective neurons (test)	33 of 525 (6%)	19 of 443 (4%)
Category index value (mean)	0.217	0.161
<i>t</i> test on index value, PFC versus ITC	$p = 0.238$	
One-way ANOVA <i>p</i> value		
Sample versus delay versus test indices:	$p = 1.6 \times 10^{-5}$	$p = 2.8 \times 10^{-5}$
Student–Newman–Keuls test for homogeneous subsets ($p < 0.05$):	Test > delay Test > sample Delay > sample	Test > delay Test > sample Sample = delay

rons convey some information about stimulus category. In the sample epoch (Fig. 4A) or the delay epoch (Fig. 4B), they showed different levels of activity to dogs versus cats, yet activity was similar to stimuli from the same category regardless of their degree of “dogness” or “catness.” A more detailed account of neuronal activity is shown in the color tile plots of Figures 4C,D and 5. Each tile shows the activity of a single neuron to each of the 42 sample stimuli that lay along the nine morph lines that crossed the category boundary (the lines that connect each cat prototype to each dog prototype in Fig. 1A). The PFC neurons shown in Figures 4 and 5 showed stronger category effects than the ITC neurons shown in Figures 4 and 5. As we will see below, these neurons are representative of their respective populations. The

PFC neurons showed sharper between-category differences and lower within-category variance, an observation confirmed by computing a category-tuning index and ROC values (see below and Figs. 4, 5).

To evaluate the strength of category effects, we computed a standard index that compared each neuron’s difference in activity between pairs of sample stimuli from different categories with its difference for pairs of sample stimuli from the same category (see Materials and Methods). Positive values indicate greater differences to samples between than within categories (i.e., stronger category effects), whereas negative values mean the opposite. We applied this analysis separately to each neuron and each epoch. We used all neurons that showed a significant difference in activity to all sample cats versus all dogs (*t* test at $p < 0.01$; sample epoch: 119 ITC and 90 PFC neurons; delay epoch: 39 ITC and 94 PFC neurons). For both the PFC and the ITC, the distribution of the index values was shifted significantly positive for the sample epoch (PFC: mean index value = 0.105, *t* test versus a mean of zero: $p = 2.8 \times 10^{-10}$; ITC: mean index value = 0.062, *t* test: $p = 4.6 \times 10^{-7}$) and delay epoch (PFC: mean index value = 0.136, *t* test: $p = 1.0 \times 10^{-11}$; ITC: mean index value = 0.081, *t* test: $p = 2.6 \times 10^{-6}$). However, the PFC distribution showed a stronger positive shift (i.e., category effect) than the ITC distribution for both the sample epoch (PFC mean index = 0.105; ITC mean = 0.062) and delay epoch (PFC mean index = 0.136; ITC mean = 0.081) that was significant at the $p < 0.05$ level (*t* test, sample epoch: $p = 0.023$; delay epoch: $p = 0.018$).

We next applied the same analysis to the entire population of neurons that showed any stimulus selectivity (one-way ANOVA with the 54 samples as a factor at $p < 0.01$; sample epoch: 102 PFC and 213 ITC neurons; delay epoch: 58 PFC and 29 ITC neurons). This revealed the same pattern of results. For both areas, the

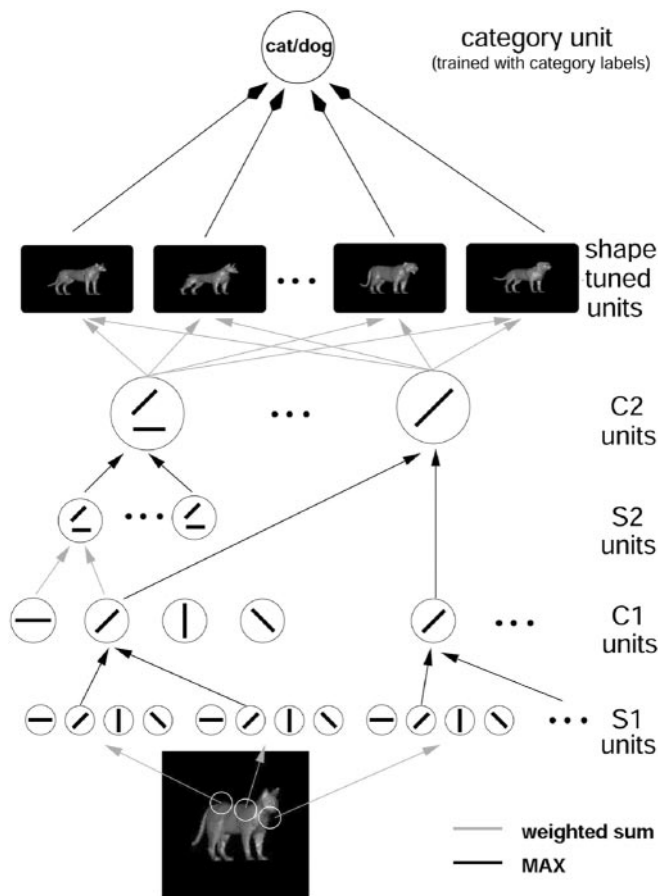


Figure 7. This figure shows a model of the recognition architecture in cortex (Riesenhuber and Poggio, 1999) that combines and extends several recent models (Fukushima, 1980; Perrett and Oram, 1993; Wallis and Rolls, 1997) and effectively summarizes many experimental findings (for review, see Riesenhuber and Poggio, 2002). It is a hierarchical extension of the classical paradigm (Hubel and Wiesel, 1962) of building complex cells from simple cells. The specific circuitry that we proposed consists of a hierarchy of layers with linear (“S” units, performing template matching, gray lines) and nonlinear units (“C” pooling units, performing a “MAX” operation, black lines, in which the response of the pooling neuron is determined by its maximally activated afferent). These two types of operations provide, respectively, pattern specificity, by combining simple features to build more complex ones, and invariance: to translation, by pooling over afferents tuned to the same feature at different positions, and scale (data not shown), by pooling over afferents tuned to the same feature at different scales. Shape-tuned model units (STU) exhibit tuning to complex shapes but are tolerant to scaling and translation of their preferred shape, like view-tuned neurons found in ITC (cf. Logothetis et al., 1995). STUs can then serve as input to task modules located farther downstream, e.g., in PFC, that perform different visual tasks such as object categorization. These modules consist of the same generic learning process but are trained to perform different tasks. For more information, see <http://www.ai.mit.edu/projects/cbd/hmax>.

mean of the category index distribution was significantly above zero for the sample epoch (PFC: mean index = 0.072, *t* test versus a mean of zero: $p = 6.2 \times 10^{-6}$; ITC: mean index = 0.025, *t* test, $p = 0.0028$) and the delay epoch (PFC: mean index = 0.141, 2.9×10^{-3} ; ITC: mean index = 0.047, $p = 0.0085$). Once again, the PFC index values showed a stronger positive shift than the ITC index values for the sample epoch (PFC mean index = 0.072; ITC = 0.025; *t* test, $p = 0.002$) and delay epoch (PFC mean index = 0.141; ITC mean index = 0.047; *t* test, $p = 0.003$). Thus, it seems that category effects are generally stronger in the PFC than ITC.

To solve this task, the monkeys needed to compare the test stimulus category with the memory of the sample’s category. To analyze category tuning (for the sample stimulus) during the test

stimulus epoch, we computed the category index for neurons that were significantly sample-stimulus selective during the test epoch (one-way ANOVA with the 54 sample stimuli as a factor; $p < 0.01$). Category index values were not computed for the test stimulus category because they were a large set of randomly generated morphs (see Materials and Methods). We used activity up to, but not including, the behavioral response (see Materials and Methods). The sample category indices for the test epoch were again shifted toward positive values for both the PFC ($n = 33$ neurons; index = 0.217; *t* test, $p = 2.0 \times 10^{-8}$) and ITC neurons ($n = 19$ neurons; index = 0.161; *t* test, $p = 2.4 \times 10^{-4}$). During the test epoch, the PFC and the ITC category index distributions were not significantly different from one another (*t* test, $p = 0.24$). Figure 6 shows distributions of the sample category index values for all stimulus-selective neurons during the sample (*a, b*), delay (*c, d*), and test (*e, f*) epochs.

We also compared the sample category index values across the three epochs. We used the index values for the neurons that were stimulus selective during the corresponding epoch (shown in Fig. 8 and discussed above). Mean index values as well as results from a one-way ANOVA and *post hoc* tests (Student–Newman–Keuls at $p < 0.05$) are shown in Table 1. They indicated the following pattern of results. In both the PFC and the ITC, the effect of the sample category was higher in the test epoch than during the sample or delay, and in the PFC, category tuning was stronger during the delay than the sample epoch. A similar pattern of results was obtained with a detailed analysis of the time course of category signals. This will be discussed below.

Comparison between neurons and model units

In principle, highly selective responses to individual stimuli could produce positive category index values. Thus, we needed to know how the results obtained above compare with that expected from stimulus selectivity alone. So, we calculated the category index using 144 units from a computational model of object recognition (Riesenhuber and Poggio, 1999; Knoblich and Freedman, 2002) (Fig. 7) that were trained to be selective for individual shapes but not carry any explicit category information. The mean index value from the model STUs was 0.03, a value that did not significantly differ from that obtained from the ITC stimulus-selective neurons [sample epoch: $n = 213$ ITC neurons, mean index value = 0.025 (Fig. 6D); *t* test vs STU, $p = 0.61$; delay epoch: $n = 29$ ITC neurons, mean index value = 0.047 (Fig. 6E); *t* test vs STU, $p = 0.39$). During the test epoch, however, sample category index values from the 19 sample stimulus-selective ITC neurons were significantly greater than the STU values [ITC mean index value = 0.161 (Fig. 6F); *t* test vs STU, $p = 1.7 \times 10^{-5}$]. This suggests that through most of the trial, the influence of category in the ITC neural population is not different from that expected from a population of stimulus-selective (and not category-selective) neurons.

By contrast, the comparison between the STU and PFC neurons suggested that the PFC neuronal population showed an influence of sample category throughout the trial. Category index values for stimulus-selective PFC neurons were shifted more positively than the STU values for the sample epoch [$n = 102$ PFC neurons; mean index value = 0.072 (Fig. 6A); *t* test vs STU, $p = 5.2 \times 10^{-4}$], delay epoch [$n = 58$ PFC neurons; mean index value = 0.141 (Fig. 6B); *t* test vs STU, $p = 2.7 \times 10^{-8}$], and test epoch ($n = 33$ PFC neurons; mean index value = 0.217 (Fig. 6C); *t* test vs STU, $p = 1.8 \times 10^{-12}$].

Time course and latency of activity

Many of the results reported above were reflected in a detailed analysis of the time course of category effects. Figure 8 shows a histogram of the average population activity for the 141 ITC (Fig. 8*A*) and 156 PFC (Fig. 8*B*) neurons that showed a significant difference to all sample cats versus all dogs during the sample and/or delay epochs (*t* test, $p < 0.01$). There are several noteworthy differences between PFC and the ITC average activity: ITC activity showed a sharper phasic burst at sample onset than PFC activity and during the delay, PFC average activity remained above baseline (fixation epoch) level, but ITC average activity returned to baseline.

We examined the latencies of neuronal responses to onset of the sample stimulus using the population of ITC and PFC neurons that showed a difference to all cats versus all dogs during the sample epoch (*t* test, $p < 0.01$; ITC, $n = 119$ neurons; PFC, $n = 90$ neurons). For each neuron, we calculated the time at which neural activity began to depart from baseline levels (see Materials and Methods). Distributions of these values are shown in Figure 8, *C* and *D*. On average, ITC neurons (mean latency = 103 msec) began responding sooner than PFC neurons (mean latency = 196 msec; Wilcoxon rank sum test, $p = 4.4 \times 10^{-4}$).

This difference in response latency between the areas can also be seen in the population average histograms shown in Figure 8, *A* and *B*. We calculated the latency for this average activity to begin to depart from baseline by determining the time bin at which it exceeded 3 SDs of the variability across the baseline time bins (see Materials and Methods) for three consecutive 1 msec bins. This occurred at 88 msec for the ITC population histogram and 124 msec for the PFC histogram.

Next, we examined the time course of category-related signals by using a sliding ROC analysis. For each neuron, we calculated ROC values of the difference in activity to all cat samples versus all dog samples at 10 msec time steps (Green and Swets, 1966; Tolhurst et al., 1983; Freedman et al., 2002) (see Materials and Methods). Figure 9 shows a plot of these values for all neurons that showed a significant difference to all cats versus all dogs in their average sample and delay epoch activity (*t* test, $p < 0.01$; 156 PFC neurons; 141 ITC neurons). In the ITC, (Fig. 9*A*), information about the sample category was relatively transient. ROC values showed a sharp increase after sample onset and offset, but ITC values were relatively low during the memory delay, only to peak again at test stimulus onset. In the PFC, ROC values also showed a sharp increase after sample and test stimulus onset, but in contrast to the ITC, the PFC values remained elevated throughout the remainder of the trial (Fig. 9*B*).

For each neuron, we calculated the latency for category signals to appear by determining at which time step its ROC values first rose above baseline levels (see Materials and Methods). Distributions of the latencies are shown in Figure 9, *C* and *D*, respectively. The mean latency values were 127 msec for the ITC (Fig. 9*C*) and 171 msec for the PFC (Fig. 9*D*). The difference in these latency




Figure 8. Average histogram of ITC and PFC neuronal activity. *A*, The average normalized activity is shown for the population of 141 category-sensitive ITC neurons ($p < 0.01$). *B*, The average histogram across 156 category-sensitive PFC neurons. *C*, Distribution of neural response latencies to sample onset for 119 ITC sample-epoch category-sensitive neurons. *D*, Distribution of neural response latencies to sample onset for 90 PFC sample-epoch category-sensitive neurons.

values did not reach significance according to a Wilcoxon rank sum test ($p = 0.13$).

Category match/non-match effects

When the test stimulus appeared, the monkeys needed to judge whether its category matched that of the sample category. To discern whether the category match/non-match status of the test stimulus was reflected in neural activity, we computed, for each neuron, a two-way ANOVA on its average activity across the test stimulus epoch (i.e., before the behavioral response; see Materials and Methods). The factors were the test stimulus match/non-match status and test stimulus category. The results are summarized in Table 2.

A similar proportion of PFC and ITC neurons showed a significant effect of one or the other factor, or both (PFC: 29% or 152 of 525; ITC: 34% or 151 of 443; any main effect or a significant interaction, $p < 0.01$). However, the patterns of effects differed between the areas. In the PFC, a relatively small proportion of neurons (20% or 31 of 152) reflected the test stimulus category (main effect, $p < 0.01$; with no interaction, $p > 0.01$). More PFC neurons reflected its match/non-match status: ~32% (48 of 152) of PFC neurons showed a main effect of match/non-match (with no interaction) and 51% (77 of 152) showed a significant interaction between category and match/non-match factors. Interactions reflect different effects of match/non-match between the categories or category selectivity for the sample stimulus. We have previously reported finding such effects in the PFC (Freedman et al., 2002). By contrast, fewer ITC neurons showed category match/non-match effects (main effect with no interaction: 7% or 10 of 151, interaction: 36% or 55 of 151). Instead, the modal group of ITC neurons (59%, $n = 89$ of 151) reflected the




Figure 9. Time course of category sensitivity. The activity of 141 category-sensitive ITC (*A*) and 156 category-sensitive PFC (*B*) neurons was computed using a sliding ROC analysis at 10 msec steps (see Materials and Methods). ROC values were sorted independently for each time bin. Higher ROC values indicate better discrimination between categories (0.5 = no discrimination, 1.0 = perfect discrimination). The black lines correspond (from right to left) to sample onset, offset, and test stimulus onset. Time is aligned to the midpoint of the sliding window (i.e., values at time = 0 indicate the ROC values during the –100 to 100 msec time epoch relative to sample onset). *C* and *D* show the distribution of the latencies for neurons to begin to distinguish between categories (as computed by ROC) for 119 ITC (*C*) and 90 PFC (*D*) neurons that were category sensitive during the sample epoch.

Table 2. Incidence of choice-period selectivity

	PFC	ITC
Number of neurons selective for:		
Match/non-match	48 of 152 (32%)	10 of 151 (7%)
Choice category	31 of 152 (20%)	89 of 151 (59%)
Match × category	77 of 152 (51%)	55 of 151 (36%)

PFC: $n = 152$ of 525 have significant effect of either factor or interaction ($p < 0.01$); ITC: $n = 151$ of 443 have significant effect of either factor or interaction ($p < 0.01$).

test stimulus category. The difference in proportion of PFC and ITC neurons showing match/non-match or category effects were both significant (χ^2 test, $p < 0.001$). Examples of a PFC neuron that showed a match/non-match effect and an ITC neuron that reflected the test stimulus category are shown in Figure 10, *A* and *B*, respectively. Thus, it seemed that during the test epoch, ITC

neurons were more involved in the visual analysis of the currently viewed (test) stimulus, whereas there was a stronger representation of behavioral factors (match/non-match status and selection of the corresponding behavioral response) in the PFC.

In the PFC, match/non-match effects were evenly split between neurons showing more activity to matches and neurons showing more activity to non-matches. Of the PFC neurons with match/non-matches effects (including those with a significant interaction), 49% (39 of 79) showed higher activity to matches than non-matches and 51% (40 of 79) showed the opposite effect. Among match/non-match selective ITC neurons (including those that had a significant interaction), 63% (12 of 19) showed more activity to matches and 37% (7 of 19) showed more activity to non-matches. These proportions were not significantly different from 0.5 according to a binomial test (PFC: $p = 0.999$; ITC: $p = 0.359$).




Figure 10. Activity of a single PFC (*A*) and ITC (*B*) neuron during the test epoch. Neuronal activity is grouped according to the category of the test stimulus and whether the category of the test stimulus matched the category of the sample. *A*, A PFC neuron that showed a match/non-match effect; it responded preferentially to when the test stimulus did not match the sample category. *B*, An example of a category-sensitive ITC neuron that showed more activity to both the sample and test stimuli when they were dogs as opposed to cats.

To examine the latency of match/non-match effects, we focused on neurons that were match/non-match selective (main effect, no interaction) during the test epoch (see Materials and Methods and above; $n = 48$ PFC; $n = 10$ ITC). For each neuron, we computed a sliding ROC analysis (100 msec width of analysis window, 10 msec steps) that compared activity on all match trials versus all non-match trials. An average histogram of sliding ROC values for the PFC and ITC is shown in Figure 11. For each neuron, the latency was defined as the point at which ROC values exceeded the mean ROC value before onset of the test stimulus by 3 SDs for three consecutive time epochs (see Materials and Methods). The mean latencies for the PFC (130 msec) and ITC (140 msec) were not significantly different from one another (Wilcoxon rank sum test, $p = 0.718$).




Figure 11. Time course of match/non-match selectivity. The results of a sliding ROC analysis that compared activity on all match trials versus all non-match trials. This figure shows the average ROC values for 48 PFC (solid line) and 10 ITC (dotted line) neurons that showed a match/non-match effect (see Table 2 and Materials and Methods). The average latencies of match/non-match selectivity in the PFC (130 msec) and ITC (140 msec) were not significantly different from one another.

Discussion

We trained monkeys to perform a visual categorization task and found that neurons in the PFC and the ITC shared many properties but also exhibited differences that may provide insight into their respective roles. Behaviorally relevant factors such as stimulus category and match/non-match status (or response selection) were encoded more strongly in the PFC than the ITC. Furthermore, the time course of effects also suggested that the ITC was more involved in rapid visual analysis. Category-related signals peaked in the ITC around the time a stimulus was being viewed, whereas the PFC showed more robust maintenance of category-related signals across a memory delay. Simply put, the PFC seems more “behavioral,” whereas the ITC seems more “visual.” We elaborate below.

First and foremost, category effects tended to be stronger in the PFC than the ITC. In fact, except for activity near the end of the trial (during the test stimulus epoch), the distribution of category effects across the population of stimulus-selective ITC neurons did not differ from that obtained from stimulus-selective model units that conveyed no category information per se. By contrast, the distribution of category effects for PFC neurons was significantly greater than that of the model units and ITC neurons. This does not mean that category membership was not reflected in any ITC neurons; a few ITC neurons did show strong category effects, particularly during the test epoch. However, our results indicate that during most of the trial, there was a stronger and more explicit representation of visual category in the PFC than ITC.

This presumably reflects a greater emphasis on analysis of features of individual stimuli in the ITC versus the PFC. Such a role for the ITC during categorization was recently suggested by a study by Sigala and colleagues (2002). In that study, monkeys grouped face or fish stimuli into two categories on the basis of several “diagnostic features” (such as the distance between the eyes and the height of the eyes on the face) that were relevant for categorization. After training, ITC neurons showed enhanced selectivity for the diagnostic features relative to other stimulus fea-

tures that were unrelated to category membership. This, along with our results, suggests that although the emphasis of diagnostic features may occur in the ITC, the process of combining those features into an explicit representation of category occurs more at the level of the PFC than the ITC. Of course, this does not mean that the PFC is never sensitive to individual stimuli; PFC neurons can be highly sensitive to small changes in stimuli in both visual tasks and non-visual tasks, particularly when such fine details are task relevant (Romo et al., 1999; Rainer and Miller, 2000).

We also found more robust maintenance of category signals across the memory delay in the PFC than the ITC. In fact, across the ITC population, average activity during the memory delay returned to baseline levels, whereas it remained above baseline in the PFC. This observation of more robust working memory-related activity in the PFC versus ITC is consistent with previous studies (Miller et al., 1996). Interestingly, the strength of sample category-related signals peaked in the PFC after the memory delay, during test stimulus presentation. At around the same time, category-related signals reappeared and peaked in the ITC after having essentially disappeared during the memory delay (Fig. 9) (ROC analysis). Thus, the reappearance of category-related signals in the ITC might reflect a “top-down” feedback from the PFC, where category information was being held in working memory. This is consistent with a demonstration by Tomita et al. (1999) of the importance of feedback from the PFC to the ITC during a visual recall task.

The notion that the PFC is more involved in encoding stimuli in terms of behavioral relevance is further supported by our observation of a greater incidence of match/non-match effects in the PFC than the ITC. These effects could reflect the process of evaluating the match/non-match status of stimuli or the selection of the corresponding behavioral response (hold versus release). Previous studies of the ITC during an identity (not category) match-to-sample task that used the same behavioral response found a greater incidence of ITC match/non-match effects than is reported here (Miller et al., 1991; Miller and Desimone, 1994). It may be that identity matching engages more ITC mechanisms than category matching; category matching may depend more on PFC mechanisms because the PFC has a more explicit representation of category than the ITC. Our results are also compatible with that of Romo and colleagues (Hernandez et al., 2000, 2002; Romo et al., 2002; Romo and Salinas, 2003), who have compared neuronal activity in different cortical areas while monkeys performed a tactile discrimination task. They found that higher-order association cortex plays a greater role in decision-making and motor planning, whereas primary and secondary somatosensory cortices show a greater emphasis on the perceptual aspects of the task.

The presence of robust category and behavioral signals in the PFC task fits well with its putative role in guiding goal-directed behaviors (Grafman, 1994; Duncan, 1996; Miller, 2000; Miller and Cohen, 2001) and its emphasis on task-relevant information (Rainer et al., 1998) and in the representation of the behavioral context of stimuli (White and Wise, 1999; Asaad et al., 2000). Categories are defined by their functional relevance and thus might be expected to be strongly represented in the PFC, a brain area that is thought to mediate the functions needed to transform sensory information into voluntary actions, such as the integration of temporally separated events (Fuster, 1990; Fuster et al., 2000), the acquisition and representation of behavior-guiding rules (Asaad et al., 1998; White and Wise, 1999; Wallis et al., 2001), and visuomotor decisions (Bichot and Schall, 1999; Kim and Shadlen, 1999). The stronger representation of category in

PFC activity is also consistent with theoretical models positing that categorization arises as appropriately weighted inputs from stimulus-selective ITC neurons converge on PFC neurons specialized for encoding behaviorally relevant variables (Riesenhuber and Poggio, 2000). Whether this division of labor is true for all visual categories (including those that are highly familiar and innate) (Kanwisher, 2000) remains to be determined. It is worth noting, however, that this pattern of results is not likely to be limited to just cats and dogs. It presumably extends to many complex visual stimuli. In previous work using these stimuli, we demonstrated that category effects in the PFC were acquired during training and therefore not dependent on an innate representation of cat or dog (Freedman et al., 2001; Op de Beeck et al., 2001; Thomas et al., 2001). In principle, a similar division of labor may occur with any complex visual stimuli for which category membership must be learned.

In sum, our results suggest that the PFC and the ITC play overlapping but distinct roles in visual categorization. Although our discussion has emphasized their putative unique contributions, it is important to stress in closing that many neurons in both areas were engaged by this task and differences between these areas is a matter of degree rather than a strict division of labor. Categorization is likely to depend on interactions between these and other structures.

References

- Asaad WF, Rainer G, Miller EK (1998) Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21:1399–1407.
- Asaad WF, Rainer G, Miller EK (2000) Task-specific neural activity in the primate prefrontal cortex. *J Neurophysiol* 84:451–459.
- Baker CI, Behrmann M, Olson CR (2002) Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nat Neurosci* 5:1210–1215.
- Bichot NP, Schall JD (1999) Effects of similarity and history on neural mechanisms of visual selection. *Nat Neurosci* 2:549–554.
- Booth MC, Rolls ET (1998) View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cereb Cortex* 8:510–523.
- Bruce C, Desimone R, Gross CG (1981) Visual properties of neurons in a polysensory area in superior temporal sulcus in the macaque. *J Neurophysiol* 46:369–384.
- Desimone R, Albright TD, Gross CG, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4:2051–2062.
- Duncan J (1996) Intelligence and the frontal lobe: the organization of goal-directed behavior. *Cognit Psychol* 30:257–303.
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291:312–316.
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2002) Visual categorization and the primate prefrontal cortex: neurophysiology and Behavior. *J Neurophysiol* 88:914–928.
- Fukushima K (1980) Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol Cybern* 36:193–202.
- Fuster JM (1990) Prefrontal cortex and the bridging of temporal gaps in the perception-action cycle. *Ann NY Acad Sci* 608:318–329.
- Fuster JM, Bodner M, Kroger JK (2000) Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature* 405:347–351.
- Grafman J (1994) Alternative frameworks for the conceptualization of prefrontal functions. In: *Handbook of neuropsychology* (Boller F, Grafman J, eds), pp 187–200. Amsterdam: Elsevier.
- Green DM, Swets JA (1966) Signal detection theory and psychophysics. New York: Wiley.
- Gross CG (1973) Visual functions of inferotemporal cortex. In: *Handbook of sensory physiology*, Vol VII/3B (Aultrium H, Jung R, Lowenstein W, Mckay D, Teuber H-L, eds), pp 451–482. Berlin: Springer.
- Hernandez A, Zainos A, Romo R (2000) Neuronal correlates of a sensory

- discrimination in the somatosensory cortex. *Proc Natl Acad Sci USA* 97:6191–6196.
- Hernandez A, Zainos A, Romo R (2002) Temporal evolution of a decision-making process in medial premotor cortex. *Neuron* 33:959–972.
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol (Lond)* 160:106–154.
- Kanwisher N (2000) Domain specificity in face perception. *Nat Neurosci* 3:759–763.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Knoblich U, Freedman DJ, Riesenhuber M (2002) Categorization in IT and PFC: model and experiments. *AI Memo 2002–007*. Cambridge, MA: MIT AI Laboratory.
- Kobatake E, Wang G, Tanaka K (1998) Effects of shape discrimination training on the selectivity of inferotemporal cells in adult monkeys. *J Neurophysiol* 80:324–330.
- Logothetis NK, Sheinberg DL (1996) Visual object recognition. *Annu Rev Neurosci* 19:577–621.
- Logothetis NK, Pauls J, Poggio T (1995) Shape representation in the inferior temporal cortex of monkeys. *Curr Biol* 5:552–563.
- Miller EK (2000) The prefrontal cortex and cognitive control. *Nat Rev Neurosci* 1:59–65.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
- Miller EK, Desimone R (1994) Parallel Neuronal mechanisms for short-term memory. *Science* 263:520–522.
- Miller EK, Li L, Desimone R (1991) A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* 254:1377–1379.
- Miller EK, Li L, Desimone R (1993) Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13:1460–1478.
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154–5167.
- Miyashita Y (1988) Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature* 335:817–820.
- Nieder A, Freedman DJ, Miller EK (2002) Representation of the quantity of visual items in the primate prefrontal cortex. *Science* 297:1708–1711.
- Op de Beeck H, Wagemans J, Vogels R (2001) Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nat Neurosci* 4:1244–1252.
- Paxinos G, Huang XF, Toga AW (2000) The rhesus monkey brain in stereotaxic coordinates. San Diego: Academic.
- Perrett D, Oram M (1993) Neurophysiology of shape processing. *Imag Vis Comput* 11:317–333.
- Perrett DI, Rolls ET, Caan W (1982) Visual neurones responsive to faces in the monkey temporal cortex. *Exp Brain Res* 47:329–342.
- Rainer G, Miller EK (2000) Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron* 27:179–189.
- Rainer G, Asaad WF, Miller EK (1998) Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393:577–579.
- Riesenhuber M, Poggio T (1999) Hierarchical models of object recognition in cortex. *Nat Neurosci* 2:1019–1025.
- Riesenhuber M, Poggio T (2000) Models of object recognition. *Nat Neurosci [Suppl]* 3:1199–1204.
- Riesenhuber M, Poggio T (2002) Neural mechanisms of object recognition. *Curr Opin Neurobiol* 12:162–168.
- Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4:203–218.
- Romo R, Brody CD, Hernandez A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473.
- Romo R, Hernandez A, Zainos A, Lemus L, Brody CD (2002) Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat Neurosci* 5:1217–1225.
- Shelton C (2000) Morphable surface models. *Int J Comp Vis* 38:75–91.
- Sigala N, Logothetis NK (2002) Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature* 415:318–320.
- Tanaka K (1996) Inferotemporal cortex and object vision. *Annu Rev Neurosci* 19:109–139.
- Thomas E, Van Hulle MM, Vogels R (2001) Encoding of categories by noncategory-specific neurons in the inferior temporal cortex. *J Cognit Neurosci* 13:190–200.
- Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23:775–785.
- Tomita H, Ohbayashi M, Nakahara K, Hasegawa I, Miyashita Y (1999) Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature* 401:699–703.
- Ungerleider LG, Gaffan D, Pelak VS (1989) Projections from inferior temporal cortex to prefrontal cortex via the uncinate fascicle in rhesus monkeys. *Exp Brain Res* 76:473–484.
- Vogels R (1999) Categorization of complex visual images by rhesus monkeys. *Eur J Neurosci* 11:1223–1238.
- Wallis G, Rolls ET (1997) Invariant face and object recognition in the visual system. *Prog Neurobiol* 51:167–194.
- Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956.
- Webster MJ, Bachevalier J, Ungerleider LG (1994) Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. *Cereb Cortex* 4:470–483.
- White IM, Wise SP (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res* 126:315–335.