Comparison of Primate Prefrontal and Premotor Cortex Neuronal Activity during Visual Categorization

Jason A. Cromer, Jefferson E. Roy, Timothy J. Buschman, and Earl K. Miller

Abstract

■ Previous work has shown that neurons in the PFC show selectivity for learned categorical groupings. In contrast, brain regions lower in the visual hierarchy, such as inferior temporal cortex, do not seem to favor category information over information about physical appearance. However, the role of premotor cortex (PMC) in categorization has not been studied, despite evidence that PMC is strongly engaged by well-learned tasks and reflects learned rules. Here, we directly compare PFC neurons with PMC neurons during visual categorization. Unlike PFC neurons, relatively few PMC neurons distinguished between categories of visual images during a delayed match-to-category task. However, despite the lack of category information in the PMC, more than half of the neurons in both PFC and PMC reflected whether the category of a test image did or did not match the category of a sample image (i.e., had match information). Thus, PFC neurons represented all variables required to solve the cognitive problem, whereas PMC neurons instead represented only the final decision variable that drove the appropriate motor action required to obtain a reward. This dichotomy fits well with PFC's hypothesized role in learning arbitrary information and directing behavior as well as the PMC's role in motor planning. ■

INTRODUCTION

Categorization is a fundamental cognitive function that allows us to group items by utility, providing a basis for abstract behavior (Seger & Miller, 2010; Miller, Nieder, Freedman, & Wallis, 2003). This functionality is disrupted in several neuropsychiatric disorders including schizophrenia and autism (Kuperberg, West, Lakshmanan, & Goff, 2008; Scherf, Luna, Kimchi, Minshew, & Behrmann, 2008; Bolte, Holtmann, Poustka, Scheurich, & Schmidt, 2007; Uhlhaas & Mishara, 2007). Categorization is likely to engage many brain areas. The specific areas engaged probably depend on the nature of categorization (e.g., modality, complexity, and learned vs. innate; Seger & Miller, 2010). Our laboratory has used a task that requires monkeys to determine whether two images, separated by a brief delay, belong to the same category or not (Freedman, Riesenhuber, Poggio, & Miller, 2001). This engages cognitive functions like STM and decision-making for which PFC is important. Prior studies have revealed that as many as 30–40% of randomly selected lateral PFC neurons showed activity that reflected the learned category groups of stimuli at the expense of their exact physical appearance (Cromer, Roy, & Miller, 2010; Roy, Riesenhuber, Poggio, & Miller, 2010; Freedman, Riesenhuber, Poggio, & Miller, 2002; Freedman et al., 2001). This differs from lower visual structures, such as inferior temporal cortex, which initially encodes physical appearance rather than category information (Freedman & Miller, 2008; Freedman, Riesenhuber, Poggio, & Miller, 2003).

Here, we directly compare the lateral PFC to the premotor cortex (PMC) during visual categorization. Previous work has found that parts of PMC are more involved in cognitive than motor processes (Picard & Strick, 2001) and that well-practiced tasks can more strongly engage the PMC than the PFC in humans (Boettiger & D'Esposito, 2005; Della-Maggiore & McIntosh, 2005; Raichle et al., 1994). Furthermore, in a different experiment employing an abstract rule task, we found that information about abstract rules ("same" vs. "different") was reflected more strongly and earlier in the PMC than in the PFC in well-practiced monkeys (Muhammad, Wallis, & Miller, 2006). Here, we aimed to determine whether visual categories were also reflected in PMC activity.

METHODS

Data Collection

Data were collected from the same two macaque monkeys (*Macaca mulatta*) that were reported in our previous study (Cromer et al., 2010). These animals were cared for in accordance with National Institutes of Health guidelines and the policies of the Massachusetts Institute of Technology Committee on Animal Care. The monkeys' eye movements were recorded using an infrared eye tracking system (Iscan, Burlington, MA) at a sampling rate of

Massachusetts Institute of Technology

240 Hz. Neural recordings were made using extracellular electrodes (FHC, Inc., Bowdoin, ME) that were lowered through the dura each day using in-house screw microdrives. Electrodes were either driven independently or in pairs, and up to 32 electrodes were used during each recording session. Recording wells were positioned based on MRIs over the lateral PFC and the dorsal, rostral PMC—the area receiving input from 9/46 and corresponding to 6DR in Petrides and Pandya (2006). Electrodes were lowered each day into the cortical cell layer (i.e., when neurons or hash was easily identifiable) and allowed to settle, after which they were adjusted to obtain neurons on one or both electrodes on each microdrive. Importantly, no prescreening of neurons took place (i.e., we did not first search for neurons that were task selective and then only record those neurons). This resulted in an unbiased sample of neurons from both brain regions. Spike waveforms were amplified, digitized, and then stored for subsequent off-line sorting using principal component analysis (Offline Sorter, Plexon, Inc., Dallas, TX). We included all well-isolated neurons that were held for a minimum of 500 correct trials in our analysis. This resulted in 455 PFC neurons (358 from monkey "ti" and 97 from monkey "lu") and 185 premotor neurons (185 from monkey "ti").

Visual Stimuli

Two independent sets of stimuli, cars and animals, were used for this study (Figure 1A). Within each set, there was a single category distinction. The first category set (cars) was categorized as either Sports cars or Sedans, whereas the second set (animals) consisted of a Cats vs. Dogs category distinction. Each category set was made up of four prototype images (two prototypes from each category, as seen in Figure 1A) and morphs between those prototypes (Figure 1B). Morph images were generated as in previous studies, ensuring that multiple features were smoothly morphed without the sudden appearance of any feature and so that all images within a category set had identical color, shading, orientation, and scale (Cromer et al., 2010; Roy et al., 2010; Jiang et al., 2007; Freedman, Riesenhuber, Poggio, & Miller, 2006; Freedman et al., 2001, 2002, 2003; Shelton, 2000; Beymer & Poggio, 1996). Each category set (Cars and Animals) had a fixed category boundary at 50%, meaning that an image was considered a member of a category if it contained more than a 50% contribution from a prototype in that category. During training, the image sets consisted of hundreds of morph images to prevent the monkeys from memorizing specific images. Recording sessions utilized only morphs at the 100:0, 80:20, 60:40, 40:60, 20:80, and 0:100 combination levels for each possible morph line to present each sample image frequently enough to obtain sufficient data for statistical analysis. It also allowed us to determine if neurons were more selective for specific morphs versus their category membership. This process was repeated for both category sets, so that the final recording set contained 40 possible sample images. An example of a morph line depicting the transition from Sports Car prototype a2 to Sedan prototype b1 with the morph steps used during recordings is shown in Figure 1B.

Behavioral Delayed Match-to-Category Task

As in previous studies of categorization (Cromer et al., 2010; Freedman et al., 2001, 2002, 2003), we employed a delayed match-to-category test during which monkeys had to categorize a sample stimulus and then determine whether a test stimulus matched (i.e., came from the same category) the sample's category (Figure 1C). Monkeys initiated the task by grabbing a response bar, which caused the onset of a white fixation square to which monkeys had to maintain fixation within 2° during the course of the entire trial. After a 1000-msec baseline period of fixation, one of the sample images (40 possibilities—20 from either category set) was presented for 600 msec. At this time, the only information available to the monkeys was the sample image itself and the category to which the monkey assigned that image. Monkeys were required to hold in mind that category information throughout a subsequent 1000-msec delay period, after which time a test image was presented. This test image could be a category match (an image from the same category as the sample) or a category nonmatch (an image from the opposite category as the sample). For example, in Figure 1C, the test image that was a sedan is a category match to the sample image that was also a sedan, whereas the sports car test image is a category nonmatch. Monkeys indicated whether the test image was a category match to the sample by releasing the response bar within 600 msec of the test image presentation. However, if the test image was a nonmatch, they were to continue to hold the response bar. On nonmatch trials, the test image was turned off after 600 msec, and the monkeys waited through a second 600-msec memory delay until the presentation of a second test image that was always a category match to the sample image when they could release the response bar. Monkeys were immediately rewarded with several drops of apple juice for correctly releasing the response bar. If the monkeys incorrectly released the response bar or made no response, the trial was immediately aborted and an extended "timeout" delay occurred before the start of the next trial.

This paradigm required the monkeys to make a response on every trial (either releasing the bar during the first or second test image), and they would receive a reward on every correct trial (reward amounts were always kept constant throughout each recording session). Furthermore, the monkeys' motor responses were dissociated from the category of the sample stimulus (i.e., knowing the category of the sample stimulus during the sample and delay phases of the trial did not tell the monkey whether they would release or hold the bar during the test

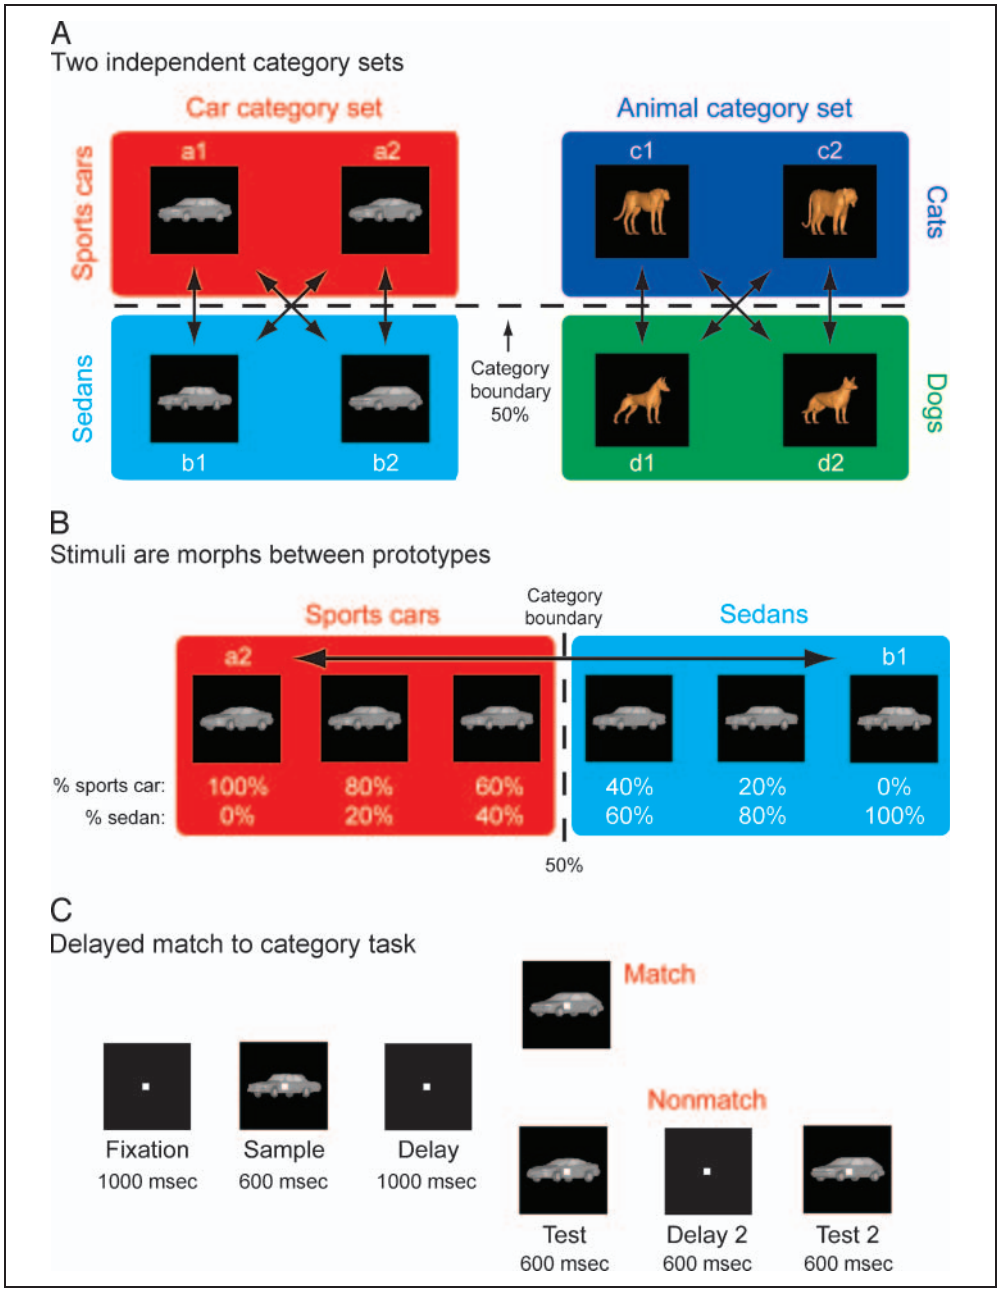

Figure 1. Stimulus set and behavioral task. (A) Visual stimuli were generated for two independent category sets, cars and animals. The car set was divided into "Sports Cars" versus "Sedans" categories, whereas the animal set had "Cats" and "Dogs" categories. Each category set was formed on the basis of two prototype images (shown) and morphs between those images along the four depicted morph lines (arrows) between all combinations of the prototypes within each set. (B) Morphing allowed parameterization of sample images. An example morph line between Sports Car prototype a2 and Sedan prototype b1 displays images at the morph steps used for recording. Intermediate images were a mix of the two prototypes. Those images comprised of greater than 50% of one category (marked by "Category Boundary") were to be classified as members of that category. Note how the 60%/40% morphs (nearest to "Category Boundary") are closer in physical similarity to each other than they are to the prototypes, yet they are categorized differently because they are on opposite sides of the category boundary. (C) The delayed match to category task required monkeys to respond to whether a test stimulus matched the category of the sample stimulus. During the sample and delay periods, the monkeys must hold in memory the category of the sample stimulus, but the outcome of the trial (i.e., the appropriate motor response) is unknown. During the test period, monkeys had to determine whether the test image was a category match to the sample image (and release the response bar) or a category nonmatch (and continue to hold the response bar).

phase). Only after the test stimulus appeared could the decision be made whether to release the response bar based on the category of the test stimulus and its relation (match or nonmatch) to the sample's category. Importantly, this allowed us to separately examine neural information about the sample category and the match/nonmatch status of the trial. Trials of each category set (i.e., cars and animals) and each response type (i.e., match vs. nonmatch) were randomly interleaved and occurred at similar frequency. Test stimuli were always from the same category set as the sample stimuli (e.g., if the sample stimulus was a car, the test stimulus was also a car) but could either be a category match (e.g., Sports Car–Sports car) or a category nonmatch (e.g., Sports Car–Sedan).

Data Analysis

Analysis Epochs

Initially, category data were analyzed over three time epochs throughout the trial. The fixation analysis epoch included the last 500 msec before the onset of the sample

image. During this time, the eyes were stable and no images were on the screen—this served as a baseline measure of neural activity. The sample analysis epoch included data from 100 to 600 msec after the sample onset and assessed the time when the sample image was present on the screen (adjusted for the typical visual delay to PFC). The delay analysis epoch was analyzed from 300 to 1100 msec after the sample offset and captured the period when no image was physically present on the screen, but the monkey was holding in working memory the category of the sample image (again adjusted for PFC neural delay). The use of these analysis epochs provided a baseline and two periods when the only information present was the sample image (and its category). The timing of these analysis epochs is consistent with our previous studies (Cromer et al., 2010; Roy et al., 2010; Freedman et al., 2001, 2002, 2003). These analysis epochs were only used for statistical tests (*t* tests), which identified the population of category-sensitive neurons. Other analyses we present (e.g., receiver operating characteristics [ROCs] and mutual information) did not restrict analysis to these periods. Additionally, a test epoch was examined for the first 500 msec after the test onset to assess the time when the match/nonmatch decision was made.

Randomization Tests

To assess latency differences between neuronal activity in PFC and PMC, we utilized a mutual information statistic (Buschman & Miller, 2007), which, similar to ROC statistic, gives a measure of how well a neuron encodes a distinction between two distributions. Specifically, we computed the actual mutual information for match versus nonmatch trials. To determine the significance of this measure, we next preformed randomization tests (Buschman & Miller, 2007). All trials of both types were randomly sorted and assigned to arbitrary "match" or "nonmatch" groups, and theoretical mutual information statistics were calculated each time for 5,000 permutations. This provided an estimate of the population under the null hypothesis that a neuron provided no real information to dissociate trial types. A p value was then determined by summing the number of theoretical mutual information values that were less than the actual mutual information value, dividing by the number of repeats, and subtracting this value from 1.

To determine the latency of neuronal encoding of the match vs. nonmatch distinction, we repeated the above mutual information calculations and randomizations over time in 25-msec bins. We identified the first time after test stimulus onset that showed significant information ($p < .01$) on the basis of the randomization tests. To ensure that this time best captured a period of robust neural information encoding, we required that two consecutive 25-msec bins met this criteria and subsequently marked the first time bin as the time of first significance. This allowed us to create a distribution of the leading (first) times after test onset when neurons showed significant match versus nonmatch information.

We next computed a cumulative first significance distribution (each time point value is the number of significant neurons at that time point and any previous time points). To compare these latency values across brain regions, we converted these values to z scores by comparing against a population mean and standard deviation created using the same techniques on 5,000 randomly generated mutual information matrices.

RESULTS

Behavior

Monkey "ti" and "lu" were both proficient on the delayed match-to-category task. Both monkeys' performance was similar, so their combined performance is shown in Figure 2. Individual performance of each monkey on the task is also reported elsewhere (Cromer et al., 2010). Each monkey correctly categorized (>80% correct in all cases) images of one category as belonging to that category on most of the trials while seldom missing classifications or incorrectly classifying images that were on the opposite side of the category boundary. Performance

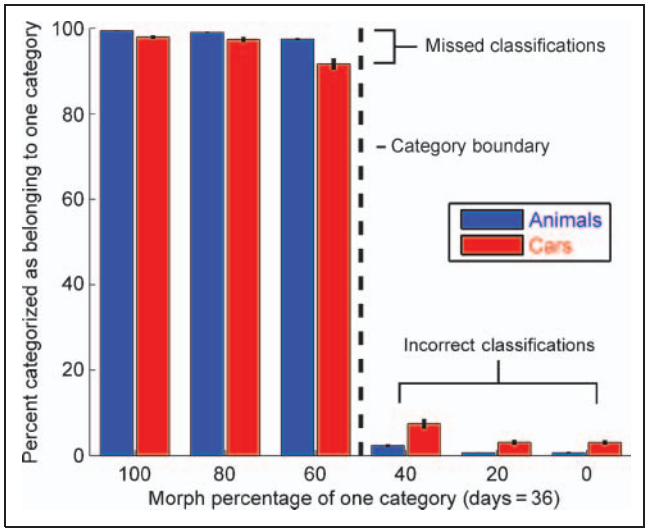

Figure 2. Behavioral results. Performance of both monkeys on the delayed match-to-category task with multiple, independent category distinctions across all recording sessions. Perfect performance would be 100% categorization of images at >50% of one category and 0% categorization of images at <50% of that one category (i.e., only classifying images at >50% of a category as members of that category). Actual performance was similar and displayed a hallmark step function in behavior at the category boundary. The majority of errors came on the 60% (missed classifications) or 40% (incorrect classifications) morphs that were closest to the category boundary and, therefore, were expected to be the hardest to correctly categorize. Error bars represent *SEM*.

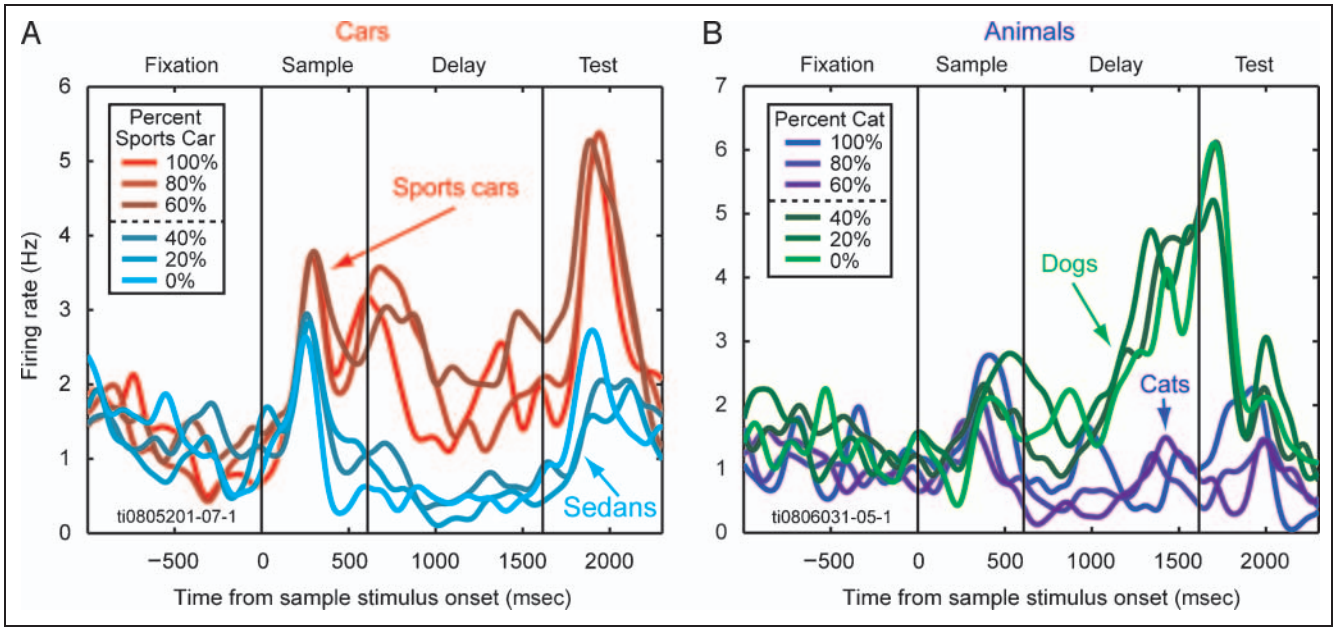

Figure 3. Single neuron examples—category sensitivity. (A) A single PFC neuron that showed distinct firing for stimuli of one category (i.e., Sports Cars) versus the other category (i.e., Sedans) starting during the robust burst of firing associated with stimulus onset and persisting throughout the delay and test epochs. Note how all morph percentages on either side of the category boundary (50%) grouped together (e.g., blue vs. red lines), despite the fact that sample images near the boundary line (60%/40%, darkest lines of each color group) were closer in physical similarity. Thus, this neuron responded to the category membership of the stimuli rather than their visual properties. (B) A second PFC neuron that categorized animals (Cats vs. Dogs) starting in the mid-delay period and into the test period.

remained high even for morphs closest to the category boundary (i.e., those images made with 60% of one category and 40% of the opposite category). For both category sets (cars and animals), the monkeys' behavior showed the hallmark of perceptual categorization: A much greater distinction between than within categories, with a sharp, "step" change in behavior across the category boundary. As previously reported, monkeys made significantly more errors and reacted slower on the car category set than the animal set, suggesting that the car set was more difficult (Cromer et al., 2010). However, behavioral performance was high and well above criterion for both sets.

PFC and PMC Neural Activity to Independent Category Sets

It was previously known that PFC neurons are sensitive to categories (Freedman et al., 2001), and in a previous article using the same PFC data used here, we reported that many PFC neurons showed category selectivity for both the cars and animals category sets. The results from PFC were similar in both monkeys, so their data were combined (Cromer et al., 2010). Examples of category selectivity in single neurons are shown in Figure 3. Figure 3A shows a neuron whose firing rate is grouped according to the car category membership of the sample image. Figure 3B depicts a second neuron that categorized Cats versus Dogs.

To quantify the number of neurons in both recorded brain areas that were category sensitive, we focused on neural activity during the sample and delay intervals, which is when the monkeys had to categorize the sample stimulus and retain that information in working memory. As in prior work (Cromer et al., 2010; Roy et al., 2010; Freedman et al., 2001, 2002), a t test was performed on each neuron's firing rate to all sample images from one category versus all images from the other (i.e., all Sports Car vs. Sedan images or all Cat vs. Dog images) to determine whether a neuron was category sensitive (i.e., showed significant category information, $p < .01$). We have previously shown in multiple studies that the population of neurons selected by this method displays the hallmarks of categorization (Cromer et al., 2010; Roy et al., 2010; Freedman et al., 2001, 2002).

Over one third of randomly selected neurons in the lateral PFC showed a significant difference between their average activity in the sample and/or delay intervals for one category set over the other (animals: 37%, 167 of 455; cars: 38%, 173 of 455). This was a sharp contrast to the PMC, which had relatively few neurons that were category sensitive (animals: 9%, 16 of 185; cars: 11%, 21 of 185). The results of these t tests to identify category-sensitive neurons are summarized in Table 1 for the PFC and in Table 2 for the PMC. These tables show the number of category-sensitive neurons for each category set (cars and animals) as well as which neurons were category sensitive for both sets, one set or the other only, or neither

Table 1. PFC Category Sensitivity

	Animal Sensitive	*Not Animal Sensitive*	*Total*
Car sensitive	104 (23%)	69 (15%)	173 (38%)
Not car sensitive	63 (14%)	219 (48%)	282 (62%)
Total	167 (37%)	288 (63%)	455 (100%)

The table lists the count (and percentage) of all recorded PFC neurons (455) based on whether each neuron was category sensitive (t tests, $p < .01$) for both the animal (Cats vs. Dogs) and car (Sports Cars vs. Sedans) category sets, only one set or the other, or neither set. The last row and column show the sums and specify neuronal selectivity for one category set independent of selectivity for the alternative category set.

set. They highlight that the PFC has a much greater percentage of neurons containing category information than the PMC.

As in prior studies (Cromer et al., 2010; Freedman et al., 2002), we examined the temporal dynamics of category effects in neural activity across the entire population of recorded neurons. This is shown for PFC and PMC in Figure 4, which details the information that each neuron carried about each category distinction as a function of time within the trial. We used a sliding window (200 msec in duration) and calculated an ROC statistic for each category contrast at 25-msec time steps throughout the entire trial. We used rectified ROC values that indicated the level of category sensitivity, but not which category was preferred. Higher ROC values (approaching 1) indicate a larger degree of difference in neuronal activity to sample images of the different categories (orange or yellow colors), whereas lower ROC values (near 0.5) indicate no or weak category sensitivity (darker colors). In Figure 4, each row corresponds to a single neuron, and neurons were sorted based on the time at which ROC values for each neuron reached an ROC threshold of 0.6 (Muhammad et al., 2006). This arbitrary threshold was chosen for comparability to previous work (Muhammad et al., 2006) and was used for display purposes only so that data could be segregated based on the timing of selectivity. The percentage of selective neurons was determined by t test (as above), although results were similar when an ROC threshold was used. Those neurons that did not reach this threshold were left unsorted at the top of each panel. Note the variability in the timing of category-related activity in PFC (Figure 4A and C). Some neurons are category sensitive only transiently (short orange bands), whereas others maintain their selectivity over a longer duration (long orange bands). Selectivity can start early in the sample period for some neurons, whereas other neurons do not show category selectivity until well into the delay or test periods. The peak of category selectivity varied similarly (see also Cromer et al., 2010). Most importantly, note how there is strong category information for approximately one third of PFC neurons during both animal (Figure 4A) and car (Figure 4C) categorization, whereas relatively few PMC neurons show high-category ROC values for either animals (Figure 4B) or cars (Figure 4D).

To further compare the strength of category selectivity of individual neurons across the two brain regions, we plotted for each neuron the mean ROC value for the animal category distinction against its mean ROC value for the car category distinction (Figure 5). Because the latency and duration of category selectivity was variable across neurons (Figure 4), we used the mean ROC computed using a 500-msec window centered at the time of the maximum ROC from either the sample or delay periods (Cromer et al., 2010). We coded the ROC values according to whether the neurons were from PFC (circles) or PMC (triangles) and color coded those neurons that were identified by t test (as above, $p < .01$) to be category sensitive (PFC, yellow circles; PMC, red triangles). In this case, we plotted nonrectified ROC values ranging from 0 to 1 to capture the category preference of each neuron (i.e., values near 0.5 indicate no category sensitivity, ROC values approaching 1 indicate greater activity for either Cats or Sports Cars, whereas ROC values approaching 0 indicate greater activity for Dogs or Sedans). In addition to category-sensitive neurons in PFC being more prevalent, PFC neurons also had higher mean ROCs than category-

Table 2. PMC Category Sensitivity

	Animal Sensitive	*Not Animal Sensitive*	*Total*
Car sensitive	4 (2%)	12 (7%)	16 (9%)
Not car sensitive	17 (9%)	152 (82%)	169 (91%)
Total	21 (11%)	164 (89%)	185 (100%)

The table lists the count (and percentage) of all recorded PMC neurons (185) based on whether each neuron was category sensitive (t tests, $p < .01$) for both the animal (Cats vs. Dogs) and car (Sports Cars vs. Sedans) category sets, only one set or the other, or neither set. The last row and column show the sums and specify neuronal selectivity for one category set independent of selectivity for the alternative category set.

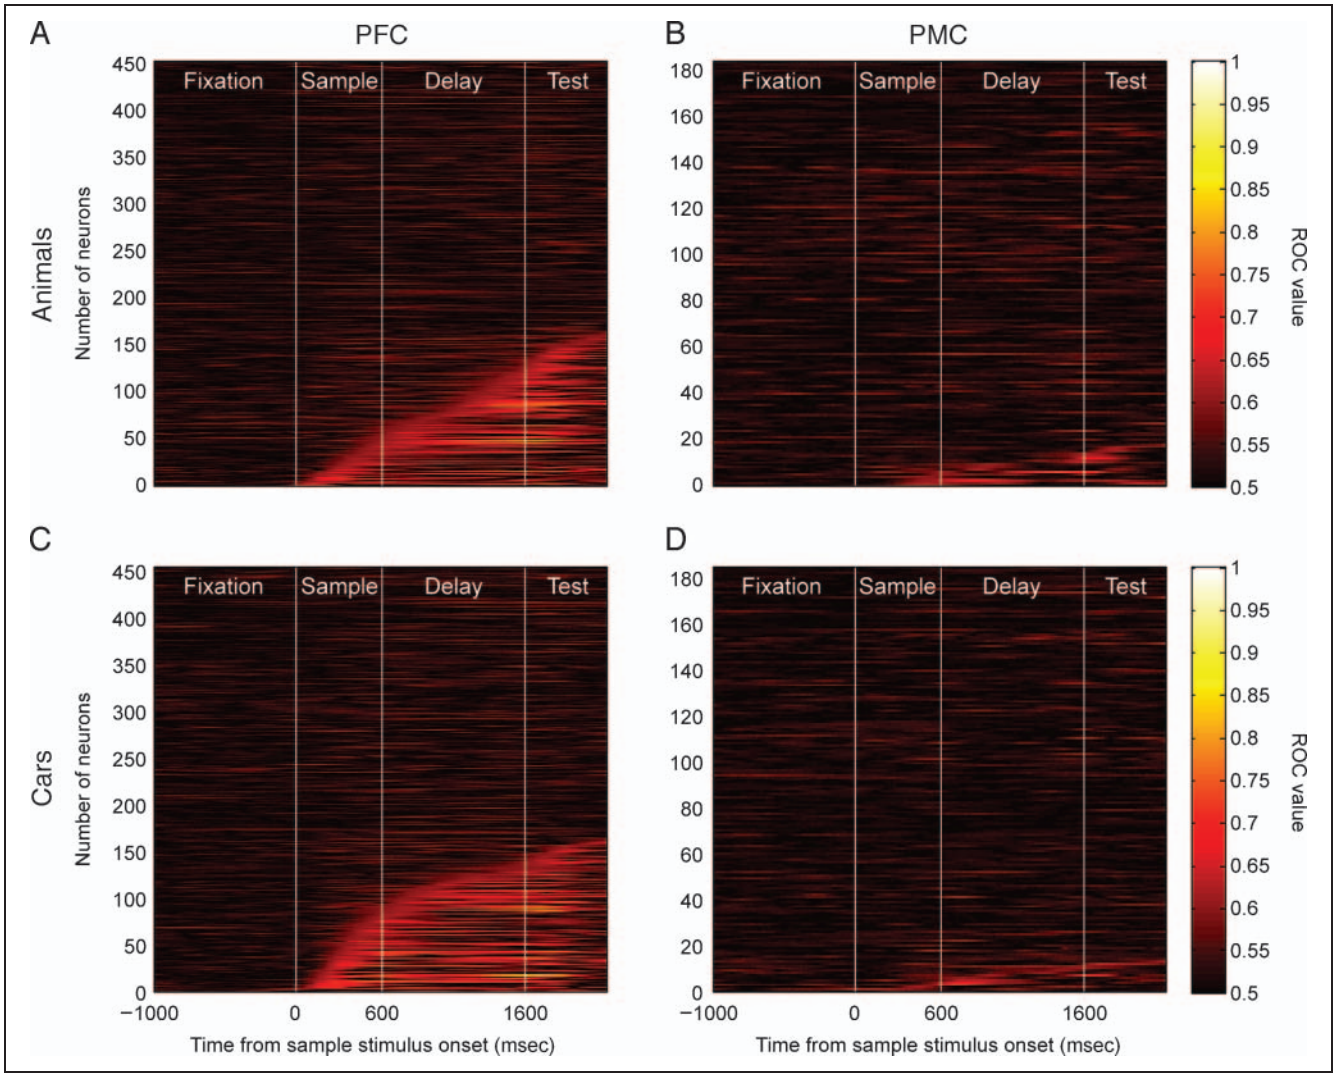

Figure 4. ROCs—category sensitivity. ROC values showing where neurons differentiate between categories for each of the recorded 455 PFC neurons (left) and 185 PMC neurons (right). Bright orange colors indicate high-category sensitivity. (A) ROCs over time for all PFC neurons showing when there was a distinction between trials with cat versus dog images (bright orange). Each row corresponds to a single neuron. Neurons were aligned if their ROC values reached 0.6 based on the earliest time of the ROC reaching this threshold. The earliest neurons show information shortly after the sample stimulus onset and approximately one third of recorded neurons reach the 0.6 threshold. (B) ROCs for all PMC neurons with high values indicating a distinction between cat versus dog trials. Note the lack of orange coloring indicating few PMC neurons that were category sensitive. (C) ROCs for all PFC neurons during car trials indicate a large percentage of PFC neurons differentiated between Sports Cars versus Sedans. Latencies are similar to animal categorization, but more neurons show activation during the sample period. (D) ROCs for all PMC neurons during car trials again indicate little category information in the PMC.

sensitive neurons from the PMC (t test, $p < .01$). In short, compared with the PFC, the PMC had a lower percentage of category-selective neurons with weaker effects.

PFC and PMC Neural Activity to Category Match Decision

After the delay epoch, the monkey saw a test image. It had to decide whether that test image was in the same category as the sample image and, if so, release a lever. Many PFC neurons have been shown to reflect this category matching decision (Cromer et al., 2010; Freedman et al., 2001), but this has not been tested in the PMC. We found that the PMC was strongly engaged during the test image epoch. Like PFC neurons, many PMC neurons distinguished between match and nonmatch trials (Figure 6). Some neurons showed match enhancement or a higher firing rate to match trials (Figure 6A), whereas others showed match suppression or a lower firing rate during match trials (Figure 6B), regardless of which category set was used in the trial.

We used t tests to assess the number of neurons that distinguished between match versus nonmatch trials in their average activity across the test stimulus epoch. In both brain regions, a large number of recorded neurons showed a significant (t test, $p < .01$) effect (PFC: 42%, 190 of 455; PMC: 62%, 115 of 185). Both regions had more

neurons that showed match enhancement (PFC: 55%, 104 of 190; PMC: 64%, 74 of 115) than match suppression (PFC: 45%, 86 of 190; PMC: 36%, 41 of 115).

To examine the temporal dynamics of the match/nonmatch effects, we again plotted rectified ROCs for each recorded neuron over time (Figure 7). This clearly demonstrates the prevalence of match-selective neurons in both brain regions. There was a greater percentage in the PMC versus the PFC. The time course of activation also appears similar in both brain regions, with some neurons becoming selective shortly after the test stimulus onset and a range of activations after that time. To quantify this relationship and identify whether either of PFC or PMC encoded this information with a shorter latency, we analyzed the time to first significant selectivity of each neuron by mutual information analysis and randomization tests (see Methods; Buschman & Miller, 2007). Figure 8A shows histograms of the time after test onset when each neuron first exhibited significant information about whether the test object was a match or a nonmatch to the sample. To compare the latency of information onset across PFC and PMC, we computed z scores (see Methods) of the cumulative first significance distribution (Figure 8B). Both brain regions first showed statistical significance ($p < .05$) information at the same latency, about 100 msec after test onset, well before the monkeys' average RT of 284 msec.

DISCUSSION

We recorded from lateral PFC and dorsal, rostral PMC neurons while monkeys categorized images from two independent category sets (cars and animals). We found that over 36% of randomly selected PFC neurons were sensitive to the category of a sample image in contrast to less than 12% category-sensitive neurons in the PMC. Furthermore, category sensitivity was weaker in the PMC than in the PFC. Thus, whereas the PFC has repeatedly been shown to be strongly involved in categorization, the PMC does not appear to maintain a similar representation of visual categories. However, close to half the neurons in both brain regions reflected the category match/nonmatch decisions that led to a behavioral response.

Although several studies have shown visual category selectivity in PFC (Cromer et al., 2010; Roy et al., 2010;

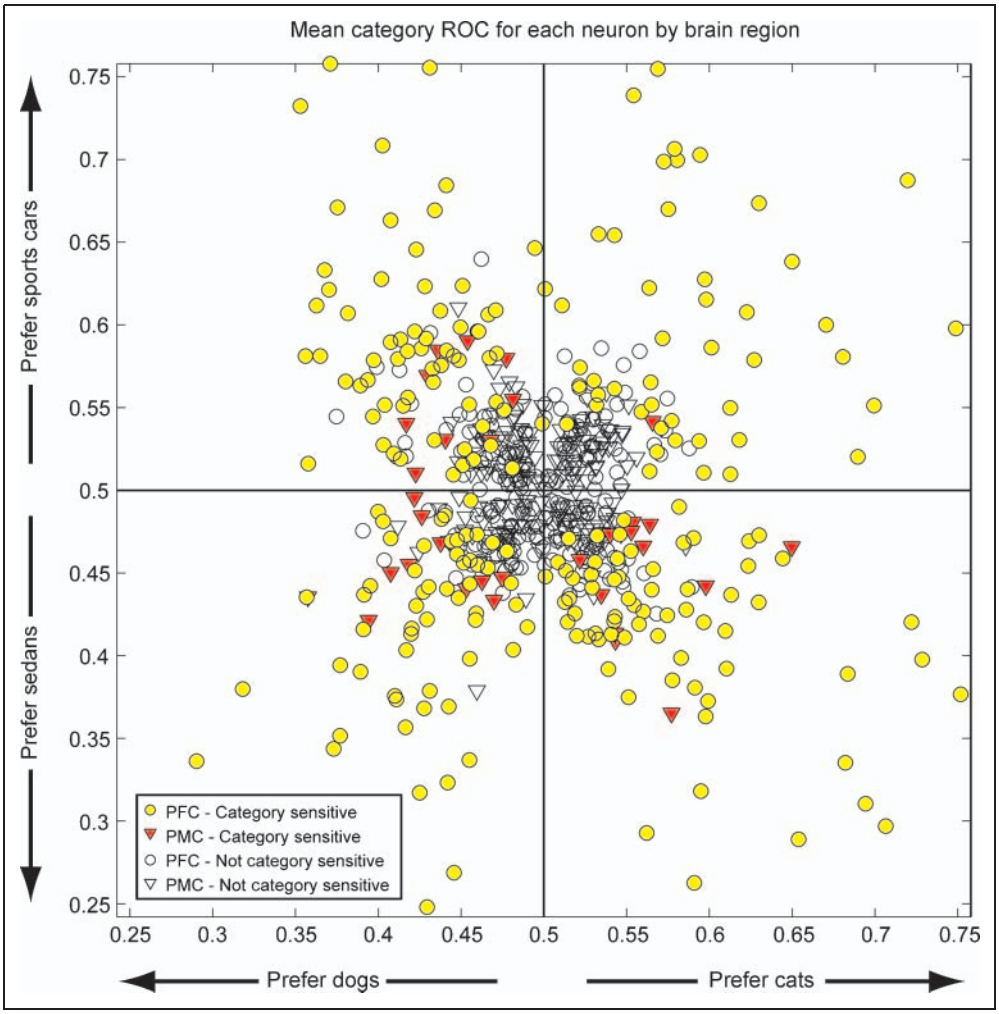

Figure 5. Mean ROCs for each category scheme. Mean ROC values for all recorded neurons in PFC (circles) and PMC (triangles) to both category distinctions. Colored points indicate significant category sensitivity as determined via t test for one or both categories. Values closer to the origin indicate weaker category sensitive for the given distinction (x values indicate the animals category distinction, and y values indicate the car distinction). Data points farthest from the origin indicate those neurons with the strongest category selectivity (yellow circles). PMC neurons that were category sensitive (red triangles) had weaker category information than the strongest PFC neurons.

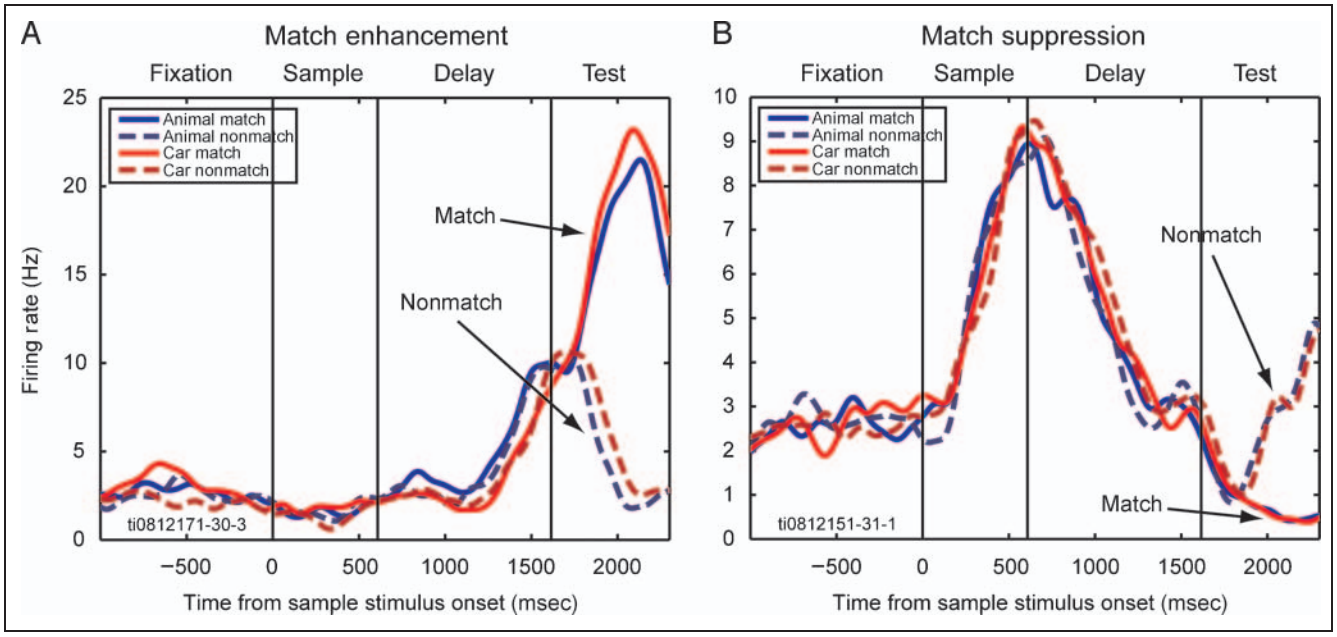

Figure 6. Single neuron examples—match sensitivity. (A) A single PMC neuron distinguished between when test images match ("match," solid lines) or did not match ("nonmatch," dashed lines) the sample images across trials. This distinction occurred regardless of whether images were from the animal or car category sets (lines cluster as solid vs. dashed rather than red vs. blue). This neuron's firing rate was higher during match trials, so it is said to have "match enhancement." (B) A single PMC neuron that distinguished match from nonmatch trials but had a lower firing rate for match trials and is, thus, characterized as having "match suppression."

Freedman et al., 2001, 2002), it was unclear whether this could be expected for the PMC. It has previously been shown that many PMC neurons strongly reflect abstract "same" vs. "different" rules (Muhammad et al., 2006). In fact, rule effects were stronger and earlier in the PMC than in the PFC. In contrast, we found little or no effect of image category in the PMC. One possible explanation for this discrepancy is that we recorded in a more dorsal part of the PMC than Muhammad et al. (2006) did. However, this region does receive direct projections from the lateral PFC (Petrides & Pandya, 2006) and has been suggested to be involved in more cognitive processes (Picard

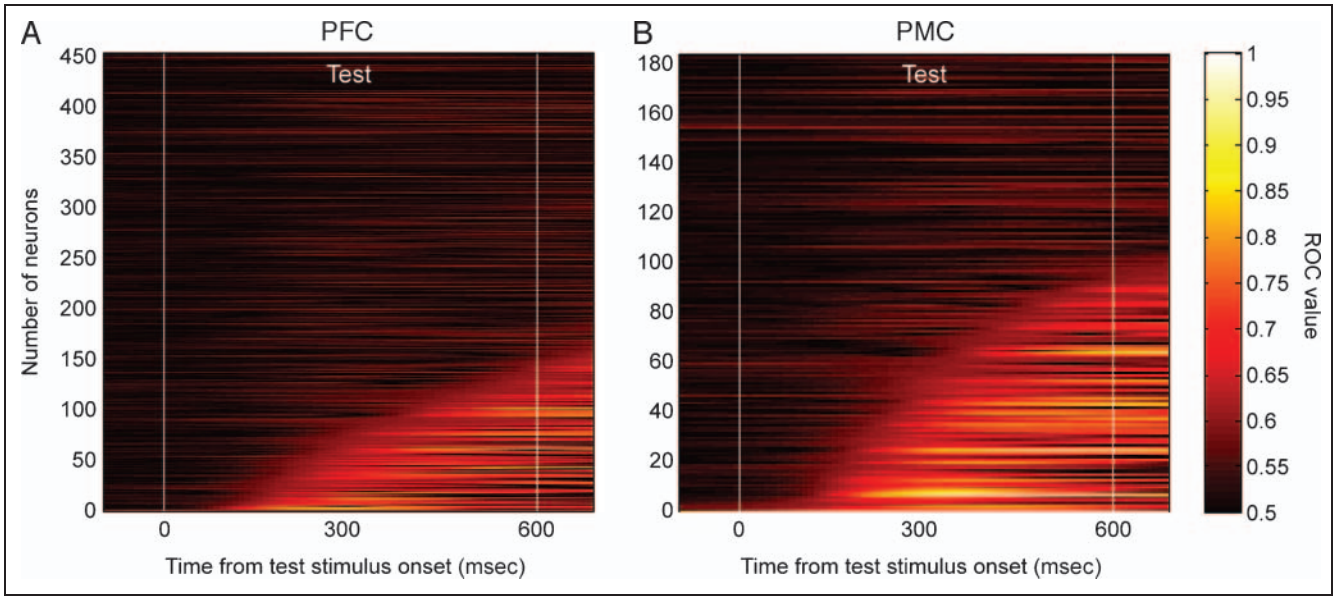

Figure 7. ROCs—match sensitivity. (A) ROC values for each of the recorded 455 PFC neurons across the test epoch with bright orange colors indicating sensitivity for match versus nonmatch trials, aligned when neurons reached a 0.6 ROC threshold. Similar to category effects, over one third of PFC neurons displayed match information. (B) ROC values for each of the 185 recorded PMC neurons show that a majority of PMC neurons dissociate match from nonmatch trials.

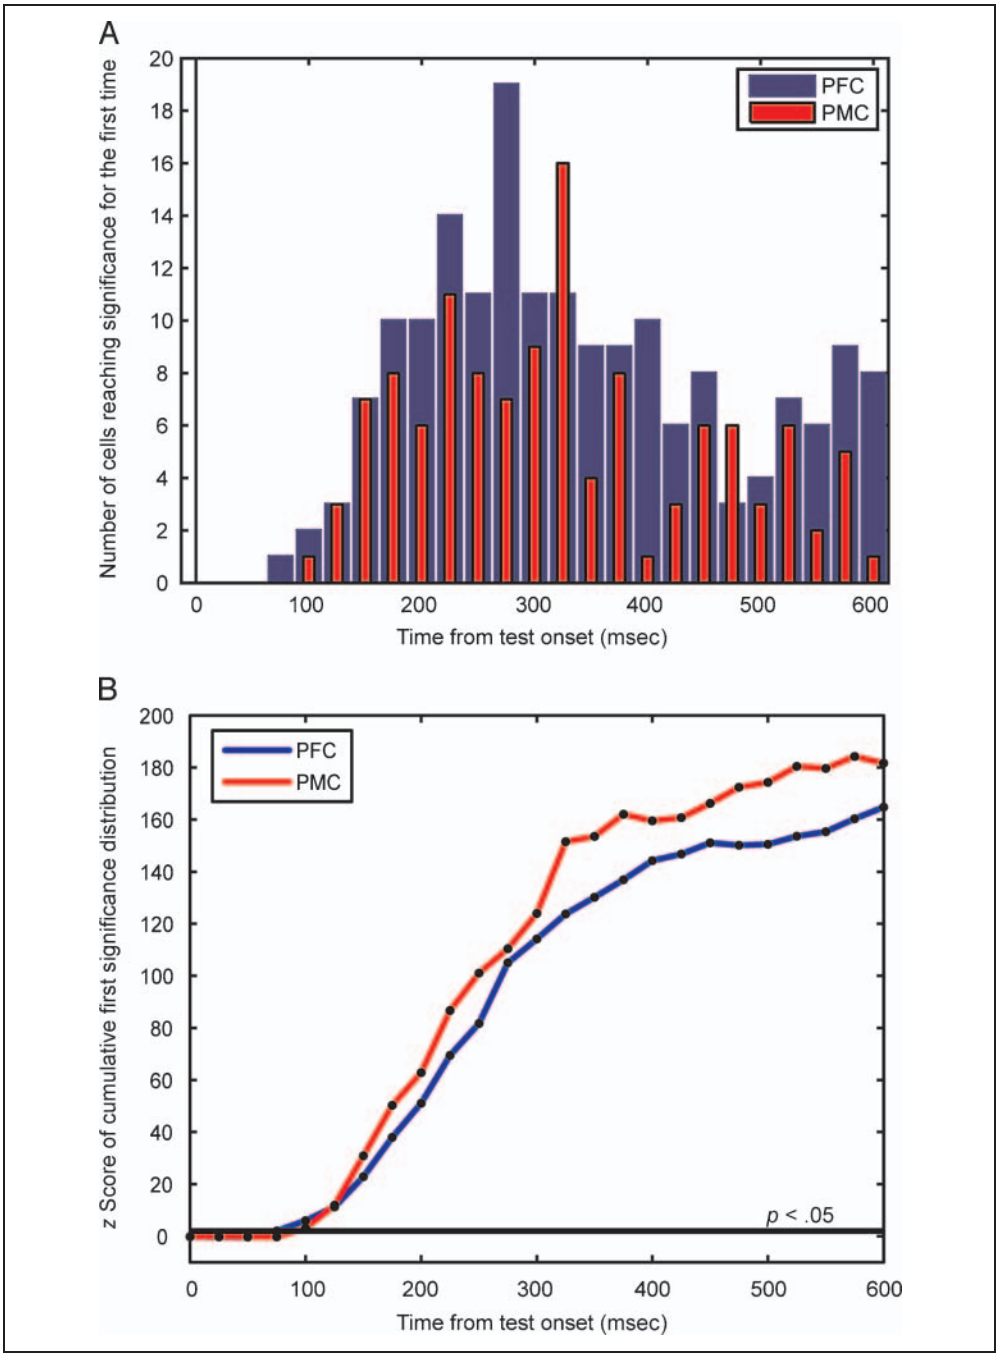

Figure 8. Latencies—match sensitivity. (A) Histograms of the number of cells reaching significance for the first time in each 25-msec time bin after test onset for PFC (blue) and PMC (red). Note that there were 455 recorded PFC neurons and 185 PMC neurons, so raw counts are not directly comparable. (B) *z* scores of the cumulative first significance distribution (see Methods) allowed comparison of latencies across PFC and PMC. Both regions reach significance ($p < .05$ or 2 *SD*s above the mean) at the 100-msec time bin.

& Strick, 2001). It seems more likely that the difference between studies is due to a division between sensory and motor processing. The delayed match-to-category task was designed so that the category membership of an image was decoupled from the motor response. The category of the sample image contained no information that could be used to resolve the motor response. It was only after the appearance of the test stimulus that the monkey could decide which motor response was necessary. The monkeys in Muhammad et al. (2006) also could not predict the motor response until test stimulus presentation. However, the rules themselves (which were presented at the beginning of a trial) contained crucial information that allowed the monkeys to map the status of the test stimulus to the motor response: For "same," the monkeys released the lever if the test image matched the sample, whereas for "different," the monkeys withheld a response to a match. Thus, the difference in PMC activity to categories versus abstract rules may be explained by the PMC only being strongly engaged when motor-related information is being retained and processed. This interpretation is further supported by the observation that many PMC neurons were strongly and selectively activated by category matching (whether sample and test image categories matched), which finally added motor information to the category task on each trial.

In summary, although PMC neurons can reflect high-order "cognitive" information above and beyond simple premotor activity, they, unlike PFC neurons, are only engaged when the information has a motor component. This fits well with PFC's hypothesized role in multimodal executive functions (Miller & Cohen, 2001) and the PMC's role in motor planning (Hoshi & Tanji, 2007; Wise, 1985).

Acknowledgments

We thank E. G. Antzoulatos, S. L. Brincat, J. R. Cromer, C. Diogo, K. MacCully, L. N. Nation, D. Ouellette, M. V. Puig, J. Rose, and M. Siegel. This work was supported by the National Institute of Mental Health (5R01MH065252-09 to E. K. M.) and the National Eye Institute (F32EY018993 to J. A. C.).

Reprint requests should be sent to Earl K. Miller, The Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, or via e-mail: ekmiller@mit.edu.

REFERENCES

Beymer, D., & Poggio, T. (1996). Image representations for visual learning. *Science, 272,* 1905–1909.

Boettiger, C. A., & D'Esposito, M. (2005). Frontal networks for learning and executing arbitrary stimulus–response associations. *Journal of Neuroscience, 25,* 2723–2732.

Bolte, S., Holtmann, M., Poustka, F., Scheurich, A., & Schmidt, L. (2007). Gestalt perception and local-global processing in high-functioning autism. *Journal of Autism and Developmental Disorders, 37,* 1493–1504.

Buschman, T. J., & Miller, E. K. (2007). Top–down versus bottom–up control of attention in the prefrontal and posterior parietal cortices. *Science, 315,* 1860–1862.

Cromer, J. A., Roy, J. E., & Miller, E. K. (2010). Representation of multiple, independent categories in the primate prefrontal cortex. *Neuron, 66,* 796–807.

Della-Maggiore, V., & McIntosh, A. R. (2005). Time course of changes in brain activity and functional connectivity associated with long-term adaptation to a rotational transformation. *Journal of Neurophysiology, 93,* 2254–2262.

Freedman, D. J., & Miller, E. K. (2008). Neural mechanisms of visual categorization: Insights from neurophysiology. *Neuroscience and Biobehavioral Reviews, 32,* 311–329.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science, 291,* 312–316.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2002). Visual categorization and the primate prefrontal cortex: Neurophysiology and behavior. *Journal of Neurophysiology, 88,* 929–941.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2003). A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *Journal of Neuroscience, 23,* 5235–5246.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2006). Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. *Cerebral Cortex, 16,* 1631–1644.

Hoshi, E., & Tanji, J. (2007). Distinctions between dorsal and ventral premotor areas: Anatomical connectivity and functional properties. *Current Opinion in Neurobiology, 17,* 234–242.

Jiang, X., Bradley, E., Rini, R. A., Zeffiro, T., Vanmeter, J., & Riesenhuber, M. (2007). Categorization training results in shape- and category-selective human neural plasticity. *Neuron, 53,* 891–903.

Kuperberg, G. R., West, W. C., Lakshmanan, B. M., & Goff, D. (2008). Functional magnetic resonance imaging reveals neuroanatomical dissociations during semantic integration in schizophrenia. *Biological Psychiatry, 64,* 407–418.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24,* 167–202.

Miller, E. K., Nieder, A., Freedman, D. J., & Wallis, J. D. (2003). Neural correlates of categories and concepts. *Current Opinion in Neurobiology, 13,* 198–203.

Muhammad, R., Wallis, J. D., & Miller, E. K. (2006). A comparison of abstract rules in the prefrontal cortex, premotor cortex, inferior temporal cortex, and striatum. *Journal of Cognitive Neuroscience, 18,* 974–989.

Petrides, M., & Pandya, D. N. (2006). Efferent association pathways originating in the caudal prefrontal cortex in the macaque monkey. *Journal of Comparative Neurology, 498,* 227–251.

Picard, N., & Strick, P. L. (2001). Imaging the premotor areas. *Current Opinion in Neurobiology, 11,* 663–672.

Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A. M., Pardo, J. V., Fox, P. T., et al. (1994). Practice-related changes in human brain functional anatomy during nonmotor learning. *Cerebral Cortex, 4,* 8–26.

Roy, J. E., Riesenhuber, M., Poggio, T., & Miller, E. K. (2010). Prefrontal cortex activity during flexible categorization. *Journal of Neuroscience, 30,* 8519–8528.

Scherf, K. S., Luna, B., Kimchi, R., Minshew, N., & Behrmann, M. (2008). Missing the big picture: Impaired development of global shape processing in autism. *Autism Research, 1,* 114–129.

Seger, C. A., & Miller, E. K. (2010). Category learning in the brain. *Annual Review of Neuroscience, 33,* 203–219.

Shelton, C. R. (2000). Morphable surface models. *International Journal of Computer Vision, 38,* 75–91.

Uhlhaas, P. J., & Mishara, A. L. (2007). Perceptual anomalies in schizophrenia: Integrating phenomenology and cognitive neuroscience. *Schizophrenia Bulletin, 33,* 142–156.

Wise, S. P. (1985). The primate premotor cortex: Past, present, and preparatory. *Annual Review of Neuroscience, 8,* 1–19.