Selective Attention and the Visual Representation of Object Attributes in the Ventrolateral Prefrontal Cortex and Anterior Cingulate Cortex of the Rhesus Monkey

by

Cynthia E. Kiddoo

B. A. Math and Computer Science
Mills College, 1991

Submitted to the Department of Brain and Cognitive Sciences in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Brain and Cognitive Science
at the
Massachusetts Institute of Technology

September, 2005

Signature of Author: _____
Brain and Cognitive Science
August 31, 2005

Certified by: _____
Earl K. Miller
Picower Professor of Neuroscience
Thesis Supervisor

Accepted by: _____
Earl K. Miller
Picower Professor of Neuroscience
Chairman, Department Graduate Committee

Selective Attention and the Visual Representation of Object Attributes in the Ventrolateral Prefrontal Cortex and Anterior Cingulate Cortex of the Rhesus Monkey

by

Cynthia E. Kiddoo

Submitted to the Department of Brain and Cognitive Sciences on June 9, 2005 in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Brain and Cognitive Science

ABSTRACT

The effects of attending to one or another of an object's attributes on neuronal representations of that object were investigated using extracellular recordings. A female rhesus monkey performed a delayed match to object attribute (DMSA) task, in which she alternately matched object orientations and object colors. In half of the task conditions, only one attribute matched the sample, forcing the animal to apply the current matching rule and ignore the irrelevant attribute. Multiple simultaneous single-unit extracellular recordings were made in the ventrolateral prefrontal cortex (VLPFC) and anterior cingulate cortex (ACC) while the monkey performed the task. Neuronal selectivity for matching rule, object attributes, attribute relevance, response choice, and congruency were assessed using multi-factor ANOVAs.

Attribute-selective responses were common in both cortical areas during the sample and delay periods, but were not significantly modulated by attribute relevance. There were few interactions between color-selective and orientation-selective responses according to the ANOVAs, suggesting that these attributes were represented independently. Significant effects of attribute relevance, response choice, and congruency appeared in both areas after the delay period, when the probe appeared onscreen. VLPFC cells were more active during incongruent and non-match conditions, when responses had to be suppressed. ACC cells were more active during congruent and match conditions, when active response suppression was not required. The results indicate that although prefrontal cortex often shows a bias for relevant information (Rainer et al, 1998), it may not do so if the task requires frequent alternation of attentional sets or active suppression of conflicting responses. The data also indicate that the VLPFC's role in managing attentional 'set' (Banich et al, 2000; Milham et al, 2001) is performed in conjunction with active stimulus comparison and response selection (e.g., Rushworth et al, 1997), not during working memory maintenance. The ACC may facilitate the reactivation of response tendencies that had been actively suppressed, possibly as part of a larger role in managing response conflict (Botvinick et al, 2004).

Thesis Supervisor: Earl K. Miller
Title: Professor of Brain and Cognitive Science

Introduction

The ability to filter incoming visual information strategically and selectively is essential to many complex cognitive behaviors. This ability depends critically on the prefrontal cortex (PFC). In studies of awake, behaving monkeys, the activity of single neurons in PFC has been shown to discriminate visual stimuli based on their features, their locations in space, or both (Funahashi et al., 1989; Fuster et al., 1982; Rao et al., 1997; Rosenkilde et al., 1981; Watanabe, 1981). These responses depend not just on the physical properties of the stimuli, but also on their significance in terms of the subject's current expectations and goals (Asaad et al., 2000; Rainer et al., 1998; White and Wise, 1999).

Our laboratory has explored the joint representation of visual objects and their context-dependent significance in PFC under a variety of behavioral demands. In one experiment, monkeys were required to locate a target within a three-object sample array, and then indicate whether the target appeared in the same location in a test array. Nearly half of the recorded PFC neurons responded differently to identical arrays, depending on which location contained the target. For many of these cells, the response to the target in the preferred location was the same amplitude, whether the target appeared alone or within an array (Rainer et al., 1998). In another experiment, monkeys had to make a categorical judgment, using stimuli that varied smoothly from one set of category exemplars ("cats") to another ("dogs"). PFC neurons showed similar responses for stimuli belonging to the *same* category, even though the stimuli were physically different; and dissimilar responses for stimuli belonging to different categories, even when those stimuli looked nearly the same (Freedman et al., 2001). Despite differences in the task

demands and stimulus sets used in these two studies, both demonstrated an important characteristic of object representations in PFC: Relevant information is preferentially encoded, and irrelevant information is not. Thus, when objects in an array are to be ignored, their presence makes little impact on neuronal responses; and when dissimilar stimuli have equivalent meaning within a category, their distinguishing features are, at best, only weakly encoded.

I extended these earlier investigations by examining the case where a subject must selectively attend to some subset of an object's attributes. This issue is of concern for describing the neural basis of adaptive behavior, and it is reflected in everyday experience. For example, imagine hunting through a toolbox for a screwdriver. If you intended to use the screwdriver to unscrew a screw, you would focus on the size and tip shape of screwdrivers you encountered in your search. If you wanted instead to use the screwdriver to pry up a lid, you would focus on the screwdrivers' shaft diameter and length. This example illustrates how different attributes of an object may become the focus of attention, depending on one's goals and needs.

In order to learn whether the representation of an object in PFC changes depending on which of its attributes is attended, I trained a monkey to perform a delayed-match-to-sample-attribute task (DMSA). In this task, the monkey had to indicate whether a probe object matched the previously viewed sample object along the relevant attribute dimension. The sample objects were characterized by two attribute dimensions, color and orientation, and the relevant attribute alternated between blocks. In each block, the monkey had to attend to the relevant sample attribute, and ignore the irrelevant sample attribute. This task is similar to the category-matching task of Freedman and colleagues

(Freedman et al., 2001), in that a set of probe objects matching along attribute dimension A, but varying along attribute dimension B, form a natural category. However, the stimuli used in this experiment were classifiable by two orthogonal dimensions rather than just one, and the locations of "category" boundaries were determined by which sample attribute was relevant, and thus changed dynamically from one block to the next. In the experiments by Freeman and colleagues, the monkeys were trained on only one set of categories at a time, and the category boundaries were fixed during each set of recording sessions (Freedman et al., 2002). Consequently, the DSMA task enabled me to investigate the neuronal effects of dynamic shifts of attention from one object attribute to another, much as the experiment by Rainer and colleagues had examined dynamic shifts of attention from one location in space to another (Rainer et al., 1998).

The DMSA task bears many similarities to the Stroop task and Wisconsin Card-Sorting Task (WCST). In the WCST, for example, the subject must sort cards according to one of three possible stimulus dimensions (number, shape, and color). Because the relevant stimulus dimension changes without prior warning, the subject must monitor his or her performance and adapt to new reward contingencies by ignoring a previously relevant stimulus dimension and attending to a different one (Grant & Berg, 1948). In the Stroop task, subjects must read from three cards, one of which has the names of various colors printed in incompatible ink colors (e.g., the word "blue" printed in red ink). When given this third card, subjects are instructed to say the color of each word while ignoring the word's (incongruent) lexical attribute (Klein, 1964; MacLeod & MacDonald, 2000). Numerous functional imaging studies in humans using tasks such as the Stroop and WCST have shown activation of dorsolateral and ventrolateral PFC as

well as anterior cingulate cortex (ACC) and lateral intraparietal cortex (Carter et al., 1995; Duncan and Owen, 2000; Durston et al., 2002; Konishi et al., 1998, 1999; Menon et al., 2001; Pardo et al., 1990; Peterson et al., 2002; Rubia et al., 2003; Rushworth et al., 2001; Van Veen et al., 2004). Few such studies, however, have been done in non-human primates. Using fMRI, Konishi and colleagues found activation in the posterior ventrolateral prefrontal cortex (VLFPC) in monkeys and humans performing a task analogous to the WCST (Konishi et al., 1999; Nakahara et al., 2002), but the response properties of single neurons remain unexplored. Because of these findings, and because VLPFC receives the majority of prefrontally directed afferents from areas TE and TEO (Webster et al., 1994), which play a critical role in processing object attributes such as shape, texture, and color (Gross et al., 1972; Ungerleider and Mishkin, 1982), I believed that VLPFC neurons likely play a role in representing relevant attribute information.

As noted above, ACC is typically activated in the Stroop task and the WCST. Many other studies have found ACC activation occurring in conjunction with a wide variety of cognitively demanding tasks (reviewed in Ridderinkhof et al., 2004). A model of ACC functioning proposed by Cohen and colleagues (Botvinick et al., 2001, 2004; Cohen et al., 1990; MacDonald et al., 2000) notes that a common feature of all of these tasks is response conflict: Each stimulus is capable of automatically eliciting two or more possible responses, only one of which will ultimately be expressed by the subject. Under these circumstances, selecting the "best" response may require engaging additional executive control processes. Cohen and colleagues suggest that the ACC monitors the need for these additional attentional resources, and activates the lateral PFC when they are required. Although this hypothesis has received substantial support from functional

imaging studies, it has not been tested extensively using single-unit electrophysiological techniques. The DMSA task is well suited to testing this hypothesis, because response conflict is present during conditions in which one of the probe attributes matches the sample object and the other does not. Because areas 12 and 47 of VLPFC are among the few lateral prefrontal areas directly connected to ACC (Barbas and Pandya, 1989; Pandya and Barnes, 1987; Pandya et al., 1981), any modulatory influence of ACC on neural activity in lateral PFC might be expected to be rapid and robust. Consequently, I chose to record from ACC as well as VLPFC.

The objective of this experiment was to test the generality of the hypothesis that the VLPFC preferentially represents relevant information. If the generality of this hypothesis was confirmed, I wanted to see whether the representation of preferred categories in VLPFC, as documented by Freedman and colleagues (Freedman et al., 2001, 2002), signifies a generalized mechanism for selectively attending to subsets of object properties. Because alternately attending to different sets of object attributes has the potential to create response conflict, I also set out to test the model of ACC and lateral PFC function proposed by Cohen and colleagues. If that model is correct, I expected to find strong effects of response conflict in ACC but not in VLPFC, and greater ACC activation when conflict was present than when it was absent.

Methods

Subject

An adult female rhesus macaque (*macaca mulata*) weighing 6 kg was used for this experiment. Recording chambers were implanted surgically over the ACC and the

VLPFC using stereotactic coordinates. A titanium post for immobilizing the head was implanted in an earlier surgery. All surgical procedures and post-operative care were conducted in accordance with the guidelines of the National Institutes of Health and the Animal Care and Use Committee of the Massachusetts Institute of Technology.

Before implanting the recording chambers, I collected coronal and sagittal structural images of the monkey's entire head using a 1.5T magnetic resonance imaging (MRI) scanner. The monkey was anesthetized with ketamine and placed in a stereotax, which was carefully leveled and aligned in the scanner using a positioning light beam. The anatomical images provided a detailed map of the ACC and lateral PFC, and enabled me to place the recording chambers as described below. The ACC recording chamber was centered directly above the anterior tip of the genu of the corpus callosum, so that its anterior edge was located above and 1-2 mm forward of the anterior terminus of the cingulate sulcus (Fig. 1B). Based on these landmarks and earlier anatomical mapping studies (Pandya and Barnes, 1987), I judged the chamber to be positioned directly above area 24. The ventrolateral recording chamber was centered over VLPFC, ventral to the principal sulcus and anterior to the ventral arm of the arcuate sulcus (Fig. 1A). A craniotomy approximately 1 cm in diameter was made within each recording chamber. The MRI scans indicated that this opening was wide enough to allow access to much of the medial/lateral extent of the cingulate sulcus in both hemispheres.

Recording Techniques

Data were collected while the monkey sat in a primate chair with her head fixed, facing a CRT monitor with a refresh rate of 60 Hz located 52 cm away. A cylindrical metal bar projected through the front panel of the primate chair. By grasping and

releasing this bar, and thereby closing or opening a circuit, the monkey could relay her behavioral decisions. A juice spout was positioned in front of the monkey's mouth for reward delivery. An infrared camera centered directly below the monitor provided video images of eye movements and enabled eye position tracking at 60-Hz resolution via dedicated video processing hardware (Iscan, Cambridge, MA). Although the experiment only required central fixation, the eye tracking system was freshly calibrated each day using both central and distal fixation targets in order to ensure that gains and offsets were accurate.

Extracellular recordings were made with thin-wire tungsten microelectrodes (FHC Instruments, Bowdoinham, ME). Multiple simultaneous penetrations were made using screw-turned miniature microdrives that controlled either a single microelectrode or a pair of microelectrodes (Nichols et al., 1998). These microdrives were attached to custom-made plastic grids with holes spaced 1 mm apart that fit snugly inside the recording chambers. Each grid rested on the lip of the chamber so that the bottom surface of the grid was always the same distance from the cortical surface, and was aligned each day within the chamber using set screws as reference points.

To minimize damage to the electrode tips as they descended through tissue, I lowered them through dura-piercing guide tubes made from 23 G needles. I tried to optimize the guide tube lengths so that they pierced the dura while doing minimal damage to the cortical surface, following a procedure developed in our laboratory by Jonathan Wallis. Before recording, the depth of the dura was mapped at multiple grid locations. In my initial recordings, I positioned the guide tubes such that their tips penetrated 2 mm below the surface of the dura. After withdrawing the electrodes at the end of the experiment, I

measured their impedances. If the impedances had dropped below 400 KΩ, or if the recording quality on a particular channel had been poor, I positioned the guide tube tip 0.5 mm lower for the following recording session. Using this approach, I was able to record successfully from the same locations over several consecutive days, suggesting that the guide tubes did minimal damage to the cortex.

The ACC recordings were mostly made in the upper bank of the cingulate sulcus. A few penetrations targeted the lower bank of the sulcus. I referred to the structural MRI images and used auditory feedback to determine the appropriate depth. I lowered the electrodes until neuronal activity was detected, and then allowed the brain to settle for 1 - 2 hours. I attempted to get the best possible isolation on each channel, but no effort was made to select neurons based on their response properties. I, therefore, obtained an unbiased sample of cells at each recording site. Recording sessions typically lasted 2 – 3.5 hours. The monkey was usually willing to work for much longer than she was willing to sit still while the brain settled, so it was sometimes necessary to cut the settling period short. On these occasions, I recorded for as long as the monkey was willing to work, and excluded the unstable portions of channels that showed evidence of drift from subsequent analysis.

Spike waveforms were amplified, filtered, digitized at 40 KHz, and stored for offline sorting (Plexon Systems, Dallas, TX). Clusters of waveforms with inter-spike intervals < 2 μsec were considered to be multi-unit, and were not included in the analysis. Before analyzing each neuron, I visually inspected its rastergram to check for drift over the course of the experiment. If drift was evident only near the beginning or end of the experiment, I discarded the corresponding trials from the analysis of that unit. If visual

inspection of the rastergram suggested that the recording was unstable over the entire session, I discarded the cell.

Behavioral Task

In the traditional delayed-match-to-sample (DMS) task, the subject is briefly presented with a sample object and, after an intervening delay, a probe object. The subject must compare the probe object to the sample, and indicate whether they are same. In my modified version of this task, the monkey had to indicate whether the probe matched the sample along the relevant attribute dimension, disregarding whether it matched along the irrelevant dimension. The probe stimuli were filled solid bars, 5° x 1° in size, whose color and orientation were pseudo-randomly varied. The monkey performed blocks of trials in which the relevant probe attribute (color or orientation) alternated between blocks. The relevant attribute assigned to the first block was alternated between recording sessions.

At the beginning of each trial, the monkey was required to grab and hold the bar in the front panel of the primate chair. If she released the bar prematurely, the response was counted as an error. Once the monkey was holding the bar, a white fixation spot appeared in the center of the screen. The monkey had to acquire fixation and maintain it within a 2° x 2° square window for 600 msec. While the monkey continued to fixate, the sample stimulus was displayed at the center of the visual field. At this point, the fixation window was expanded to 3.5° x 3.5° and remained this size for the remainder of the trial. Enlarging the fixation window was necessary because when stimuli were displayed and the monkey's pupils suddenly contracted, the eye tracking system produced a small (approximately 0.5°), transitory upward deflection of the eye position trace. Visual

inspection of the camera image indicated that this deflection did not correspond to an actual eye movement. The eye tracking system calculates eye position by comparing the location of the corneal reflection with respect to the center of the pupil. I speculate that there might have been a small degree of perspective distortion due to the camera's placement below the animal's sightline, tilting upward toward the eye. During the period of rapid pupil contraction, this distortion might have produced an error in computing the exact location of the pupil center. Once the pupil's aperture reached a steady state, the eye tracking system recovered and no further problems were noted.

The sample stimulus remained on the screen for 500 msec, and was followed by a 1000-msec delay period. After the delay, a probe object was displayed. If the relevant attribute of the probe matched the sample, the monkey had to release the bar within 800 msec to receive a reward of four drops of apple juice. If the relevant attribute of the probe did not match the sample, the monkey had to continue fixating and holding the bar. After the offset of the probe and an intervening 200-msec delay, an object matching the relevant sample attribute was displayed. The monkey was rewarded for releasing the bar when she detected this object. If the monkey made an error, the reward was withheld and a timeout delay of 1500 msec was imposed as punishment.

Each block of 16 trials was preceded by a set of "cue" trials that told the monkey which attribute was relevant. When color was relevant (color-matching rule), the sample object was a solid filled circle with a diameter of 3° (Fig. 2A). Because the only attribute that could be compared with the probe was color, such a trial signaled that subsequent trials would also be color-matching trials. When orientation was relevant (orientation-matching rule), the sample object was an elongated (6.25° x 0.50°) filled grey bar. The

monkey quickly learned to compare this sample object's orientation to that of the probe. None of the probe objects ever matched the sample's hue (Fig. 2B), and subsequent trials in the block were all orientation-matching trials. To ensure that the monkey adequately processed the cue trials and shifted her attention to the relevant attribute, she was required to perform three cue trials correctly before proceeding to the rest of the trials in the block.

After completing the cue trials, the monkey proceeded to the DMSA trials. The sample objects in DMSA trials were solid, filled bars the same size and shape as the probe stimuli. Each day a new set of four sample objects was generated using two different color attribute values and two different orientation attribute values. The relevant sample attributes for the cue trials used the same color and orientation values as the DMSA trials.

Although the sample set consisted of only four stimuli, the set of probe objects was much larger. For each sample orientation, I defined a set of 48 possible non-matching probe orientations. To ensure accurate perceptual discrimination, all 48 non-matching orientations differed from the sample orientation by at least 22.5°, a difference sufficient to produce consistently high performance. The range of possible non-matching orientations, then, consisted of the set of angles between 0° and 180°, excluding the sample angle and a 45° gap surrounding it. Dividing this range by 48, I obtained a set of non-match angles that varied from one another by about 2.7°, and evenly covered the non-matching orientation space.

Just as each sample orientation value had 48 possible non-matching orientations, each sample color value had 48 possible non-matching colors. As with non-matching

orientations, non-matching colors were easily discriminated from the sample color, and included a large number of perceptually distinct hues. I defined a circular palette of 65 colors using 8-bit HSV color mapping in Matlab (The MathWorks, Natick, MA). By using the maximal saturation and intensity values and varying only the hue, I was able to adapt the procedure I had used for selecting non-matching orientations to selecting non-matching colors: For each sample hue, I excluded the eight hues to either side of it on the color wheel from the set of possible non-matching probe colors. The remaining 48 hues made up the non-match set for that sample color. I recognized, however, that while a fixed angular difference between two orientations is always perceived as the same distance, no matter what the specific orientation values are, the same is not true for fixed angular differences between hues on a color wheel. For instance, it might be possible for a subject to discriminate a pair of hues near the red or blue portions of the color wheel, but not two hues near the yellow or green portions of the color wheel that are equally far apart. I considered doing psychophysical testing to determine the monkey's discrimination thresholds empirically, but was concerned that such testing might have caused her to become more experienced at color discrimination than orientation discrimination, with possible consequences for the respective attentional demands of the two tasks. Instead, I visually inspected each set of non-matching probe hues and shifted the set boundaries as needed so that the non-matching hues closest to the sample value would still be perceptually distinct. This procedure assumed that my color discrimination capabilities would be a reasonable model for those of the monkey. Research comparing human and macaque photopic sensitivity and hue discrimination thresholds provides support for this assumption (De Valois and Jacobs, 1968; De Valois et al., 1974), and the

monkey's subsequent performance on the color-matching task indicated that she could easily discriminate these non-matching colors from the sample.

Thus, for any sample object, there were always 49 possible matching probe objects (a total that includes probe objects that were identical to the sample), each of which had a distinct irrelevant attribute value. My intention in making the irrelevant attribute value so difficult to predict was to encourage the monkey to attend to, and retain in working memory, only the value of the relevant attribute.

A key feature of the experiment was that the conditions used in each DMSA block were exactly the same, and only the preceding cue instruction trials indicated which attribute was relevant. When both of the probe attributes matched the sample or neither of them did, the correct behavioral response was the same for both attribute-matching rules (Fig. 3). I referred to these trials as congruent trials. When one of the probe attributes matched the sample and the other did not, the correct behavioral response differed depending on which attribute was relevant. These trials were labeled incongruent trials, because the behavioral responses suggested by each attribute dimension were incongruent (Fig. 3). During the experiment, all these trial conditions were randomly interleaved and balanced to provide approximately equal numbers of congruent match, congruent non-match, incongruent match, and incongruent non-match trials for each of the two matching rules.

The lack of cues during the DMSA trial indicating which attribute was relevant increased the difficulty of the task. I sometimes noted during training that although the monkey performed certain blocks extremely well, on other blocks she appeared to become confused, making several errors in a row. To keep the performance for each

matching task as consistent as possible across trials and blocks, a single instruction trial was run immediately after the monkey made a mistake in order to remind her which attribute was relevant. These "hint" instruction trials used the same sample stimuli as the "cue" instruction trials; they differed only in that they were presented during the DMSA block rather than at the beginning. After a hint trial, DMSA trials resumed.

Data Analysis

Behavior

My analysis of the monkey's behavior had two objectives: First, I wanted to verify that the monkey had mastered the task, and to rule out the possibility that she might have achieved success by guessing or using an unorthodox behavioral strategy. Second, I wanted to determine whether the monkey's performance was impaired on incongruent DMSA trials, and/or enhanced on congruent DMSA trials.

To address the first objective, I compared the monkey's accuracy in each recording session to that expected by chance (i.e., 50% correct, because each trial had two response alternatives), using binomial tests with $\alpha = 0.01$. In addition, I compared her accuracy on just the congruent and incongruent subsets of DMSA trials to the 50% "guessing" rate. Failure to perform above this rate on incongruent trials would indicate that the monkey had failed to attend to the relevant attribute. If, by contrast, the monkey performed significantly better than chance for incongruent but not congruent conditions, it would suggest that she had adopted an incorrect strategy of attending to one attribute and alternating between two different stimulus-response mapping rules: match-to-sample and non-match-to-sample. For example, suppose the monkey had paid attention *only* to the sample object's color, and had used a match-to-sample rule when color was relevant and

a non-match-to-sample rule when orientation was relevant. She would then be expected to perform well on incongruent conditions because the non-match-to-sample-color rule would produce the same behavioral output as a match-to-sample-orientation rule. When color was the relevant sample attribute, she would also perform well on congruent conditions; but when orientation was the relevant attribute, she would perform poorly on congruent conditions. If the monkey performed better than chance for both congruent and incongruent DMSA conditions, however, and did so consistently from one block to the next (as well as from one experiment to the next), it would support my contention that the monkey had mastered the task and performed it using the desired behavioral strategy (i.e., by alternately matching sample color and sample orientation).

To determine whether the monkey exhibited any pre-potent response tendencies, such as a preference for releasing the bar when the probe appeared, I aggregated the trial outcomes over all the experiments and computed the monkey's accuracy on match and non-match conditions. I compared these percentages using two-tailed χ^2 tests for 2 x 2 contingency tables, for which the square root of the χ^2 statistic is compared to the 0.005 quantile of the standard normal distribution (Conover, 1999)[1]. In an effort to detect any attentional bias that might have favored one attribute over the other, I used the same test procedure to compare performance on DMSA trials during color-matching and orientation-matching blocks.

To determine whether the monkey's performance was impaired on incongruent DMSA trials or enhanced on congruent DMSA trials, compared to instruction trials, I

[1] In every case where I made a statistical comparison of two proportions, I used the χ^2 method for 2 x 2 contingency tables described in (Conover, 1999), because the test statistic is simpler to calculate. For 2 x 2 contingency tables, the number of degrees of freedom is always one. For large sample sizes, Conover recommends using a normal approximation, so the resulting test statistic is z rather than χ^2. All subsequent references to the χ^2 test in which two proportions were compared used this procedure.

analyzed accuracy and mean reaction times. Reaction times for "match" conditions were calculated as the time from the onset of the first choice stimulus (the probe) until bar release. Reaction times for "non-match" conditions were calculated as the time from the onset of the second choice stimulus until bar release. Over many weeks of training prior to recording, I used t tests ($\alpha = 0.01$) to verify that average reaction times were significantly faster for non-match conditions than for match conditions. This finding had been anticipated due to the distinct cognitive demands these conditions placed on the monkey at the moment the matching choice stimulus appeared. On non-match trials, the monkey could reliably expect that the second choice object would match the sample, and could therefore release the bar as soon as she detected this object. On match trials, however, the decision to release the bar was always preceded by an evaluation of the probe object. Given these different behavioral demands, I performed separate analyses for match and non-match conditions.

Comparing DMSA trials to instruction trials allowed me to assess the effect of the irrelevant attribute on accuracy. I treated the monkey's performance on instruction trials as a baseline, reasoning that the sample objects used in instruction trials had only one attribute available for comparison with the probe. For each set of conditions, I aggregated trial outcomes as before and computed the percentage of correct trials for congruent DMSA trials, instruction trials (including both cue and hint trials), and incongruent DMSA trials. To first establish whether there was *any* effect of congruency on accuracy, I compared the monkey's performance on congruent DMSA trials and incongruent DMSA trials using the χ^2 test. If such an effect were found, it would suggest facilitation by the irrelevant attribute during congruent trials, interference from the

irrelevant attribute during incongruent trials, or both. Using χ^2 tests, I looked for evidence of facilitation by comparing accuracy for congruent DMSA and instruction trials, and for evidence of interference by comparing incongruent DMSA and instruction trials.

I used the same approach to analyze the effects of the irrelevant attribute on reaction times: First I compared congruent and incongruent DMSA trials, and then I looked for specific evidence of facilitation or interference by comparing each set of DMSA conditions with instruction trials. Reaction times for each subset of conditions were calculated using correct trials only. Trial data were aggregated over all recording sessions, and comparisons were made using two-tailed *t* tests.

In principle, the cognitive demands of the task were the same regardless of which sample attribute was relevant, and I had no reason to believe *a priori* that an influence of the irrelevant attribute on behavior would depend on which particular attribute was irrelevant. Nevertheless, I ran the analyses described above on separate subsets of color-matching and orientation-matching conditions to see whether differences in the pattern of putative effects would be apparent.

Neuronal activity

The two peri-stimulus time histograms (PSTH) in Fig. 4 show how one VLPFC cell's mean firing rate was modulated by task events. In these histograms, trials are aligned on the time of sample onset. Match trials are plotted in the top PSTH; non-match trials, in which the monkey must wait for the second choice stimulus before releasing the bar, are plotted separately on the bottom. These histograms were generated and inspected for

every cell analyzed. They allowed me to assess visually the cell's range of firing rate modulation and its responsiveness to task events. By comparing the two histograms after the time of probe onset, I could see whether the monkey's behavioral decision had any effect on the cell's activity. Also, I could check for excessive response variability by comparing the two histograms up until the point of probe onset.

I. Definition of task periods

The histograms in Fig. 4 illustrate how trials were subdivided into separate time periods for analysis. The 1500 msec before the onset of the fixation point was taken as the inter-trial interval (ITI). A 600-msec fixation period began when the monkey acquired fixation. To allow for visual response latency, the 500-msec sample period was defined as beginning 100 msec after sample onset, and lasting until 100 msec after sample offset. The delay period began 100 msec after sample offset and lasted 900 msec, until the time of probe onset.

Because the probe was visible on-screen for a longer period of time during non-match trials than during match trials, I distinguished between "early probe" and "late probe" periods. The early probe period represented the cells' initial response to the probe object, and began at the time of probe onset. For each cell, I determined the start and end times of the early probe period as follows: I calculated the monkey's average reaction time to release the bar during congruent match conditions, which I expected to produce the fastest bar-release responses. The early probe period ended 50 msec before this mean bar-release time. Across all the cells I analyzed, this time period ranged from 234 msec to 311 msec in duration. During non-match trials, the cells continued to respond from the end of the early probe period until the onset of the second choice object several hundred

msec later. I defined this period, which included the 200-msec delay period between the two choice objects, as the late probe period. The duration of the late probe period ranged from 689 msec to 766 msec. Also for non-match trials, an additional period was defined that began at the onset time of the second choice object, and lasted until one SD before the average congruent non-match reaction time. The duration of the second choice object period ranged from 106 msec to 188 msec.

Activity during the bar-release period was averaged over a 250-msec period beginning 50 msec before the monkey released the bar. The pre-release interval was restricted to 50 msec in an attempt to exclude transitory visual responses evoked by the onset of the second choice stimulus on non-match trials. Activity during the post-release period was averaged over an 1800 msec period beginning 200 msec after bar release.

For statistical analyses of neuronal activity, each task period was analyzed independently. Only correct trials were used for these analyses, and all tests were evaluated at $\alpha = 0.01$ unless otherwise noted.

II. Evaluation of responsiveness

The cell whose activity is illustrated in Fig. 4 clearly showed different levels of activation during different task periods. To quantify these effects, I compared each cell's mean firing rate during the fixation (baseline) period to its mean firing rate during each of the other task periods using paired, two-tailed *t* tests. Because there were nine task periods to evaluate, the criterion level used for each test was $\alpha = 0.001$; this ensured that the family-wise error rate for all nine tests was less than 0.01. These analyses included all correct match and non-match trials, including instruction trials. If a cell's firing rate during any task period was significantly different from the baseline rate, it was classified

as task-responsive. Within each epoch, a cell's response was classified as excitatory if it was significantly greater than baseline, and inhibitory if it was significantly less than baseline.

For each cortical area, I compared the proportions of excitatory and inhibitory responses during each epoch using binomial tests. χ^2 tests compared the percentage of task-responsive cells and the proportions of excitatory and inhibitory responses in ACC and VLPFC.

Preliminary comparisons of PSTHs for different subsets of trials suggested that some cells responded more robustly to sample stimuli during instruction trials than during DMSA trials (e.g., Fig. 5). This finding might have indicated a preference for instruction sample stimuli, or it might have reflected the relative frequency with which they were viewed. Alternatively, it might have resulted from an enhanced response during "hint" trials, which used the same sample stimuli as the block-cueing instruction trials. To address this issue, and to determine whether hint trial and cue trial responses differed, I compared neuronal firing rates during hint, cue, and DMSA trials for the sample and delay periods. Because DMSA trials were the most numerous trial type and hint trials were by far the least numerous, I used a non-parametric test (Wilcoxon rank sum test, $\alpha = 0.0033$) that made no assumptions about the underlying distributions of mean firing rates. Three rank sum tests (family-wise error rate < 0.01) were performed for each task period: DMSA vs. hint, DMSA vs. cue, and hint vs. cue. To ensure adequate sample sizes, I performed statistical comparisons involving hint trials only if at least 20 correct hint trials had been obtained during the period of stable recording for that cell. The proportion of VLPFC and ACC cells showing significant effects of trial type was compared for each

task period using χ^2 tests. For each cortical area and task period, I also calculated the frequency of particular significant contrasts to compare the overall pattern of effects. The 'DMSA' category included cells that showed significantly higher firing during DMSA trials for a particular task period; 'hint' and 'cue' categories included cells that showed significantly higher firing during hint and cue trials, respectively.

III. Evaluation of selectivity

Because the animal had to retain the relevant matching rule in working memory during the ITI and fixation periods, I tested for rule selectivity during these periods using two-tailed *t* tests. During the sample and delay periods, two additional factors were of interest besides the matching rule: the color and orientation of the sample stimulus. For both of these periods, I performed a three-way ANOVA on matching rule, sample color, and sample orientation. For the early probe period, bar-release period, and post-release periods, I performed three-way ANOVAs with matching rule, response condition (i.e., match or non-match), and congruency as factors. Because the late probe and second choice object periods were defined only for non-match trials, I performed two-way ANOVAs on matching rule and congruency for these epochs. I used a two-tailed *t* test to compare the average firing rate during the late probe period of non-match trials to that of a "post-probe" period for match trials, which I specified as starting at the end of the early probe period and lasting 50 msec.

I analyzed overall cell activity and selectivity for sample attributes during instruction trials in addition to DMSA trials, in order to determine whether effects were similar across trial types. Selectivity for sample attributes during instruction trials was assessed using a two-way *t* test on sample color (for color-matching instruction trials) or sample

orientation (for orientation-matching instruction trials), and the presence of neuronal selectivity and its conformity to the attribute value preferences seen in DMSA trials was noted.

Results

Behavioral performance

The monkey performed this cognitively demanding task extremely well (Table 1), achieving accuracy well above chance (50% correct) on all groups of conditions (multiple binomial tests: $z > 58, p < 0.01$). As expected, the monkey's reaction time was significantly slower on match trials than on non-match trials (compare Fig. 6B and Fig. 7B). For instruction trials (t test, match vs. non-match: $t = 85.9, p < 0.01$), the mean reaction time difference was 146.4 msec. For DMSA trials (t test, match vs. non-match: $t = 142.6, p < 0.01$), the mean reaction time difference was 161.4 msec. As noted earlier, these mean reaction time differences reflected the different cognitive demands that were operative at the time the stimulus matching the sample was displayed. I, therefore, performed separate analyses for match and non-match conditions.

Match conditions

On DMSA trials, the monkey performed color matching and orientation matching equally well (binomial test: $z = 0.8, p > 0.01$), suggesting that she was not biased for either attribute when comparing probe to sample (Fig. 6A). On instruction trials, the monkey did slightly better identifying matching orientations than matching colors

(binomial test: $z = 5.1, p < 3.8 \times 10^{-7}$), but the difference in accuracy was small (2.5%; Fig. 6A). While attribute relevance had a minimal effect on accuracy, it had a strong effect on reaction times (Fig. 6B). When color was the relevant attribute, the monkey took 32.3 msec longer to respond to the probe on instruction trials (t test, $t = 15.0, p < 5.0 \times 10^{-49}$), and 29.6 msec longer on incongruent DMSA trials (t test, $t = 11.3, p < 5.8 \times 10^{-29}$). Even on congruent DMSA trials, in which the probe was identical to the sample, the monkey took significantly longer (7.2 msec) to respond when color was relevant (t test: $t = 4.8, p < 1.6 \times 10^{-6}$).

On incongruent DMSA trials, interference by the irrelevant attribute produced striking deficits in both accuracy and reaction time. Compared to her near-perfect performance on instruction trials (98.4% correct), the monkey's performance on incongruent trials was significantly worse, only 84.0% correct (binomial test: $z = 18.6, p < 0.01$; Fig. 6A). Not only did she make more mistakes on incongruent trials, she also responded more slowly (Fig. 6B). On instruction trials, her mean reaction time was 344.3 ± 0.04 msec, whereas on incongruent DMSA trials, it was 383.7 ± 0.05 msec, nearly 40 msec longer (t test: $t = 22.1, p < 8.9 \times 10^{-104}$). To determine whether the reaction time deficits due to attribute relevance and to trial type were independent, I performed an ANOVA on reaction times using these two factors. Strong effects were observed for each factor alone, but there were no significant interactions between them ($F = 0.61, p = 0.44$).

Whereas the irrelevant attribute produced interference on incongruent DMSA trials, it produced facilitation on congruent DMSA trials. The monkey's excellent performance on instruction trials (98.4% correct; Table 6A) left little room for improvement in

accuracy (congruent DMSA trials, 99.1% correct; binomial test: $z = 2.5, p > 0.01$), but her average reaction time was 20.4 msec faster (t test: $t = 15.7, p < 5.1 \times 10^{-54}$).

Non-match trials

Interference from the irrelevant attribute caused the monkey to make more errors on matching incongruent DMSA trials, and the same proved to be true for non-matching incongruent DMSA trials (Fig. 7A). Compared to non-matching instruction trials, the monkey was significantly less accurate (binomial test: $z = 12.2, p < 0.01$). By contrast, she made significantly fewer errors on congruent DMSA trials than on instruction trials (binomial test: $z = 5.2, p < 2.0 \times 10^{-7}$), an enhancement of performance accuracy that I had not observed for match trials. Examining the non-matching instruction trials more closely, I found that the monkey's performance was nearly perfect when orientation was relevant (98.1%), but only 92.2% when color was relevant (Fig. 7A). Because of the ceiling effect for orientation-relevant trials, no facilitation of performance accuracy was seen (binomial test: $z = 1.2, p = 0.24$). For color-relevant trials, however, the effect was highly significant (binomial test: $z = 7.0, p < 2.1 \times 10^{-12}$).

It would be reasonable to expect the monkey's reaction times on correct non-match trials always to be the same, regardless of task condition, because the second choice object always matched the sample. Instead, I found small but significant differences in average reaction time depending, as for match trials, on attribute relevance and congruency. Although match trials revealed a reaction time *increase* when color was relevant, non-match trials showed a small reaction time *decrease* (Fig. 7B). Only for incongruent DMSA trials was this decrease significant (t test: $t = 7.2, p < 5.5 \times 10^{-13}$); on

these trials, the monkey's response was 15 msec faster when color was relevant. The other trial types, however, showed a similar trend: a 5.8 msec decrease in mean reaction time on instruction trials, and a 3.2 msec decrease on congruent DMSA trials (Fig. 7B).

Just as for match trials, reaction times on congruent DMSA trials were significantly faster compared to instruction trials (t test: $t = 6.7, p < 2.5 \times 10^{-11}$). Unlike match trials, however, reaction times did not increase on incongruent DMSA trials. Indeed, when color was relevant, the monkey's responses on incongruent DMSA trials were significantly faster than on instruction trials (t test: $t = 3.8, p < 1.6 \times 10^{-4}$; Fig. 7B).

Neuronal Responses

Despite the fact that I did not specifically seek out neurons with task-related responses, I found that 107 of 114 VLPFC cells (94%) and 79 of 80 ACC cells (99%) were responsive to task events such as stimulus presentation and bar release (paired, two-tailed t tests for each task epoch: $p < 1.0 \times 10^{-3}$). Cortical areas did not differ significantly in terms of overall responsiveness (χ^2 test: $z = 1.7, p = 0.09$) or in responsiveness during any single task period when epochs were compared individually (multiple binomial tests: $p > 0.48$ for all). In each task epoch, roughly half of the responsive cells had firing rates below the baseline rate (Fig. 8). Binomial tests comparing the proportion of excitatory and inhibitory responses in each epoch showed no significant differences for either VLPFC or ACC. I found no between-area differences in the proportion of cells with firing rates above baseline for any task epoch (χ^2 tests: $p > 0.24$ for all).

In addition to comparing average firing rates in each task period to the baseline firing rate, I compared the effect of trial type (DMSA trial, hint trial, or cue trial) on neuronal responsiveness during each task period. To ensure an adequate number of hint trials for these comparisons (see *Methods*), I restricted the analysis to 103 VLPFC cells and 68 ACC cells. The proportion of these cells whose responses were modulated by trial type is shown in Fig. 9. In VLPFC, significant effects of trial type occurred most frequently during the sample, delay, and probe periods, but for ACC these effects were also commonly seen during the ITI and fixation periods.

I compared areas with respect to the frequency of particular significant contrasts. Only seven VLPFC cells and eight ACC cells showed any significant differences between hint and cue trials in any task period (binomial tests: VLPFC, $p = 0.17$; ACC, $p = 0.05$). Hint trials occurred immediately after errors, and error-related signals have been proposed to represent a mechanism for sharpening attention and executive control. Such signals did not appear to be prevalent in either cortical area.

The most commonly occurring firing rate difference in VLPFC was between hint/cue trials and DMSA trials during the sample and delay periods. In nearly all of these cases (Fig. 10), firing rates were higher during hint/cue trials (χ^2 tests: sample period, $z = 4.6$, $p < 3.4 \times 10^{-6}$; delay period, $z = 3.3$, $p < 8.6 \times 10^{-4}$). In ACC, however, responses during the sample and delay periods of DMSA trials were as likely to be enhanced as suppressed compared with hint and cue trials (χ^2 tests, $p = 0.08$).

Attribute Value Selectivity

Many cells in both VLPFC and ACC were selective for particular sample attribute values during one or more task epochs. Using the criteria described above, I classified 80 VLPFC cells (70%, $N = 114$) and 30 ACC cells (38%, $N = 80$) as attribute-selective during the sample or delay periods. I found significantly more attribute-selective cells in VLPFC cells than in ACC (χ^2 test: $z = 4.5, p < 6.2 \times 10^{-6}$).

Of the 80 attribute-selective VLPFC cells, 69 (86%) were selective for sample orientation, and 42 (53%) were selective for color. An example of an orientation-selective cell is shown in Fig. 11, and an example of a color-selective cell is shown in Fig. 12. The proportion of orientation-selective cells in VLPFC was significantly higher than the proportion of color-selective cells (χ^2 test: $z = 4.6, p < 3.7 \times 10^{-6}$). When instruction and DMSA trials were analyzed separately, higher percentages of orientation-selective cells were found in both cases. Because of the properties of the sample stimuli used in the three condition groups (color-matching instruction trials, orientation-matching instruction trials, and DMSA trials), analyzing instruction conditions required comparing selective responses during two different sets of trials, while analyzing DMSA trials entailed examining the effects of two different factors on responses during a single set of trials. During instruction conditions, 49 of 114 VLPFC cells (43%) were selective for sample attributes. Significantly more of these cells (41 of 49, or 84%) were selective for sample orientation than for sample color (16 of 49, or 33%; χ^2 test: $z = 5.1, p < 3.1 \times 10^{-7}$). During DMSA conditions, 69 of 114 VLPFC cells (61%) were selective for sample attributes. Of these cells, 58 (84%) were selective for sample orientation while just 35 (51%) were selective for sample color; as was the case with instruction conditions, this

bias for representing sample orientation was statistically significant (χ^2 test: $z = 4.2$, $p < 3.0 \times 10^{-5}$).

Of the 69 VLPFC cells that were attribute-selective during DMSA trials, 26 (38%) were selective for both attributes during the same task epoch, but only 7 (10%) showed an interaction between sample color and sample orientation in the three-way ANOVA. Fig. 13A shows one such cell, which was selective for a particular sample item during both the sample and delay periods. For the most part, sample objects appeared to be represented in VLPFC by two essentially independent neural codes, one for each attribute dimension. An example of a cell that was both color-selective and orientation-selective but showed no interaction between the two factors is shown in Fig. 14.

In ACC, 21 of 80 cells (26%) were orientation-selective, and 19 (24%) were color-selective. In addition to having fewer attribute-selective cells, ACC differed from VLPFC by showing no bias for sample orientation over sample color (χ^2 test: $z = 0.5$, $p = 0.60$). Of the 25 ACC cells that were attribute-selective during DMSA trials, 9 (36%) were selective for both attributes during the same task epoch, but only 2 of these featured a significant interaction between sample orientation and color in the three-way ANOVA, and could thus be considered item selective (Fig. 13B). An additional 2 cells featured a significant interaction among all three factors, such that their optimal stimulus conditions were impossible to classify in terms of simple attribute value preferences. For the other 5 ACC cells that were selective for both sample attributes during the same task period, I found no interaction between orientation-selectivity and color-selectivity. Thus in both cortical areas, over 80% of the attribute-selective cells appeared to maintain independent representations of the two DMSA sample attributes.

To assess whether attribute-selectivity was driven visually or by the working memory demands of the task, I compared cells' activities during the sample and delay periods. In both VLPFC and ACC, a majority of attribute-selective cells made differential responses while the monkey viewed the sample object onscreen. In VLPFC, 47 of the 69 orientation-selective cells (68%) and 27 of the 42 color-selective cells (64%) were selective during the sample period of instruction trials, DMSA trials, or both. In ACC, 15 of the 21 orientation-selective cells (71%) and 12 of the 19 color-selective cells (63%) were selective during the sample period.

In VLPFC, a majority of attribute-selective cells were also selective when information about the sample object had to be maintained in working memory. Of the 69 orientation-selective cells and 42 color-selective cells, 51 (74%) and 27 (64%), respectively, were selective during the delay period. In ACC, smaller but still substantial fractions of the attribute-selective cells were selective during the delay period: 13 of the 21 orientation-selective cells (62%), and 9 of the 19 color-selective cells (47%). For each cortical area, the total number of cells with attribute-selective responses was the equivalent for both task periods: In VLPFC, 60 cells were attribute-selective during the sample period, and 58 were attribute-selective during the delay period (χ^2 test: $z = 0.4$, $p = 0.70$), while in ACC, 20 cells were attribute-selective during the sample period, and 19 were attribute-selective during the delay period (χ^2 test: $z = 0.3$, $p = 0.80$). This result suggests that in both areas, information about sample attributes was represented as effectively when the sample object was absent as when it was present.

To determine whether presumed working memory activity reflected the sustained activation of attribute-selective responses to the sample object, I compared cell

preferences during the sample and delay periods. In VLPFC, 24 cells were orientation-selective during both periods, and all but 4 had the same preferred orientation during the delay period as during the sample period. Thus, 20 of the 51 (39%) orientation-selective "memory" cells in VLPFC behaved as though they were simply continuing to respond to the sample object. By contrast, 11 of the 42 (26%) color-selective "memory" cells in VLPFC were also color-selective during the sample period, of which only 7 (17%) had the same preferred color during both periods. Thus, the proportion of delay-period activity in VLPFC that could be interpreted as a sustained response to the sample object was slightly higher for orientation-selective responses than for color-selective responses (χ^2 test: $z = 2.4$, $p = 1.7 \times 10^{-3}$). In ACC, 5 of the 7 cells that were orientation-selective during both task periods maintained the same preferred orientation, a ratio similar to that observed in VLPFC (χ^2 test: $z = 1.2$, $p = 0.20$). None of the 5 ACC cells that were color-selective during the delay period was also color-selective during the sample period.

Because sample orientation and color were, for the most part, encoded independently during DMSA trials, I wondered whether these preferences were generalized and whether, therefore, they would apply to any stimulus of the appropriate orientation or color. If so, then the sample objects presented during instruction trials should elicit similarly selective responses. Only 30 of the 69 of the orientation-selective VLPFC cells (43%) and 10 of the 42 of the color-selective VLPFC cells (24%) exhibited selective responses during both instruction trials and DMSA trials. In ACC, only 4 of the 21 orientation-selective cells (19%) and 3 of the 19 color-selective cells (16%) showed selective responses during both types of trial. These results indicated that most attribute-selective cells did *not* respond in a generalized fashion; instead, other factors (e.g., shape,

size, or block type) modulated their activity. Among the sub-population of cells that were attribute-selective during both instruction and DMSA trials, however, most preferred the same attribute value. In VLPFC, only 4 cells preferred different orientations during instruction and DMSA trials, and only 1 cell preferred different colors. The rest (26 orientation-selective cells and 9 color-selective cells) consistently preferred the same attribute. By contrast, only 3 ACC cells consistently preferred the same orientation in instruction trials as in DMSA trials, and only 1 ACC cell preferred the same color. Because neuronal representations of each attribute were almost entirely independent, I additively combined attribute response profiles for statistical comparison. Significantly more VLPFC cells preferred the same attribute value in both instruction trials and DMSA trials (χ^2 test: $z = 3.0, p < 3.0 \times 10^{-3}$), indicating that categorical responses were more common in VLPFC than ACC.

Effect of attribute relevance on attribute value selectivity

During DMSA trials, the same sample stimuli were used in color-matching and orientation-matching blocks, allowing me to assess the affect of attribute relevance on attribute selectivity in VLPFC and ACC. If single neurons only represented information about an attribute when that attribute was relevant, I expected to find significant interactions between the factors of matching rule and sample orientation or sample color in a three-way ANOVA. By contrast, if cells represented information about sample attributes even when that information was *not* relevant, then interactions between the matching rule and sample attribute values would not be significant.

Whereas a large proportion of cells in both VLPFC and ACC was attribute-selective according to the ANOVAs, very few cells in either area showed an interaction between attribute selectivity and matching rule. In VLPFC, I found only 3 cells with a significant Rule x Sample Color effect, and only 7 cells with a significant Rule x Sample Orientation effect. In ACC, only 2 cells showed a significant Rule x Sample Color effect, and only 1 cell had a significant Rule x Sample Orientation effect. Additionally, 2 ACC cells featured a significant interaction of all three factors. In neither area, therefore, did I find that only *relevant* information about sample attributes was represented in cells' activity (binomial tests: ACC, $p = 0.52$; VLPFC, $p = 0.15$).

Effects of matching rule

The use of a blocked design required that the monkey remember the current matching rule during each task period of the trial. Prior to the onset of the probe object, knowledge of the current matching rule might have played a role in selectively attending to and remembering the relevant sample attribute. When the probe object was displayed, knowledge of the current rule could have been applied in response selection. Even after a response was made, and the monkey's behavior was rewarded (if correct) or punished (if incorrect), she had to remember the rule for the next trial. I, therefore, tested each task epoch for rule selectivity, and noted how frequently such activity occurred before and after the probe was displayed.

I. Prior to probe display

During DMSA trials, only 6 VLPFC cells and 9 ACC cells showed a main effect of matching rule prior to the onset of the probe object. In VLPFC, 1 cell was rule-selective

during the ITI period, 1 was rule-selective during the sample period, and 4 were rule-selective during the delay period (e.g., Fig. 15A). In ACC, 1 cell was rule-selective during the fixation period (Fig. 15B), 3 cells were rule-selective during the sample period, and 5 cells were rule-selective during the delay period. Interestingly, one of the cells that preferred orientation-matching blocks during the delay period also showed a significant main effect of matching rule during the ITI, but preferred color-matching blocks. The percentage of rule-selective cells was significant in ACC (binomial test: $p < 6.0 \times 10^{-3}$) but not in VLPFC (binomial test: $p = 0.30$).

II. After probe display

Rule effects occurred more frequently in VLPFC after the onset of the probe object. During the task periods that followed probe onset, 36 VLPFC cells showed an effect of rule, and half of those showed a main effect (binomial test: $p < 2.7 \times 10^{-5}$). During the early probe period, 4 cells showed a main effect of matching rule, and 9 cells showed significant interactions between rule and the other two factors (response type and congruency). An example of a cell with a main effect of rule during the early probe period is shown in Fig. 15C. During the late probe period, 6 cells showed a main effect of rule, of which 2 also showed a significant interaction between rule and congruency. One additional cell did not show a main effect of rule, but did show an interaction effect. During the display of the second choice object, 2 cells showed a main effect of rule, and 3 cells showed a significant interaction between rule and congruency. During the bar-release period, 7 cells showed a main effect of rule and 11 cells showed significant interactions between rule and the other two factors (one of the 11 also had a significant main effect of matching rule). During the post-release period, 3 cells showed a main

effect of rule and 2 cells showed significant interactions between rule and the other factors. Among the cells with a main effect of rule, half preferred orientation matching and half preferred color matching.

I found a main effect of matching rule during one of the task periods following the onset of the probe object in 11 ACC cells (binomial test: $p < 8.0 \times 10^{-3}$). During the early probe period, 9 cells showed an effect of rule, of which 3 had a main effect. During the late probe period, 4 cells showed an effect of rule, of which 2 had a main effect. During the display of the second choice object, only 1 cell was rule-selective. During the bar-release period, 8 cells showed an effect of rule, of which 2 had main effects. An ACC neuron with a significant Rule x Decision interaction is shown in Fig. 15D. During the post-release period, 9 cells had significant effects of the matching rule, of which 4 showed main effects. Of the 11 ACC cells with a main effect of rule, 8 preferred orientation matching and 3 preferred color matching. I found no significant difference between the proportions of ACC and VLPFC cells selective for rule after probe onset (χ^2 test, $z = 0.2$, $p = 0.81$).

Effect of congruency

Over one third of VLPFC cells (47 of 114) showed an effect of congruency during the early probe period, late probe period, second test object period, bar release period, or post-release period. Of those cells, 21 (18%) showed main effects during one or more of those task periods (binomial test: $p < 5.4 \times 10^{-7}$). During the early probe period, 11 cells showed congruency effects, of which 3 had a main effect. During the late probe period, 12 cells were congruency-selective, and 9 of those showed a main effect. During the

second choice object period, 7 cells showed an effect of congruency, of which 4 had main effects. During the bar-release period, 21 cells showed an effect of congruency, most of which had a significant interaction between congruency and response type (match or non-match). Only 2 cells showed a main effect of congruency during the bar-release period. During the post-release period, 9 cells showed congruency effects, and 4 of those had a main effect of congruency. Fig. 16A shows an example of a congruency-selective VLPFC cell.

In ACC, 25 of 80 cells (31%) showed a significant effect of congruency, and 12 of those cells had a main effect (binomial test: $p < 8.5 \times 10^{-4}$). Specifically, 8 cells showed an effect during the early probe period, 7 during the late probe period, 2 during the second choice stimulus period, 11 during the bar-release period, and 3 during the post-release period. Fig. 16B shows an example of a congruency-selective ACC cell.

I found no significant differences between VLPFC and ACC in the percentage of cells with congruency effects (χ^2 test, $z = 0.9$, $p = 0.35$). What was striking, however, was that among cells with main effects of congruency, those in VLPFC generally responded more robustly during *incongruent* conditions (62%), whereas those in ACC overwhelmingly preferred *congruent* conditions (83%). This difference came close to reaching significance (χ^2 test, $z = 2.5$, $p = 1.2 \times 10^{-3}$), and might have done so had the sample sizes been larger.

For cells with a main effect of congruency, I calculated effect sizes using Cohen's f (Kirk, 1995)[2]. Effect sizes ranged from 0.10 to 0.36 in VLPFC, and from 0.08 to 0.31 in

[2] The square of Cohen's f is equal to the treatment variance divided by the error variance. The square root of this ratio, f, is therefore a measure of effect size. Cohen considered an f of 0.10 to be a small effect, an f of .25 to be a medium-sized effect, and an f of 0.40 to be a large effect (Kirk, 1995). The formula for converting the effect size f into the proportion of the total variance explained by the treatment is $f^2/(1+f^2)$.

ACC. In both areas, the strongest effects were observed during the late probe and second choice object periods. In VLPFC, effect sizes during these two task periods ranged from 0.14 to 0.36, and the mean value for *f* was 0.24; by contrast, the average *f*-value during the early probe period was 0.12, and during the bar-release and post-release periods it was 0.11. In ACC, effect sizes during the late probe and second choice object periods ranged from 0.18 to 0.31. As in VLPFC, the mean value for *f* during these periods was 0.24, whereas it was 0.12 during the early probe period and 0.13 during the bar-release and post-release periods.

Effect of response choice

A large number of cells in both VLPFC and ACC showed effects of the monkey's behavioral choice in their activity. Some cells showed a stronger visual response to the probe object, depending on whether or not it matched the sample (Fig. 17A). Many cells showed increases in their firing rates during the late probe period of non-match trials. Some of these cells produced a strong burst of firing that terminated before the offset of the probe object (Fig. 17B), while others displayed a gradual ramping up of their firing rates (Fig. 17C-D). Several cells displayed different levels of activity immediately preceding bar release, depending on whether the release occurred after the probe object or the second choice object. I, therefore, found it useful to analyze the effect of response type during all of the task periods that came after the onset of the probe object.

In VLPFC, 59 of the 114 cells (52%) showed an effect of response type during the early probe period, bar-release period, or post-release period. Using *t* tests to compare

---

Thus, a small effect ($f = 0.10$) explains about 1% of the variance, a medium-sized effect ($f = 0.25$) explains about 6% of the variance, and a large effect ($f = 0.40$) explains about 14% of the variance.

the average firing rate during the late probe period of non-match trials to the average firing rate immediately after the early probe period of match trials (see *Methods*), I identified an additional 18 VLPFC cells with significant decision-related responses. Of the 59 VLPFC cells I analyzed using ANOVAs, 47 (80%) showed a main effect of response type. Throughout all task periods analyzed, VLPFC cells with main effects of response type consistently preferred non-match trials: During the early probe period, only 3 of 10 cells showing main effects preferred match trials. During the late probe period, 27 of 41 cells (66%) preferred non-match trials, and 29 of 38 cells (76%) preferred non-match trials during the bar-release period (only 3 cells showed a main effect of response type during the post-release period).

In ACC, 29 of 80 cells (36%) showed an effect of response type during the early probe period, bar-release period, or post-release period. Of these 29 cells, 19 (66%) showed a main effect. An additional 19 ACC cells had a significant effect of response type during the late probe period of non-match trials, according to *t* tests. I found no significant difference between the proportions of decision-related ACC and VLPFC cells (χ^2 test, $z = 1.1, p = 0.28$). As in VLPFC, cells in ACC with a main effect of response type tended to prefer non-match trials during the early and late probe periods. Only 2 of 8 cells showed a preference for match trials during the early probe period, and only 8 of 33 cells (24%) preferred match trials during the late probe period. During the bar release period, however, 11 of the 15 ACC cells (73%) with a main effect of response type preferred match trials. Thus, although the percentage of cells showing a main effect of response type was statistically the same for ACC and VLPFC (χ^2 test, $z = 1.4, p = 0.17$),

the proportions favoring match and non-match trials were significantly different (χ^2 test, $z = 3.4, p < 7.9 \times 10^{-4}$).

Discussion

In this experiment, I attempted to confirm and extend the hypothesis that VLPFC neurons preferentially represent relevant information (Rainer et al., 1998). I also attempted to learn whether the representation of relevant information and the representation of categorical membership described by Rainer, Freedman, and co-workers (Freedman et al., 2001, 2002; Rainer et al., 1998) were manifestations of a general mechanism in VLPFC for directing attention selectively to one or another of an object's attributes. Additionally, I evaluated a model of ACC and lateral PFC functioning advanced by Cohen and colleagues (Botvinick et al., 2001, 2004; Cohen et al., 1990; MacDonald et al., 2000) that has garnered support from human functional imaging, but has not been tested extensively using single-unit electrophysiology methods. Below, I review the behavioral and neurophysiological results of this experiment, and discuss their implications for our understanding of VLPFC and ACC functioning.

Behavioral performance

Verifying that the monkey was performing the task using the appropriate strategy was a major concern for this experiment. Because she performed well above chance on both congruent and incongruent trials, and did so consistently from day to day, it seems unlikely that she could have used any other strategy to solve the task. Still, I considered the possibility that she might have memorized all the specific sample/test object combinations for each attribute-matching block. That possibility seems highly unlikely

for the following reasons: First, the task design discouraged memorization because half of the sample/test object pairs required a match response in one block, and a non-match response in the other block. The monkey still might have memorized specific combinations of instruction-trial cues, sample objects, and test objects; but this explanation seems implausible because the instruction trials appeared at the beginning of a block of trials, and so were remote in time from many of the incongruent trials that the monkey performed correctly. Second, I used a large number of sample/test object combinations for each block. For each sample object, there were 145 different probe objects that could follow. Besides making it much more difficult to memorize all the conditions, this approach also encouraged the monkey to use a retrospective memory strategy (remembering the relevant sample attribute during the delay) rather than a prospective memory strategy (anticipating the matching test object or a small set of matching objects). One benefit of encouraging the monkey to adopt a retrospective memory strategy was that it made the cognitive demands during the sample and delay periods equivalent for both types of attribute-matching block. Third, each day I chose two new sample orientations and two new sample colors to generate a novel set of four sample stimuli. During the animal's training, I verified that she performed as well with novel sample sets as with familiar ones, with no evidence of a learning curve. For these reasons, I rejected the hypothesis that the monkey had memorized all the task conditions, and concluded that she could only have performed the task by attending to the relevant attribute.

Analysis of the monkey's behavior on match trials revealed two interesting results. First, despite the monkey's unmistakable mastery of the task, her performance was

powerfully and adversely affected by interference from the irrelevant attribute. This interference produced consistent and significant deficits in both accuracy and reaction time. The irrelevant attribute, however, could also have facilitated performance, as I observed when I compared instruction trials (single-attribute comparison) to congruent DMSA trials (dual-attribute comparison). This facilitation was apparent in the monkey's reaction times but not her accuracy, probably because of a ceiling effect.

Second, the monkey required significantly more time to identify matching colors than matching orientations. This result seems unlikely be due to difficulty ignoring the probe's orientation when color was relevant, because I observed an effect of equivalent size during instruction trials, when the sample objects had only one attribute available for comparison with the probe. The reaction-time cost associated with color matching, however, was much smaller (though still significant) for congruent DMSA trials. Because I believe that the irrelevant attribute facilitated performance in these conditions, it is possible that the orientation attribute significantly improved performance in congruent color-matching trials, helping to close the gap. It is also possible that on some fraction of these trials, the monkey was paying attention to the probe's orientation rather than its color, thereby lowering the reaction time average. When I plotted the distribution of reaction times during congruent color-matching trials, however, it was unimodal rather than bimodal, suggesting this explanation was not valid.

On non-match trials, I again observed a significant decrease in response accuracy for incongruent DMSA conditions, and (when color was relevant) an improvement in response accuracy for congruent conditions. These results were consistent with what I found for match trials, but other behavioral metrics were different. Instead of a reaction-

time increase when color was the relevant attribute, I saw a small reaction-time decrease. Also, reaction times were longer on incongruent trials than on congruent trials. These results were somewhat surprising because all the monkey needed to do on non-match trials was detect the second choice object, and she had several hundred milliseconds to prepare. These reaction-time effects must have resulted from prior events in the trial, possibly reflecting inhibitory processes that were deactivated more slowly for certain sets of conditions.

Neuronal responses

Although the monkey had extensive experience with the task prior to recording, her pattern of mistakes and slower reaction times on incongruent conditions suggested that she still found the task challenging. The same was also suggested by the fact that the firing rates of nearly all the cells I sampled (94% in VLPFC, and 99% in ACC) were significantly shifted from baseline during various task epochs. Both cortical areas appeared to be strongly activated by the cognitive demands of the task.

Strong attribute selectivity was found in VLPFC and, somewhat surprisingly, in ACC as well, although attribute-selective responses were less common in ACC than VLPFC. ACC is not a recipient of direct visual cortical afferents, but it shares connections with numerous prefrontal areas, as well as perirhinal cortex (Vogt and Pandya, 1987), that could relay visual information. VLPFC does receive direct visual input from inferior temporal cortex (Webster et al., 1994), and also shares connections with medial temporal cortex and other prefrontal areas. I found significantly more orientation-selective responses than color-selective responses in VLPFC, the reason for which is unclear. A difficult task might be expected to activate more cortical neurons than an easy task, but

the monkey's behavioral data suggest that she was slightly better at matching orientation than matching color.

Approximately one third of the attribute-selective neurons in both areas were selective both for both attributes during the same task period, but very few neurons were selective for specific combinations of sample color and orientation. Instead, it seemed for the most part that information about the two attributes was encoded separately. Two thirds of the attribute-selective cells, then, behaved like the category-selective cells described by Freedman and colleagues (Freedman et al., 2001, 2002) because one set of visual distinctions was discriminated and the other set was ignored. I found little evidence, however, that these categorical representations of attribute value were altered by changes in the relevance of that attribute. Only a handful of cells in either area had significant interaction effects with the matching rule during the sample and delay periods, as measured by three-way ANOVAs. In contrast to the experiment by Rainer and co-workers, in which selective attention to a spatial location changed the neural response to a particular array of objects (Rainer et al., 1998), selective attention to one or another of an objects' attributes did not change the object's neural representation. It appears, then, that sometimes irrelevant information about visual objects *is* represented in lateral PFC. Additional work is needed to specifically determine how and when irrelevant information is excluded from neuronal representations of visual stimuli in PFC.

It is certainly possible that the attribute-selective responses I observed were driven strictly by the visual properties of the sample stimuli, and that high-order cognitive processing involving attribute relevance occurred elsewhere. Many neurons, however, that were selective for a particular color or orientation during the DMSA task were not

selective for stimuli with those same attribute values during instruction trials, and vice versa. I also observed differences in neurons' overall firing rate during DMSA trials and instruction trials, even during the ITI and fixation periods. Thus, factors other than visual stimulus properties did affect these neurons' responses.

During each block of DMSA trials, the monkey had to remember the current matching rule to perform well. Because the monkey's accuracy on incongruent trials was so high, I can assume that this information was maintained in working memory and updated whenever the matching rule changed. I found no evidence of active maintenance of the current matching rule during the ITI and fixation periods in either VLPFC or ACC, however. By contrast, Asaad and co-workers, using a similar blocked experimental design, reported that task-selective baseline activity was common (Asaad et al., 2000). It is possible that the difference between their findings and ours is due to the fact that their blocks were much longer (100-200 trials), or that their lateral PFC recording site was slightly dorsal and posterior to ours, encompassing more of the cortical area around the principal sulcus. In a study by Wallis and colleagues using a DMS task with two matching rules, rule-selective neurons tended to be mapped around the principal sulcus rather than on the surface of the inferior prefrontal convexity (IPC; Wallis et al., 2001). Dorsolateral PFC and possibly pre-motor cortex (Wallis and Miller, 2003) were probably more critical for maintaining information about the matching rule in this task.

During the sample and delay periods, I found a few more cells in each area with a main effect of rule in the three-way ANOVAs. Only in ACC did the number of rule-selective neurons represent a significant fraction of the total; but the overall proportion was still small (10%). Based on these results, and on the finding that attribute relevance

had little effect on attribute selectivity in either area, it seems unlikely that the contributions of VLPFC and ACC to task performance depended greatly on attribute relevance, at least during the sample and delay periods.

Once the probe object appeared after the delay, the monkey had to compare it to the sample object and decide whether or not to release the bar. Knowledge of the current matching rule was essential to that process. In ACC, the percentage of cells having main effects of rule was the same after probe onset as before, but in VLPFC the proportion of rule-selective cells increased significantly (χ^2 test, $z = 2.6, p < 0.01$). Both areas also contained a large number of cells whose activity was strongly modulated by the monkey's response choice; indeed, the strongest task-related effects I observed, in addition to cells' attribute-selectivity, were these decision-related responses. These results suggest a role in response selection for both areas. In particular, it suggests a more active role for VLPFC than simply representing the visual properties of stimuli and maintaining that information on-line. A study by Rushworth and colleagues found that following lesions of the VLPFC that included the posterior portion of the IPC, monkeys were impaired at re-learning a simultaneous match-to-sample task, and adding delays to the task did not degrade performance further (Rushworth et al., 1997). Not only did this indicate that "ventral prefrontal cortex is not essential for working memory," it suggested that VLPFC was important for stimulus comparison and selection, a view that is consistent with other lesion and imaging studies in monkeys and humans (Eacott and Gaffan, 1992; Passingham et al., 2000; Petrides et al., 2002).

As noted earlier, Cohen and colleagues (Botvinick et al., 2001, 2004; Carter et al., 1998; Cohen et al., 1990; MacDonald et al., 2000) have proposed that the lateral PFC

mediates the allocation of selective attention, while the ACC monitors behavior and activates the lateral PFC when it detects a need for increased executive control. If so, I expected that ACC neurons would discriminate between congruent and incongruent conditions, and that they would produce higher rates of firing when incongruent conditions are detected. Further, I expected this discrimination between congruent and incongruent conditions to be a cardinal property of ACC, given its well-documented activation in a variety of response-conflict paradigms. The current study confirms that a significant percentage (31%) of ACC cells were indeed congruency-selective, but an equally large percentage of VLPFC cells were congruency-selective as well, and the range of effect sizes observed was similar for both areas. Moreover, the majority of cells that preferred incongruent conditions were in VLPFC, not ACC. The increase in VLPFC activation during incongruent trials may reflect a role in imposing an attentional 'set' as opposed to resolving conflicting response impulses (Banich et al., 2000; Milham et al., 2001). The finding that ACC tended to be more active during congruent conditions, however, is puzzling and difficult to resolve with most of the imaging and electrophysiological literature (see Ridderinkhof et al., 2004, for review). Interpreting the role of ACC strictly in terms of monitoring response conflict, however, may be taking too narrow a view of its function. It may also be required for the "resolution of prior inhibition" (i.e., when a response that had been actively suppressed must be permitted once again; Ruff et al., 2001). In the DMSA task, there is a 50% chance on any DMSA trial that the trial condition will be incongruent, so the monkey must always be ready to suppress a response to the irrelevant attribute. Of course, because the matching rule changes from one block to the next, the suppression of responses to a formerly irrelevant

attribute will frequently be reversed. I speculate that if a mechanism for re-activating previously inhibited responses exists in the ACC, then perhaps it is sometimes triggered inadvertently by the occurrence of a congruent trial. The tendency to respond to the irrelevant attribute would then be facilitated early, before the matching rule switches. If so, this might produce faster reaction times for congruent trials than instruction trials, which is consistent with the behavior that I observed. A report by Swainson and co-workers has described two foci in ACC that appeared to play different roles in a task-switching paradigm that involved alternately making and withholding responses to visual stimuli (Swainson et al., 2003). The more anterior ACC focus was activated only when the rule switched from 'go' to 'wait,' and thus seemed to be involved in suppressing an automatic response tendency (a condition characteristic of all response-conflict paradigms). The more posterior ACC area was activated only when the rule switched from 'wait' to 'go,' suggesting that it enabled the release of the previously suppressed response. The fact that so many of the cells in ACC preferred congruent conditions may indicate that I recorded from a subdivision of ACC whose function is more closely related to undoing prior inhibition than detecting response conflict. A functional role in undoing prior inhibition might also account for the fact that among ACC cells with a main effect of decision during the bar-release period, the majority had higher firing rates during match trials; by contrast, the majority of decision-selective cells in VLPFC had higher levels of activity during non-match trials.

Conclusions

Neuronal activity in VLPFC and in ACC was strongly modulated by task demands, as evidenced by shifts in firing rates from baseline frequencies in relation to task events, and by the large percentage of neurons in both areas that were selective for sample attribute values, attribute-response congruency, and response choice. Sample attribute dimensions were represented almost entirely independently of one another, even in cells that were selective for both sample orientation and sample color. For VLPFC, these results are consistent with a proposed functional role in comparing visual stimuli and selecting an associated response; the fact that attribute dimensions were represented independently might facilitate making categorical discriminations along one dimension but not the other. Attribute selectivity neither vanished nor diminished during blocks in which the attribute was irrelevant, however, meaning that the allocation of attention to that attribute was either reflected in other brain regions that I did not explore, or that it was accomplished in a different manner than by suppressing neuronal responses to the unattended phenomenon. Others have proposed neuronal synchronization as a mechanism for selective attention (Fries et al., 2001; Tallon-Baudry, 2004); I did not analyze correlations between different neurons, or between spike trains and local field potentials, so I cannot address this possibility. Neither VLPFC nor ACC played a direct role in maintaining the current matching rule in memory, but the frequency of rule-selective responses in VLPFC increased after the probe appeared, which is consistent with the idea that VLPFC analyzes and compares visual stimuli in the context of ongoing task demands and behavioral decision-making.

The role of the ACC in performing this task remains somewhat difficult to interpret, based on the data described here. Although I found fewer attribute-selective cells there than in VLPFC, such cells were by no means rare. The part these stimulus representations play in ACC functioning is not yet clear. Assuming VLPFC is directly involved in comparing stimulus attributes and activating associated responses, I can imagine there might be distinct "go" networks for the various possible attribute matches, and that viewing particular sample objects might pre-potentiate one or more of these networks. In that case, the recruitment of ACC to monitor potential response conflicts might also begin when the sample is presented, and it might reflect the particular response conflicts already suggested by the sample object's attributes, producing "attribute-selective" neuronal activity. While VLPFC may play a more direct role in selecting or inhibiting responses based on visual comparisons, the ACC may contribute to these processes by helping to re-activate previously suppressed responses.

Acknowledgements

I would like to thank all the people who have helped me over the years here in the Department of Brain and Cognitive Sciences at M. I. T. First and foremost, my advisor, Prof. Earl Miller, who has generously supported me these many years, and whose "tough love" was essential in helping me finish this project! I also would like to thank the present and past members of the Miller lab who, together with Earl, trained me in the nuts and bolts of working with monkeys, designing behavioral experiments, and recording and analyzing data; and who have always been willing to give me opinions, suggestions, and advice: Kathleen Anderson, Wael Asaad, Tim Buschman, Cinira Diogo, David Freedman, Mark Histed, Rahmat Muhammad, Andreas Nieder, Anitha Pasupathy, Gregor Rainer, Jefferson Roy, Jonathan Wallis, Melissa Warden, and Richard Wehby. I also have received a lot of support from our former lab manager, Emily Le; our current lab manager, Kristin MacCully; and our computer support guru, Christopher Irving.

I am very grateful to Prof. Matt Wilson and Prof. Suzanne Corkin for serving on my thesis committee and for any and all suggestions they may have to improve this thesis. Sue, in particular, reviewed my thesis proposal very meticulously, and suggested

numerous changes that made it a much better proposal. I very much appreciate the time she took to do that, as I appreciate the time all three of my committee members will take in reviewing this document. Many warm thanks!

The work we do with non-human primates in the Miller lab would not be possible without the invaluable assistance we receive every day from the staff of the Department of Comparative Medicine. I would like to thank Dr. Bob Marini, our unflappable head veterinarian and chief of surgery at the D. C. M., who has had to patch up the subject of this experiment on more than one occasion; Katie Madden, our vet tech, who is always ready to come back into work late at night to take care of emergencies, and who has always been extremely patient with my many concerns and requests for help; Dr. Mary Patterson, whose number-one concern is, has been, and always will be the welfare of the monkeys; and all of the veterinary post-docs and animal caretakers who have taken care of my monkeys over the years. Thank you very much.

Finally, I wish to thank my family, friends, and neighbors, who have given me support and encouragement: my parents, Duane and Suzie; my sisters and brothers-in-law, Jen and Mike, and Kathi and James; and Guillermo Ameer, Amy Ashbacher, Mike Berry, Cinira Diogo, Amy Ehntholt, Mark Felton, Brenda Frye, Gadi Geiger, Joe Kewekordes, Nat Kuhn, Cathy Leamey, Heath Lukatch, Peggy and George Mamigonian, Tara McHugh, Sayan Mukherjee, Itamar Ronin, Adrienne Sacatos, Javid Sadr, Mike Saginaw, Atomu Sawatari, James Schummers, Kevin Thompson, and Louis Toth. Thank you one and all, from the bottom of my heart.

References

Asaad, W. F., Rainer, G., and Miller, E. K. Task-specific neural activity in the primate prefrontal cortex. *J. Neurophysiol.* 84: 451-459, 2000.

Banich, M. T., Milham, M. P., Atchley, R., Cohen, N. J., Webb, A., Wszalek, T., Kramer, A. F., Liang, Z.-P., Barad, V., Gullett, D., Shah, C., and Brown, C. Prefrontal regions play a predominant role in imposing an attentional 'set': evidence from fMRI. *Cogn. Brain Res.* 10: 1-9, 2000.

Barbas, H. and Pandya, D. N. Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *J. Comp. Neurol.* 286: 353-375, 1989.

Botvinick M. M., Braver, T. S., Barch, D. M., and Cohen, J. D. Conflict monitoring and cognitive control. *Psychol. Rev.* 108(3): 624-652, 2001.

Botvinick, M. M., Cohen, J. D., and Carter, C. S. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cog. Sci.* 8(12): 539-546, 2004.

Carter, C. S., Mintun, M., and Cohen, J. D. Interference and facilitation effects in an $H_2{}^{15}O$ PET study of Stroop task performance. *NeuroImage* 2: 264-272, 1995.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., and Cohen, J. D. Anterior cingulate cortex, error detection, and the on-line monitoring of performance. *Science* 280: 747-749, 1998.

Cohen, J. D., Dunbar, K., and McClelland, J. L. On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychol. Rev.* 97(3): 332-361, 1990.

Conover, W. J. *Practical Nonparametric Statistics.* Wiley Series in Probability and Statistics: Applied Probability and Statistics Section, 3rd ed. U.S.: John Wiley & Sons, Inc., chp. 4, pp.179-195, 1999.

De Valois, R. L. and Jacobs, G. H. Primate color vision. *Science* 162: 533 – 540, 1968.

De Valois, R. L., Morgan, H. C., Polson, M. C., Mead, W. R., and Hull, E. M. Psychophysical studies of monkey vision I. Macaque luminosity and color vision tests. *Vision Res.* 14: 53-67, 1974.

Duncan, J. and Owen, A. M. Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends Neurosci.* 23: 475-483, 2000.

Durston, S., Thomas, K. M., Worden, M. S., Yang, Y., and Casey, B. J. The effect of preceding context on inhibition: an event-related fMRI study. *NeuroImage* 16: 449-453, 2002.

Eacott, M. J. and Gaffan, D. Inferotemporal-frontal disconnection: the uncinate fascicle and visual associative learning in monkeys. *Eur. J. Neurosci.* 4: 1320-1333, 1992.

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

Freedman, D. J., Riesenhuber, M., Poggio, T., and Miller, E. K. Visual categorization and the primate prefrontal cortex: neurophysiology and behavior. *J. Neurophys.* 88: 929-941, 2002.

Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291: 1560-1563, 2001.

Funahashi, S., Bruce, C. J., Goldman-Rakic, P. S. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* 61(2): 331-349, 1989.

Fuster, J. M., Bauer, R. H., and Jervey, J. P. Cellular discharge in the dorsolateral prefrontal cortex of the monkey in cognitive tasks. *Exp. Neurol.* 77: 679-694, 1982.

Grant, D. A. & Berg, E. A. A behavioral analysis of degree of reinforcement and case of shifting to new responses in a Weigl-type card-sorting problem. *Exp. Psychol.* 38: 404-411, 1948.

Gross, C. G., Rocha-Miranda, C. E., and Bender, D. B. Visual properties of neurons in inferotemporal cortex of the macaque. *J. Neurophysiol.* 35(1): 96-111, 1972.

Kirk, RE. *Experimental Design: Procedures for the Behavioral Sciences.* Pacific Grove, CA: Brooks/Cole Publishing Company, 1995.

Klein, G. S. Semantic power measured through the interference of words with color-naming. *Am. J. Psychol.* 77: 576-588, 1964.

Konishi, S., Nakajima, K., Uchida, I., Kameyama, M., Nakahara, K., Sekihara, K., and Miyashita, Y. Transient activation of inferior prefrontal cortex during cognitive set shifting. *Nature Neurosci.* 1: 80-84, 1998.

Konishi, S., Kawazu, M., Uchida, I., Kikyo, H., Asakura, I., and Miyashita, Y. Contribution of working memory to transient activation in human inferior prefrontal cortex during performance of the Wisconsin Card Sorting Test. *Cereb. Cortex* 9: 745-753, 1999.

MacDonald, A. W., Cohen, J. D., Stenger, V. A., and Carter, C. S. Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science* 288:1835-1838, 2000.

MacLeod, C. M. and MacDonald, P. A. Interdimensional interference in the Stroop effect: uncovering the cognitive and neural anatomy of attention. *Trends Cog. Sci.* 4(10): 383-391, 2000.

Menon, V., Adleman, N. E., White, C. D., Glover, G. H., and Reiss, A. L. Error-related brain activation during a Go/NoGo response inhibition task. *Hum. Brain Mapp.* 12:131-143, 2001.

Milham, M. P., Banich, M. T., Webb, A., Barad, V., Cohen, N. J., Wszalek, T., and Kramer, A. F. The relative involvement of anterior cingulate and prefrontal cortex in attentional control depends on nature of conflict. *Cogn. Brain Res.* 12: 467-473, 2001.

Nakahara, K., Hayashi, T., Konishi, S., and Miyashita, Y. Functional MRI of macaque monkeys performing a cognitive set-shifting task. *Science* 295: 1532-1536, 2002.

Nichols, A. M., Ruffner, T. W., Sommer, M. A., and Wurtz, R. H. A screw microdrive for adjustable chronic unit recording in monkeys. *J. Neurosci. Methods* 81: 185-188, 1998.

Pardo, J. V., Pardo, P. J., Janer, K. W., and Raichle, M. W. The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. *PNAS* 87: 256-259, 1990.

Pandya, D. N. and Barnes, C. L. Architecture and connections of the frontal lobe. In: *The Frontal Lobes Revisited,* Perecman, E., ed. U.S.: IRBN Press, pp. 41-72, 1987.

Pandya, D. N., Van Hoesen, G. W., and Mesulam, M.-M. Efferent connections of the cingulate gyrus in the rhesus monkey. *Exp. Brain Res.* 42:319-330, 1981.

Passingham, R. E., Toni, I., and Rushworth, M. F. S. Specialisation within the prefrontal cortex: the ventral prefrontal cortex and associative learning. *Exp. Brain Res.* 133: 103-113, 2000.

Peterson, B. S., Kane, M. J., Alexander, G. M., Lacadie, C., Skudlarski, P. Leung, H.-C., May, J., and Gore, J. C. An event-related functional MRI study comparing interference in the Simon and Stroop tasks. *Cog. Brain Res.* 13: 427-440, 2002.

Petrides, M., Alivisatos, B., and Frey, S. Differential activation of the human orbital, mid-ventrolateral, and mid-dorsolateral prefrontal cortex during the processing of visual stimuli. *PNAS* 99(8): 5649-5654, 2002.

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

Rao, S. C., Rainer, G., and Miller, E. K. Integration of what and where in the primate prefrontal cortex. *Science* 276: 821-824, 1997.

Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., and Nieuwenhuis, S. The role of the medial frontal cortex in cognitive control. *Science* 306: 443-447, 2004.

Rosenkilde, C. E., Bauer, R. H., and Fuster, J. M. Single cell activity in ventral prefrontal cortex of behaving monkeys. *Brain Res.* 209: 375-394, 1981.

Rubia, K., Smith, A. B., Brammer, M. J., and Taylor, E. Right inferior prefrontal cortex mediates response inhibition while mesial prefrontal cortex is responsible for error detection. *NeuroImage* 20: 351-358, 2003.

Ruff, C. C., Woodward, T. S., Laurens, K. R., and Liddle, P. F. The role of the anterior cingulate cortex in conflict processing: evidence from reverse Stroop interference. *NeuroImage* 14: 1150-1158, 2001.

Rushworth, M. F. S., Nixon, P. D., Eacott, M. J., and Passingham, R. E. Ventral prefrontal cortex is not essential for working memory. *J. Neurosci.* 17(12): 4829-4838, 1997.

Rushworth, M. F. S., Paus, T., and Sipila, P. K. Attentional systems and the organization of the human parietal cortex. *J. Neurosci.* 21(14): 5262-5271, 2001.

Swainson, R., Cunnington, R., Jackson, G. M., Rorden, C., Peters, A. M., Morris, P. G., and Jackson, S. R. Cognitive control mechanisms revealed by ERP and fMRI: evidence from repeated task-switching. *J. Cogn. Neurosci.* 15(6): 785-799, 2003.

Tallon-Baudry, C. Attention and awareness in synchrony. *Trends Cogn. Sci.* 8(12): 523-525, 2004.

Ungerleider, L. G. and Mishkin, M. Two cortical visual systems. In: *Analysis of Visual Behavior,* Ingle, D. J., Goodale, M. A., and Mansfield, R. J. W., eds. Cambridge, MA: MIT Press, pp. 549-586, 1982.

Van Veen, V., Holroyd, C. B., Cohen, J. D., Stenger, V. A., and Carter, C. S. Errors without conflict: implications for performance monitoring theories of anterior cingulate cortex. *Brain Cog.* 56:267-276, 2004.

Vogt, B. A. and Pandya, D. N. Cingulate cortex of the rhesus monkey: II. Cortical afferents. *J. Comp. Neurol.* 262: 271-289, 1987.

Wallis, J. D., Anderson, K. C., and Miller, E. K. Single neurons in prefrontal cortex encode abstract rules. *Nature* 411: 953-956, 2001.

Wallis, J. D. and Miller, E. K. From rule to response: neuronal processes in the premotor and prefrontal cortex. *J. Neurophysiol.* 90: 1790-1806, 2003.

Watanabe, M. Prefrontal unit activity during delayed conditional discriminations in the monkey. *Brain Res.* 225: 51-65, 1981.

Webster, M. J., Bachevalier, J., and Ungerleider, L. G. Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. *Cereb. Cortex* 5: 470-483, 1994.

White, I. M., and Wise, S. P. Rule-dependent neuronal activity in the prefrontal cortex. *Exp. Brain Res.* 126: 315-335, 1999.

Table 1A: Accuracy and Reaction Times for Match Trials

	Number Of Trials	Accuracy	Mean Reaction Time (msec)
Instruction Trials			
Orientation Relevant	1333	99.6%	328.9 ± 0.07
Color Relevant	1254	97.1%	361.2 ± 0.09
Either Attribute	*2587*	*98.4%*	*344.3 ± 0.04*
DMSA Trials			
Congruent			
Orientation Relevant	1778	99.6%	320.4 ± 0.04
Color Relevant	1740	98.6%	327.6 ± 0.06
Either Attribute	*3518*	*99.1%*	*323.9 ± 0.02*
Incongruent			
Orientation Relevant	1739	83.8%	368.6 ± 0.09
Color Relevant	1815	84.2%	398.2 ± 0.10
Either Attribute	*3554*	*84.0%*	*383.7 ± 0.05*

Table 1B: Accuracy and Reaction Times for Non-Match Trials

	Number Of Trials	Accuracy	Mean Reaction Time (msec)
Instruction Trials			
Orientation Relevant	1185	98.1%	200.8 ± 0.11
Color Relevant	1307	92.2%	195.0 ± 0.10
Either Attribute	*2492*	*95.0%*	*197.9 ± 0.05*
DMSA Trials			
Congruent			
Orientation Relevant	1599	97.4%	188.0 ± 0.08
Color Relevant	1633	97.7%	184.8 ± 0.08
Either Attribute	*3232*	*97.6%*	*186.4 ± 0.04*
Incongruent			
Orientation Relevant	1703	86.4%	201.5 ± 0.08
Color Relevant	1742	83.8%	186.5 ± 0.07
Either Attribute	*3445*	*85.1%*	*194.0 ± 0.04*

A

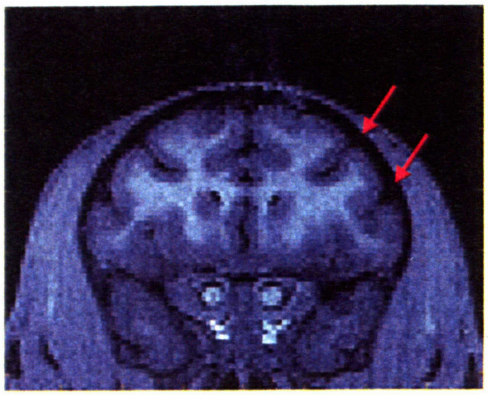

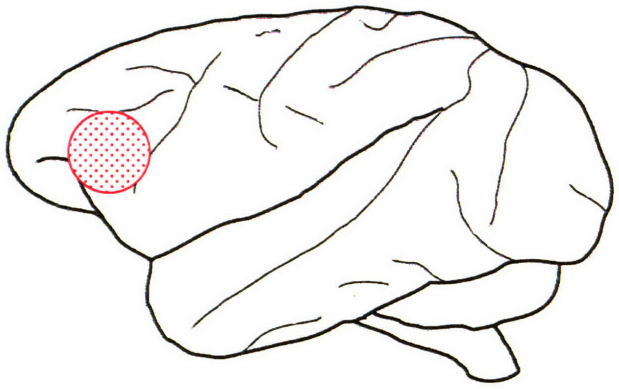

B

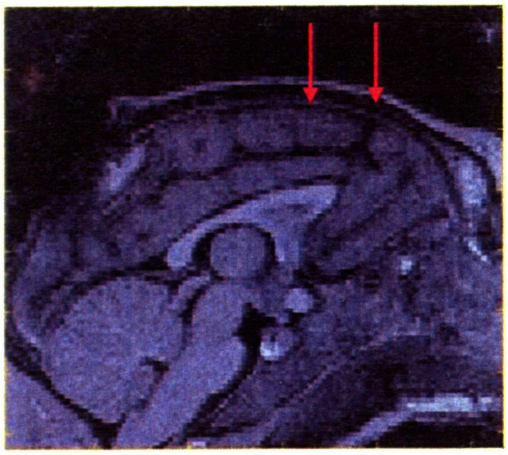

Fig. 1. MRI images showing the placement of the recording chambers: **(A),** coronal section showing where the chamber was centered over the VLPFC, as indicated by the two red arrows. The diagram on the right shows how the well was placed over the inferior prefrontal convexity, below the principal sulcus. **(B),** mid-sagittal section showing the sulcal and callosal landmarks used to center the ACC chamber. The red arrows indicate the anterior and posterior extent of the well.

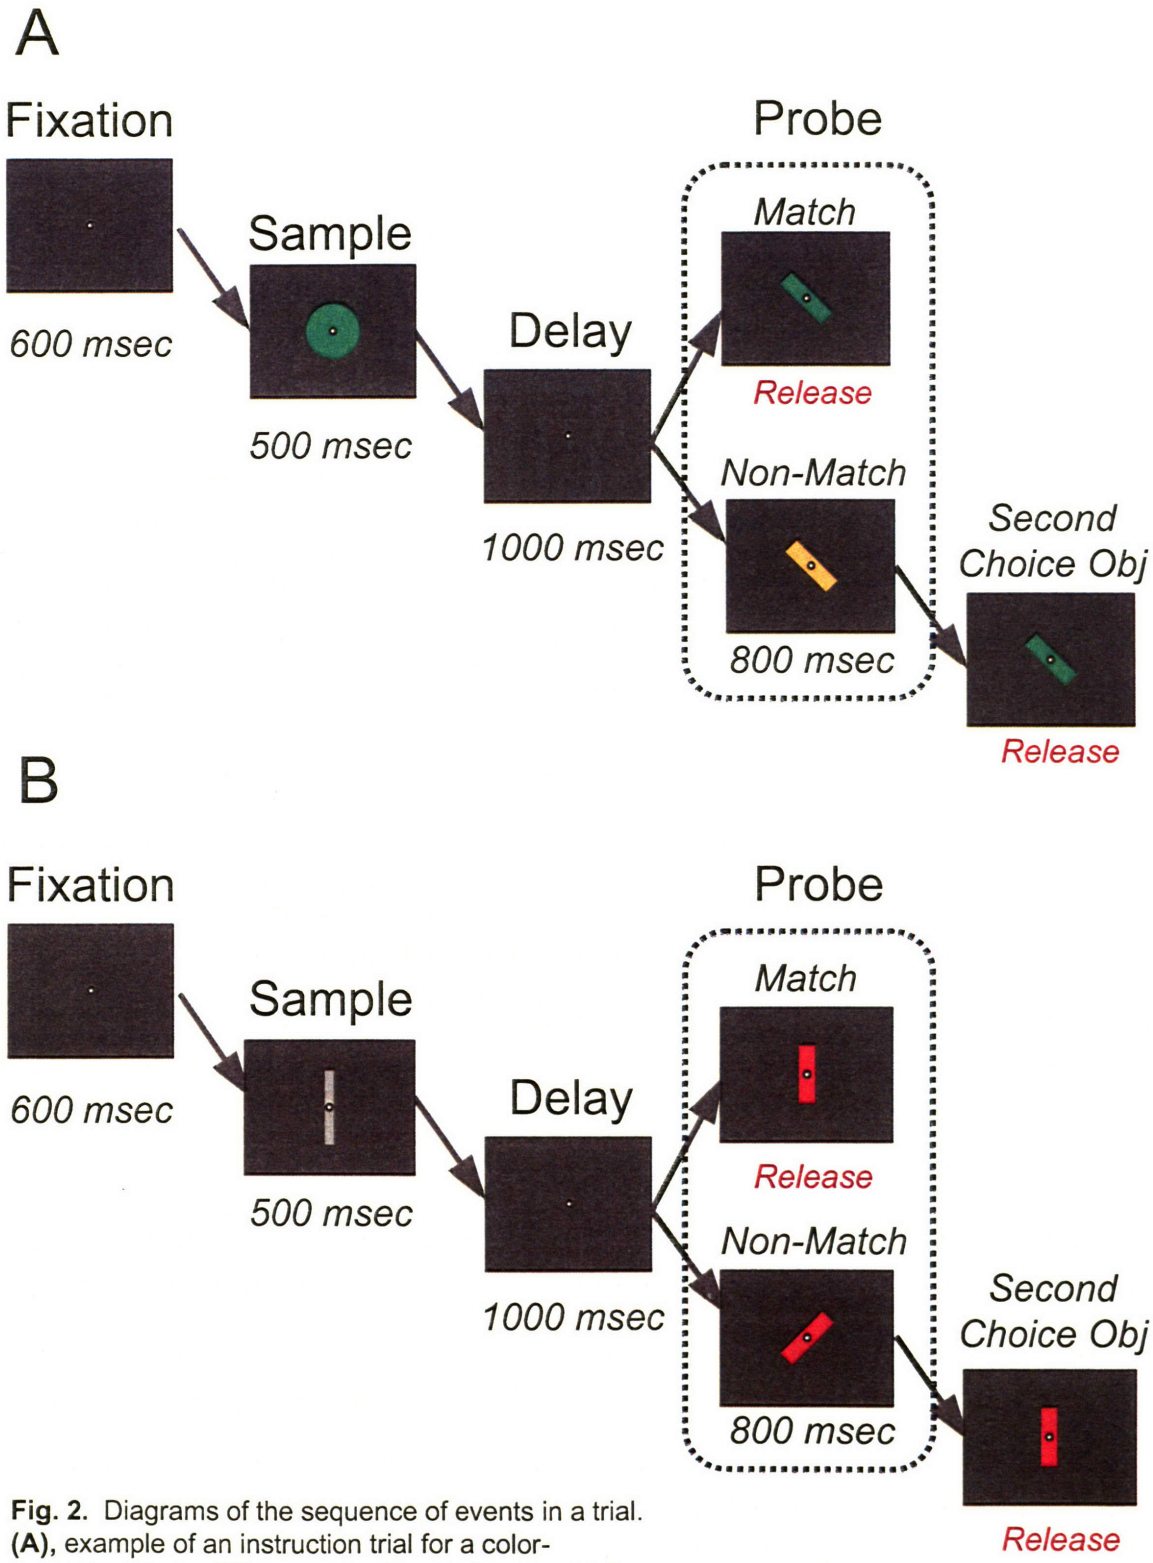

Fig. 2. Diagrams of the sequence of events in a trial. **(A),** example of an instruction trial for a color-matching block. **(B),** example of an instruction trial for an orientation-matching block.

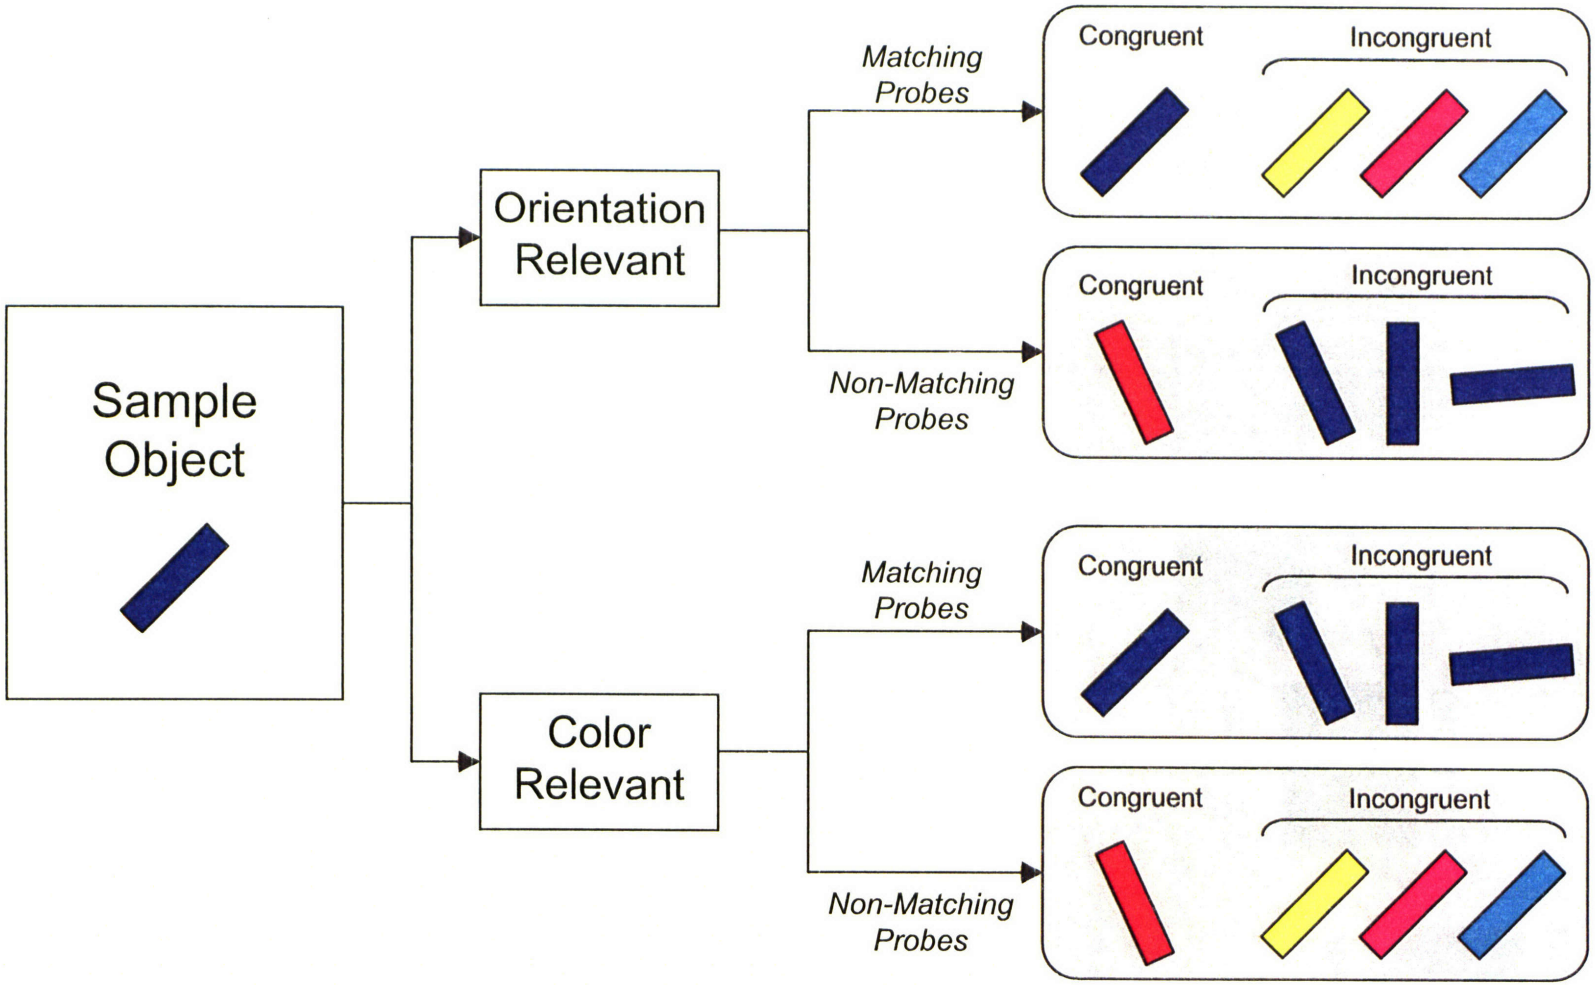

Fig. 3. Illustration of stimuli used in the DMSA task. When orientation is relevant, the monkey must match the probe's orientation to the sample, ignoring color (**top**). When color is relevant, the monkey must match the probe's color to the sample, ignoring orientation (**bottom**). Incongruent probes match the sample along only one attribute dimension, and thus can either be matches or non-matches, depending on the current matching rule. Congruent probes are either always a match or never a match, regardless of the matching rule.

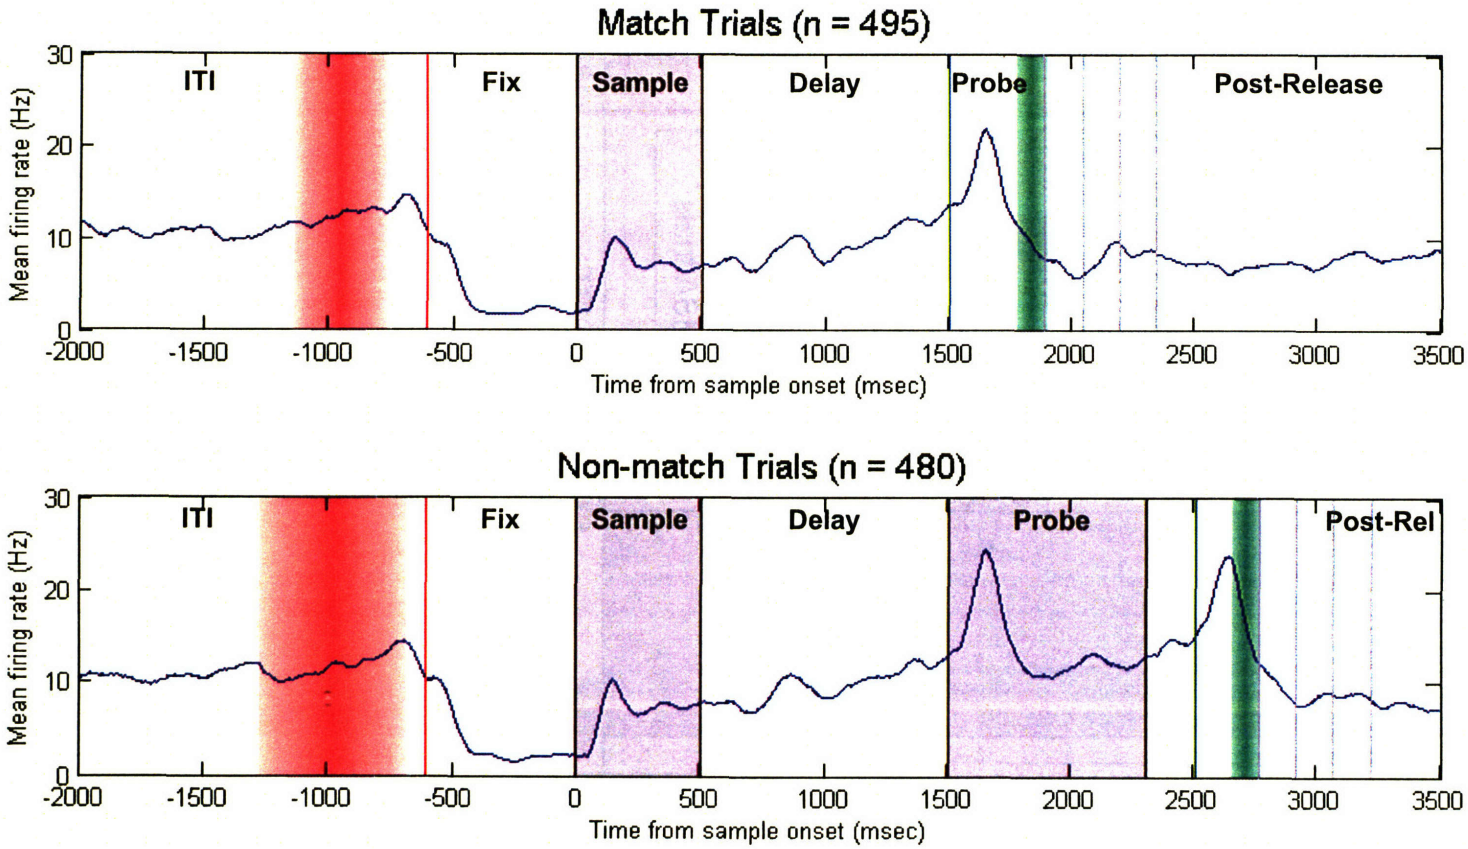

Fig. 4. PSTHs illustrating trial epochs and the task-related responses of a single VLPFC neuron. Match (**top** histogram) and non-match (**bottom** histogram) trials are plotted separately. Trials are aligned on sample onset (t = 0). The areas shaded in gray indicate when the sample object was displayed, and also when the probe object was on-screen during non-match trials. Fixation events are shown in red: the vertical red line indicates the time of fixation acquisition, while the red patch to its left indicates the average time of the fixation spot onset, ± 1 SD. Bar release events are shown in green. The vertical green line indicates the when the matching choice object appeared. For match trials, this event was the onset of the probe object, and for non-match trials, it was the onset of the second choice object. The green patch to the right of the green line shows the monkey's average reaction time to release the bar, ± 1 SD. After bar release, the delivery of four drops of liquid reward is illustrated by four vertical gray lines.

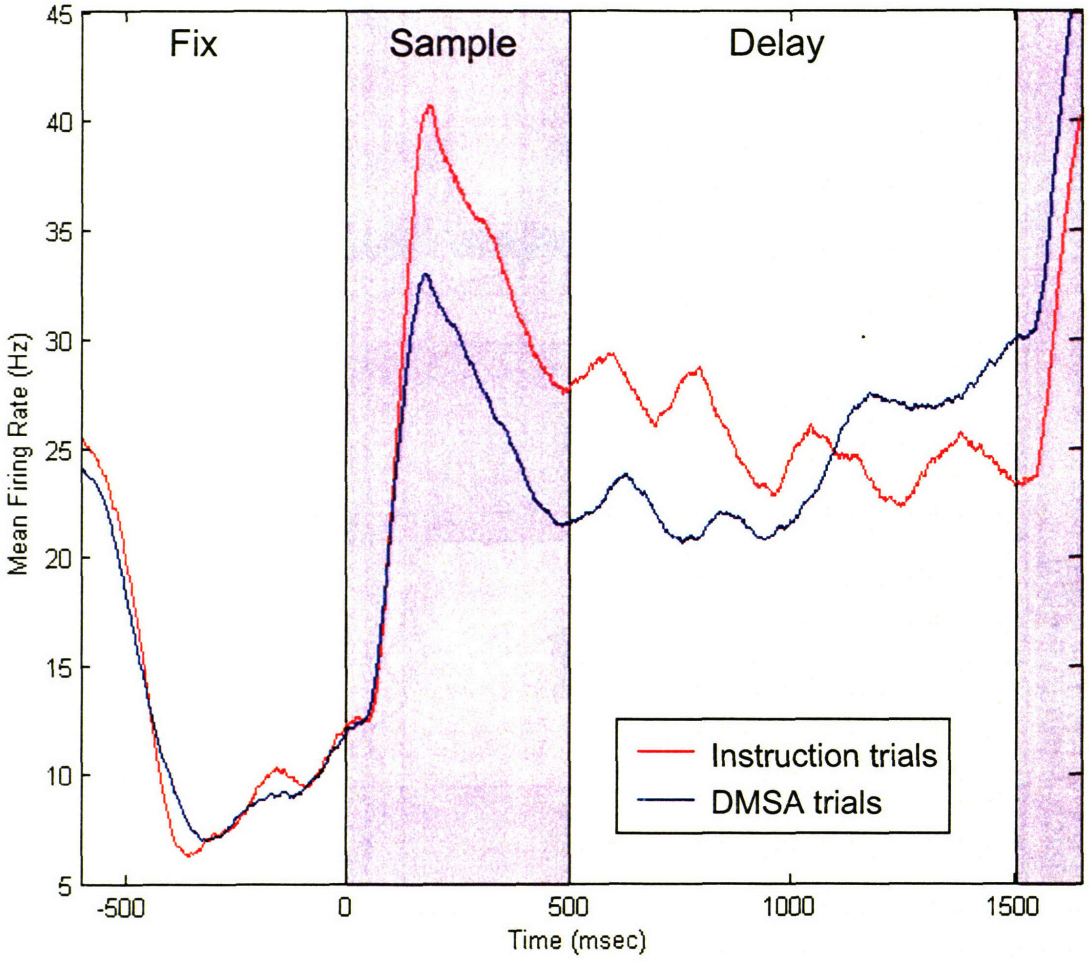

Fig. 5. A VLPFC cell with stronger responses to sample stimuli during instruction trials (n = 145) than during DMSA trials (n = 356): Wilcoxon rank sum test, $p < 2.6 \times 10^{-8}$.

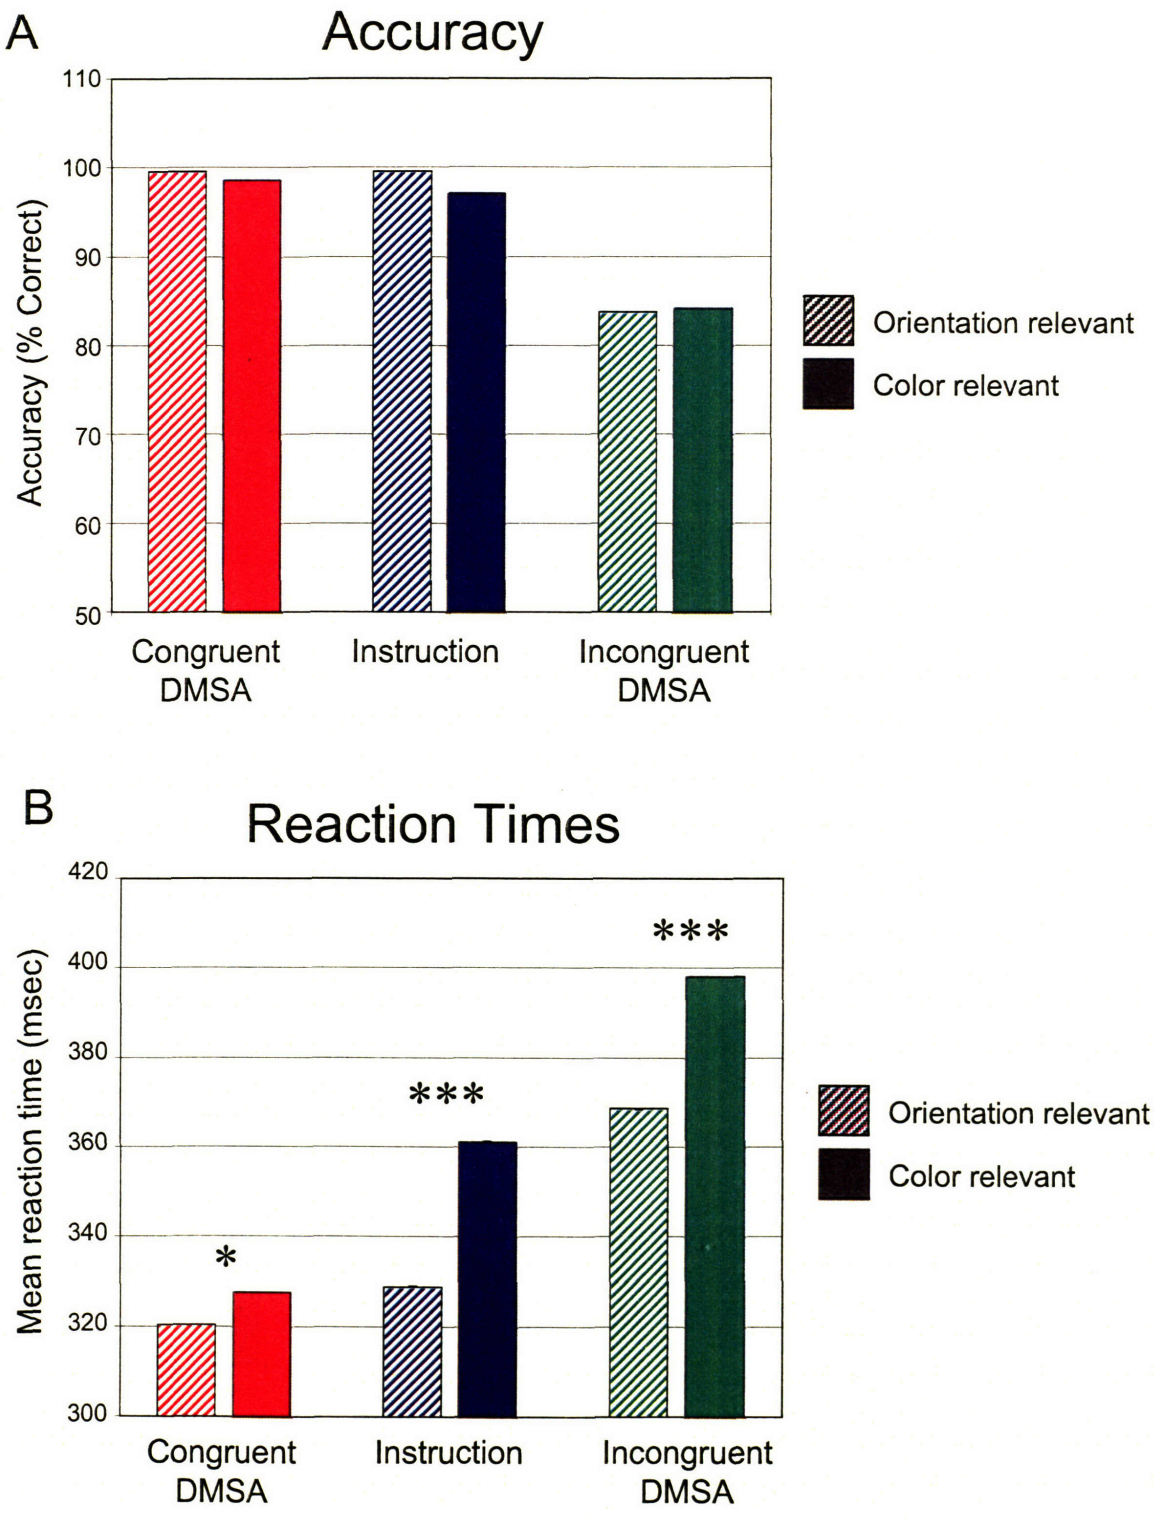

Fig. 6. Behavioral data for match trials: **(A),** performance accuracy; **(B),** mean reaction times. Data for congruent DMSA trials, instruction trials, and incongruent DMSA trials are grouped separately and broken down by matching rule. Striped bars: orientation is relevant. Solid bars: color is relevant. Significant differences based on matching rule are indicated by asterisks: * $p < 10^{-5}$; *** $p < 10^{-25}$.

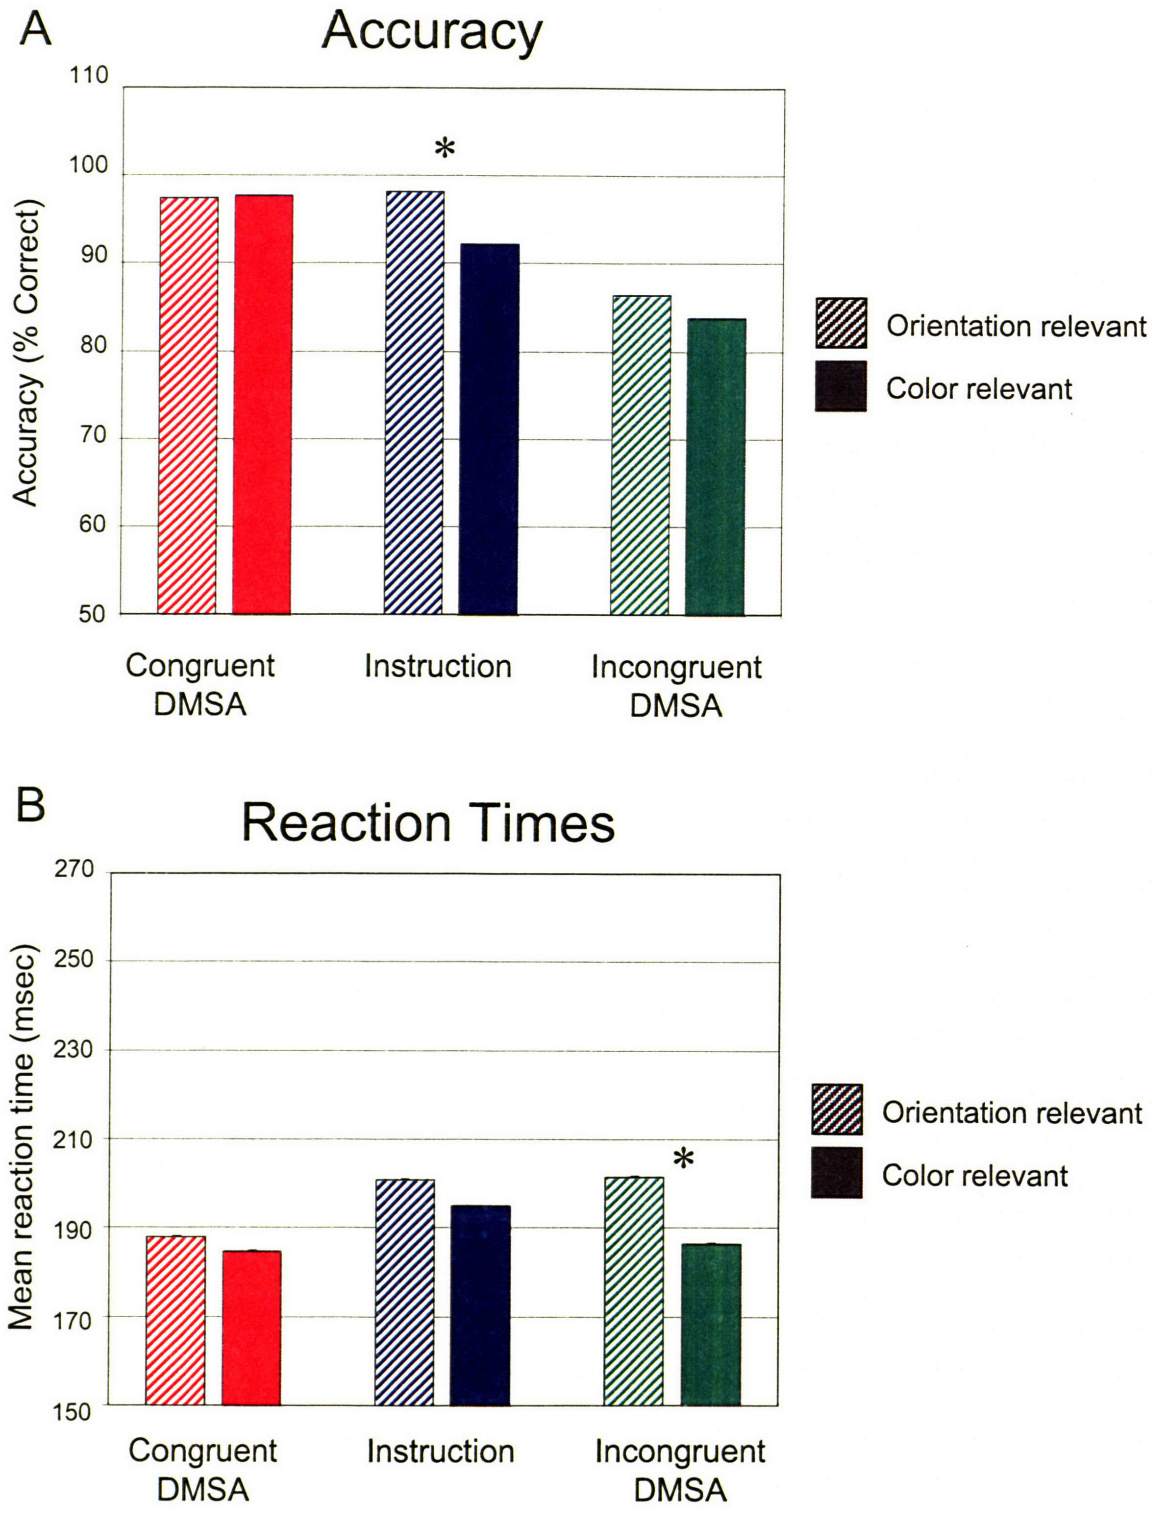

Fig. 7. Behavioral data for non-match trials: **(A)**, performance accuracy; **(B)**, mean reaction times. Data for congruent DMSA trials, instruction trials, and incongruent DMSA trials are grouped separately and broken down by matching rule. Striped bars: orientation is relevant. Solid bars: color is relevant. Significant differences based on matching rule are indicated by asterisks: * $p < 10^{-5}$; *** $p < 10^{-25}$.

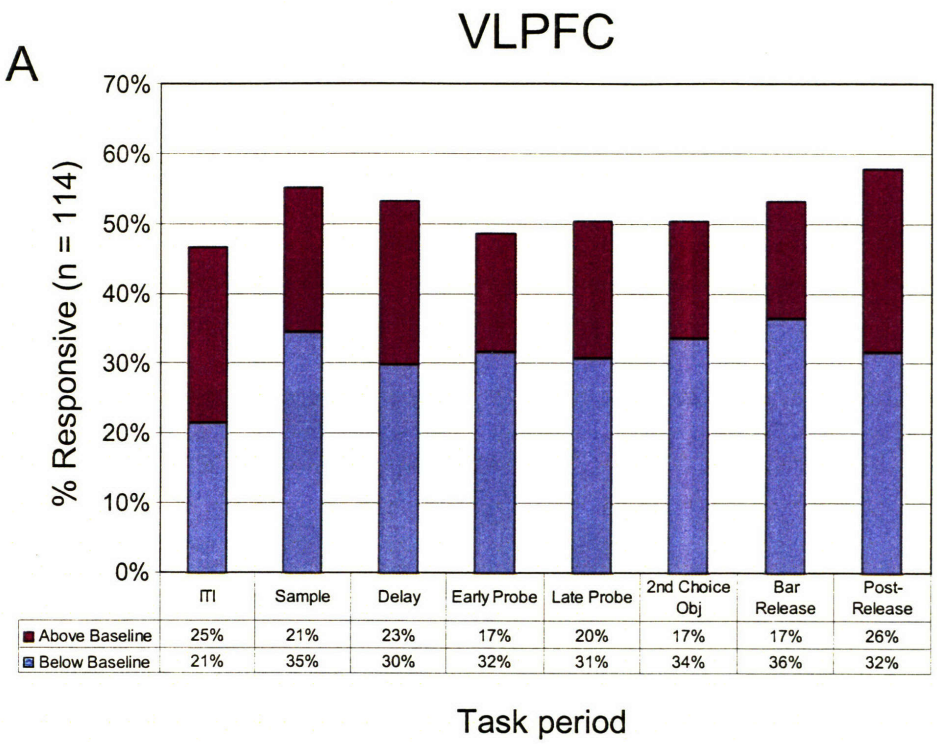

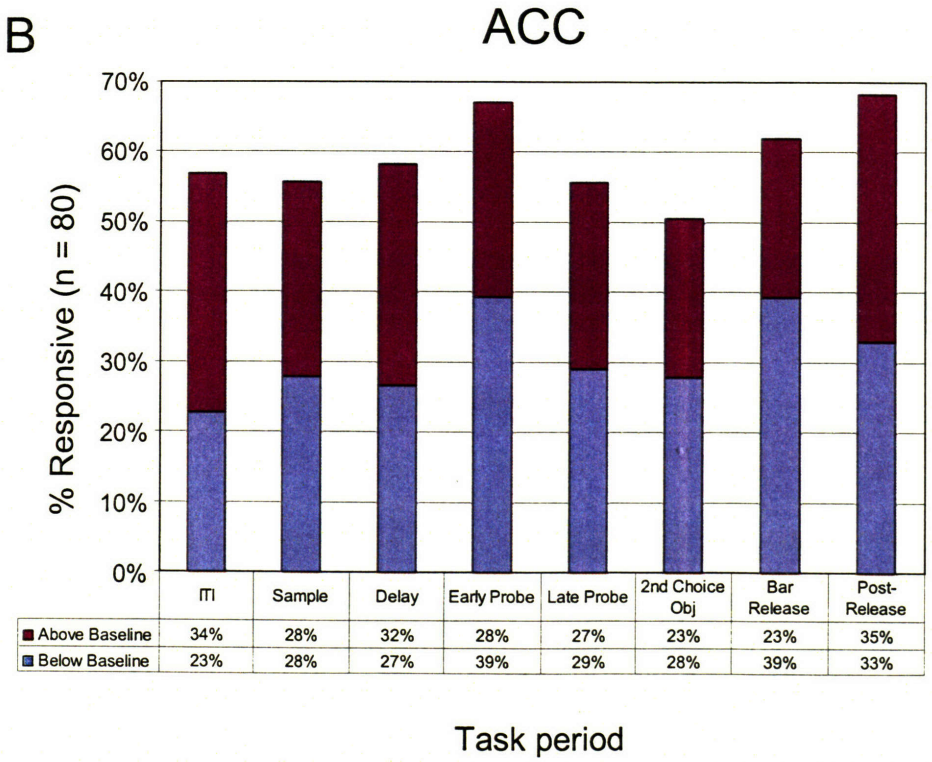

Fig. 8. Percentage of task-responsive neurons in **(A)** VLPFC and **(B)** ACC during each task period. Each bar is sub-divided to show the ratio of cells with mean firing rates above (magenta) and below (gray) baseline, with exact percentages presented in the tables below each chart.

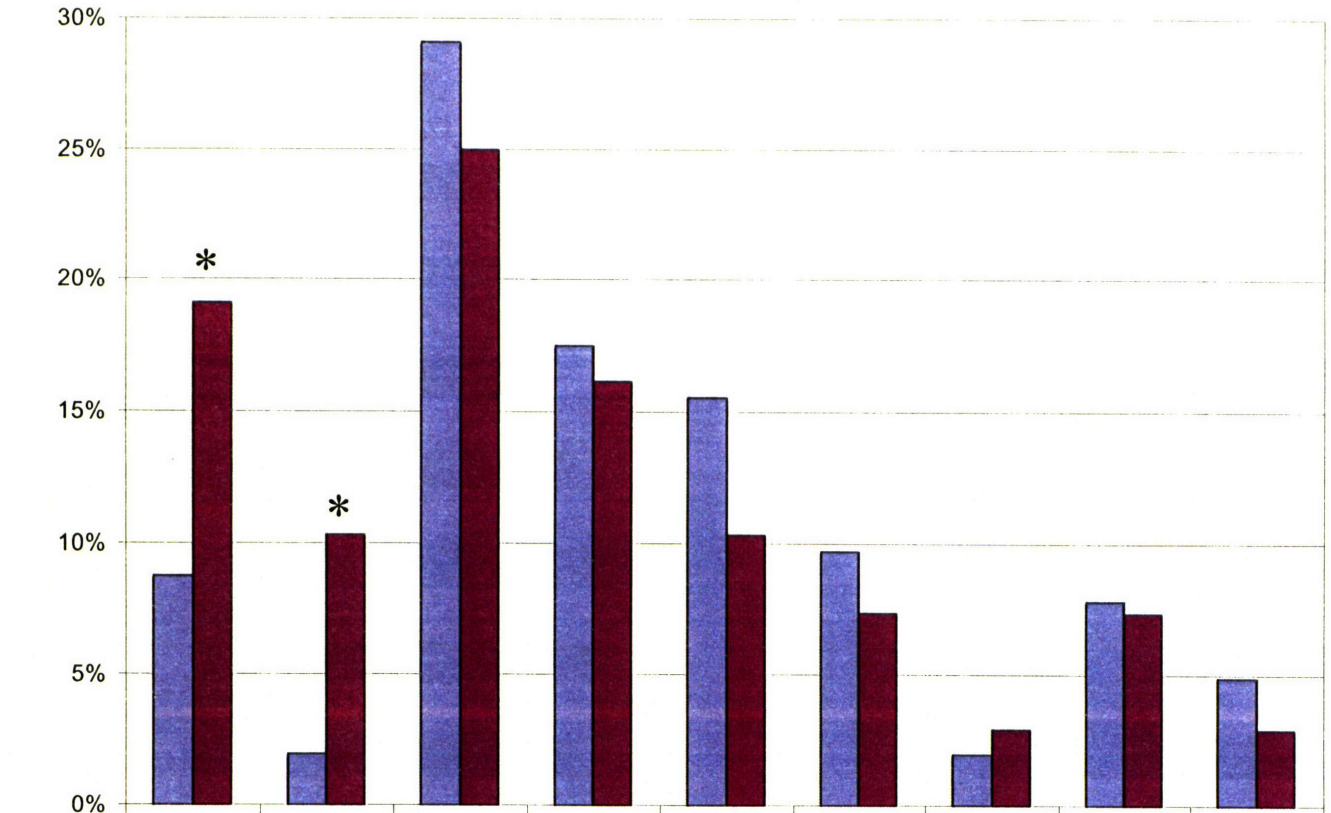

	ITI	Fix	Sample	Delay	Early Probe	Late Probe	2nd Choice Obj	Bar Release	Post-Release
VLPFC	8.7%	1.9%	29.1%	17.5%	15.5%	9.7%	1.9%	7.8%	4.9%
ACC	19.1%	10.3%	25.0%	16.2%	10.3%	7.4%	2.9%	7.4%	2.9%

Fig. 9. Percentage of VLPFC cells (n = 103; shown in gray) and ACC cells (n = 68; shown in magenta) in each task period whose firing rates for instruction and DMSA trials were significantly different. In both areas, these effects were most common early in the trial, before the monkey made a behavioral response. Significant differences between VLPFC and ACC cells (χ^2 test, $p < 0.05$) are denoted by asterisks.

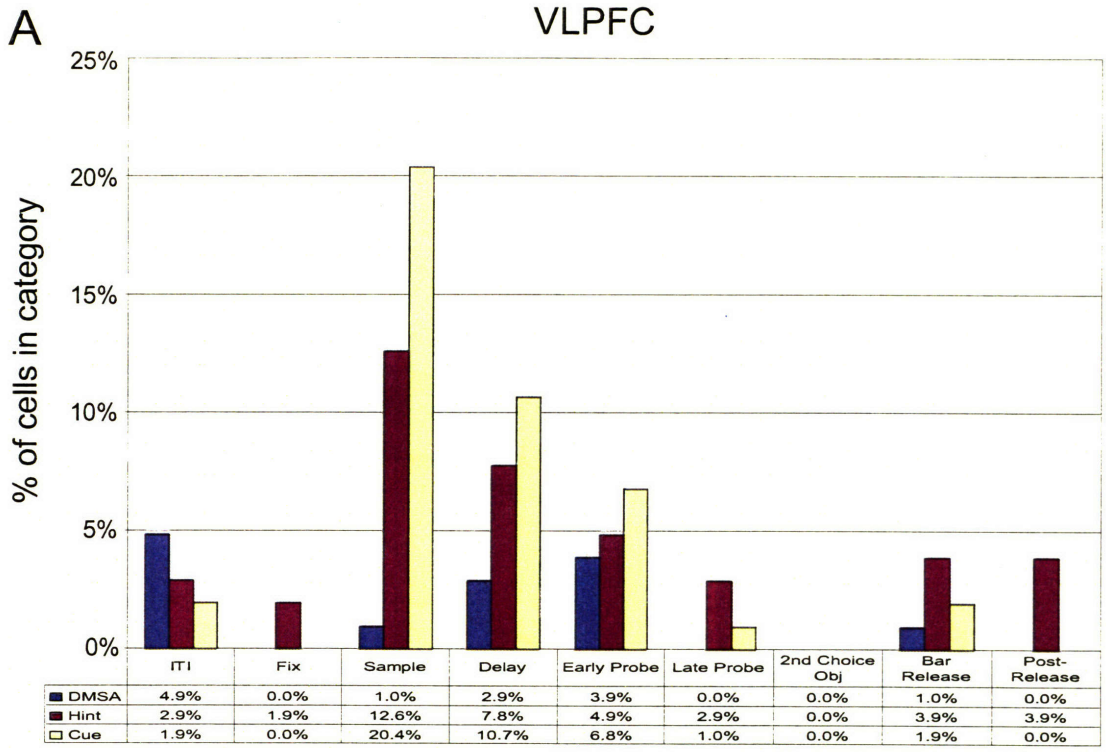

	ITI	Fix	Sample	Delay	Early Probe	Late Probe	2nd Choice Obj	Bar Release	Post-Release
DMSA	4.9%	0.0%	1.0%	2.9%	3.9%	0.0%	0.0%	1.0%	0.0%
Hint	2.9%	1.9%	12.6%	7.8%	4.9%	2.9%	0.0%	3.9%	3.9%
Cue	1.9%	0.0%	20.4%	10.7%	6.8%	1.0%	0.0%	1.9%	0.0%

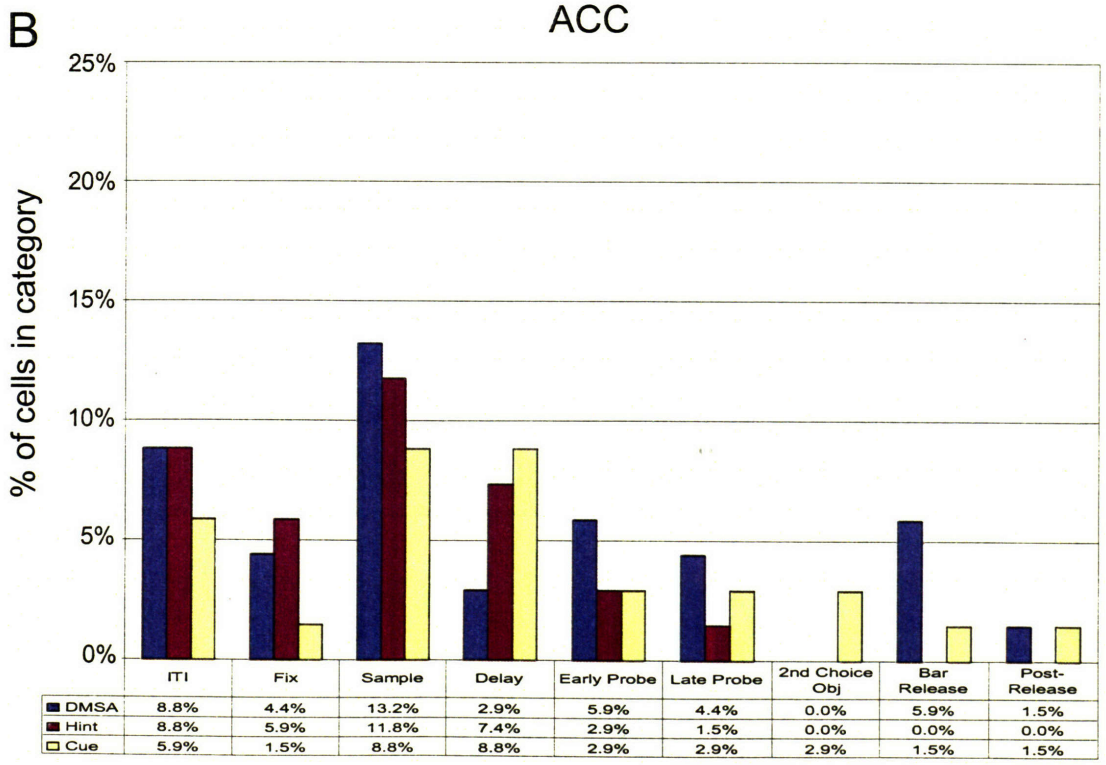

	ITI	Fix	Sample	Delay	Early Probe	Late Probe	2nd Choice Obj	Bar Release	Post-Release
DMSA	8.8%	4.4%	13.2%	2.9%	5.9%	4.4%	0.0%	5.9%	1.5%
Hint	8.8%	5.9%	11.8%	7.4%	2.9%	1.5%	0.0%	0.0%	0.0%
Cue	5.9%	1.5%	8.8%	8.8%	2.9%	2.9%	2.9%	1.5%	1.5%

Fig. 10. Task-related responses that were modulated by trial type: Percentage of **(A)** VLPFC cells (n = 103) and **(B)** ACC cells (n = 68) that had significantly stronger responses for DMSA trials (purple), hint trials (magenta) and cue trials (beige). Data are grouped by task period. Exact percentages are shown in the tables below each chart.

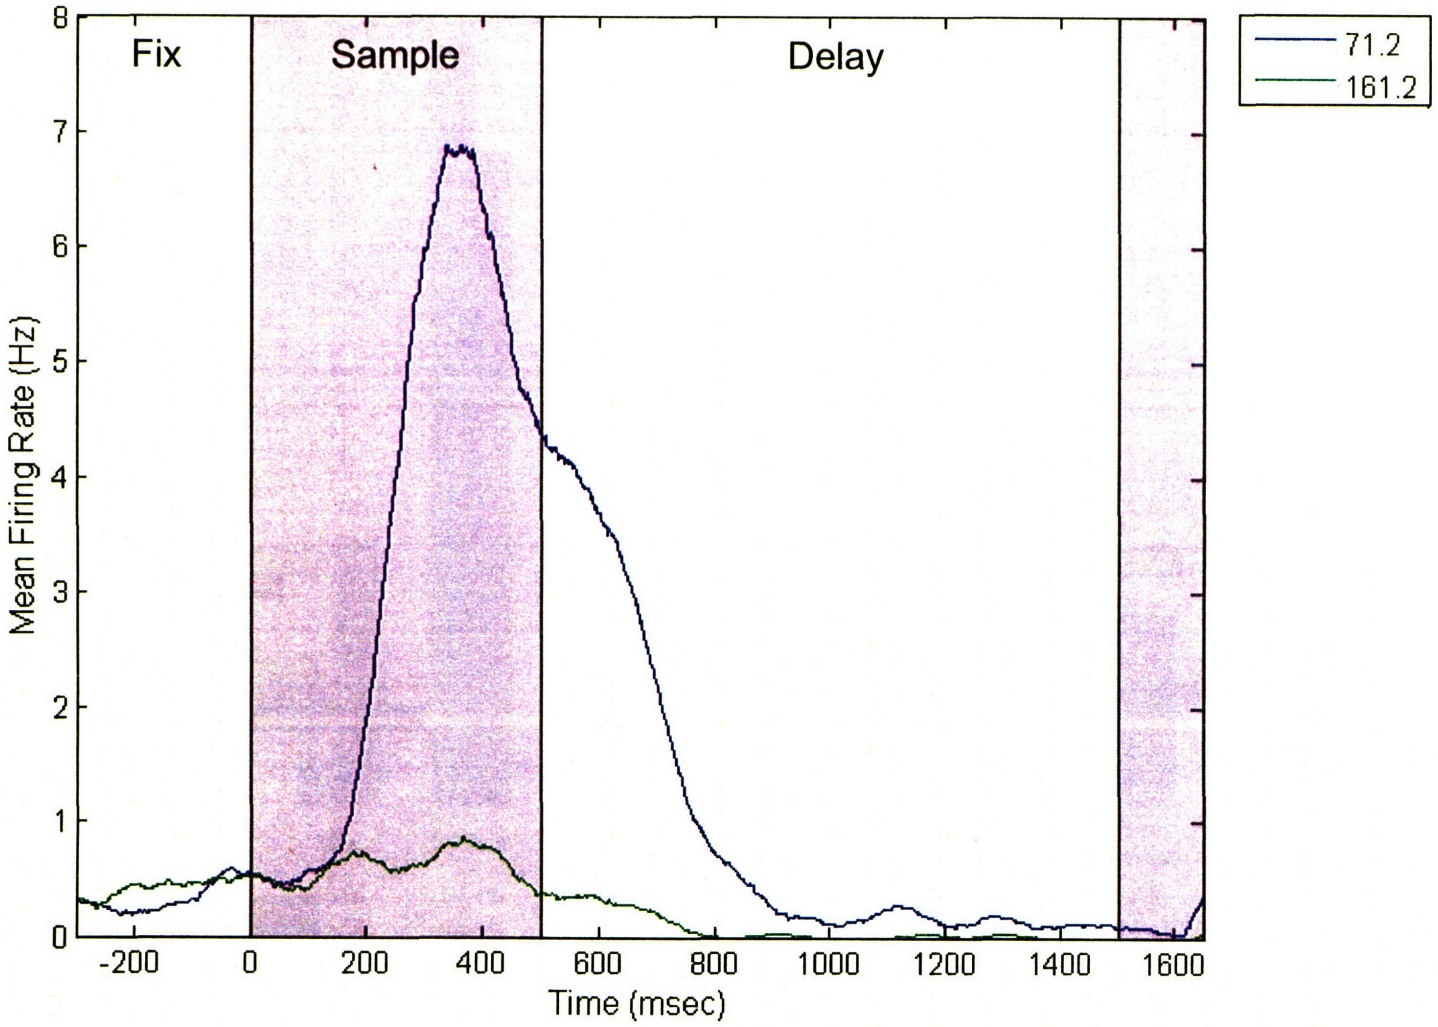

Fig. 11. Example of an orientation-selective VLPFC cell. This cell preferred 71.2° sample objects to 161.2° sample objects. Neuronal activity during all correct DMSA trials is shown, and is aligned on the time of sample onset.

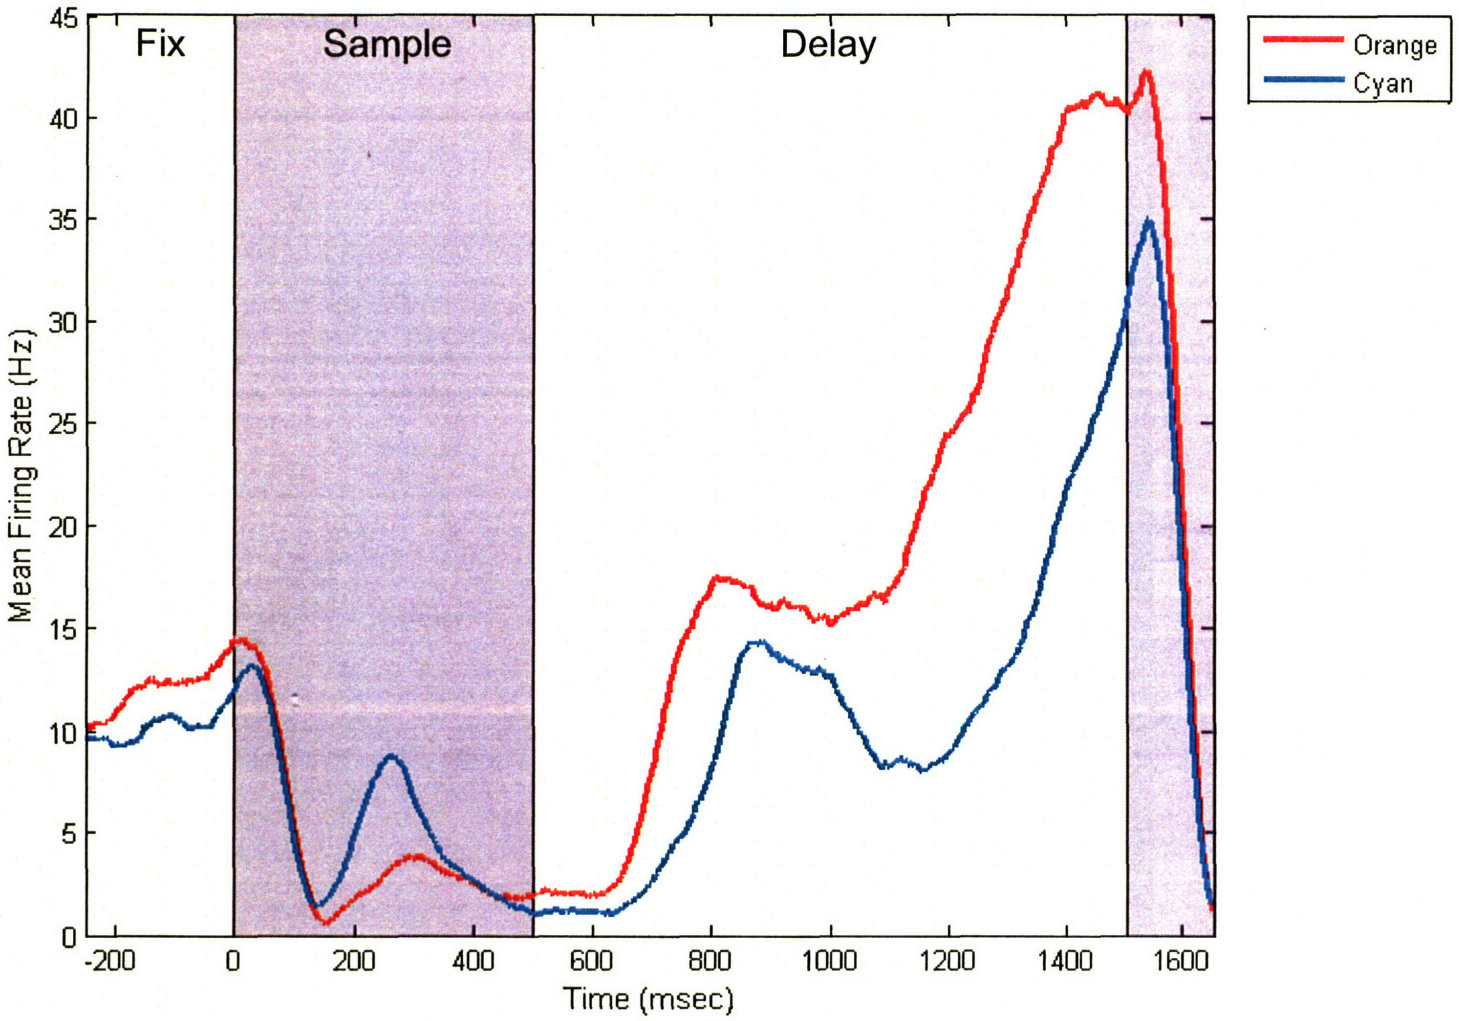

Fig. 12. Example of a color-selective VLPFC cell. This cell showed stronger delay-period activity when the sample color was orange. Neuronal activity during correct DMSA trials is shown, and is aligned on the time of sample onset.

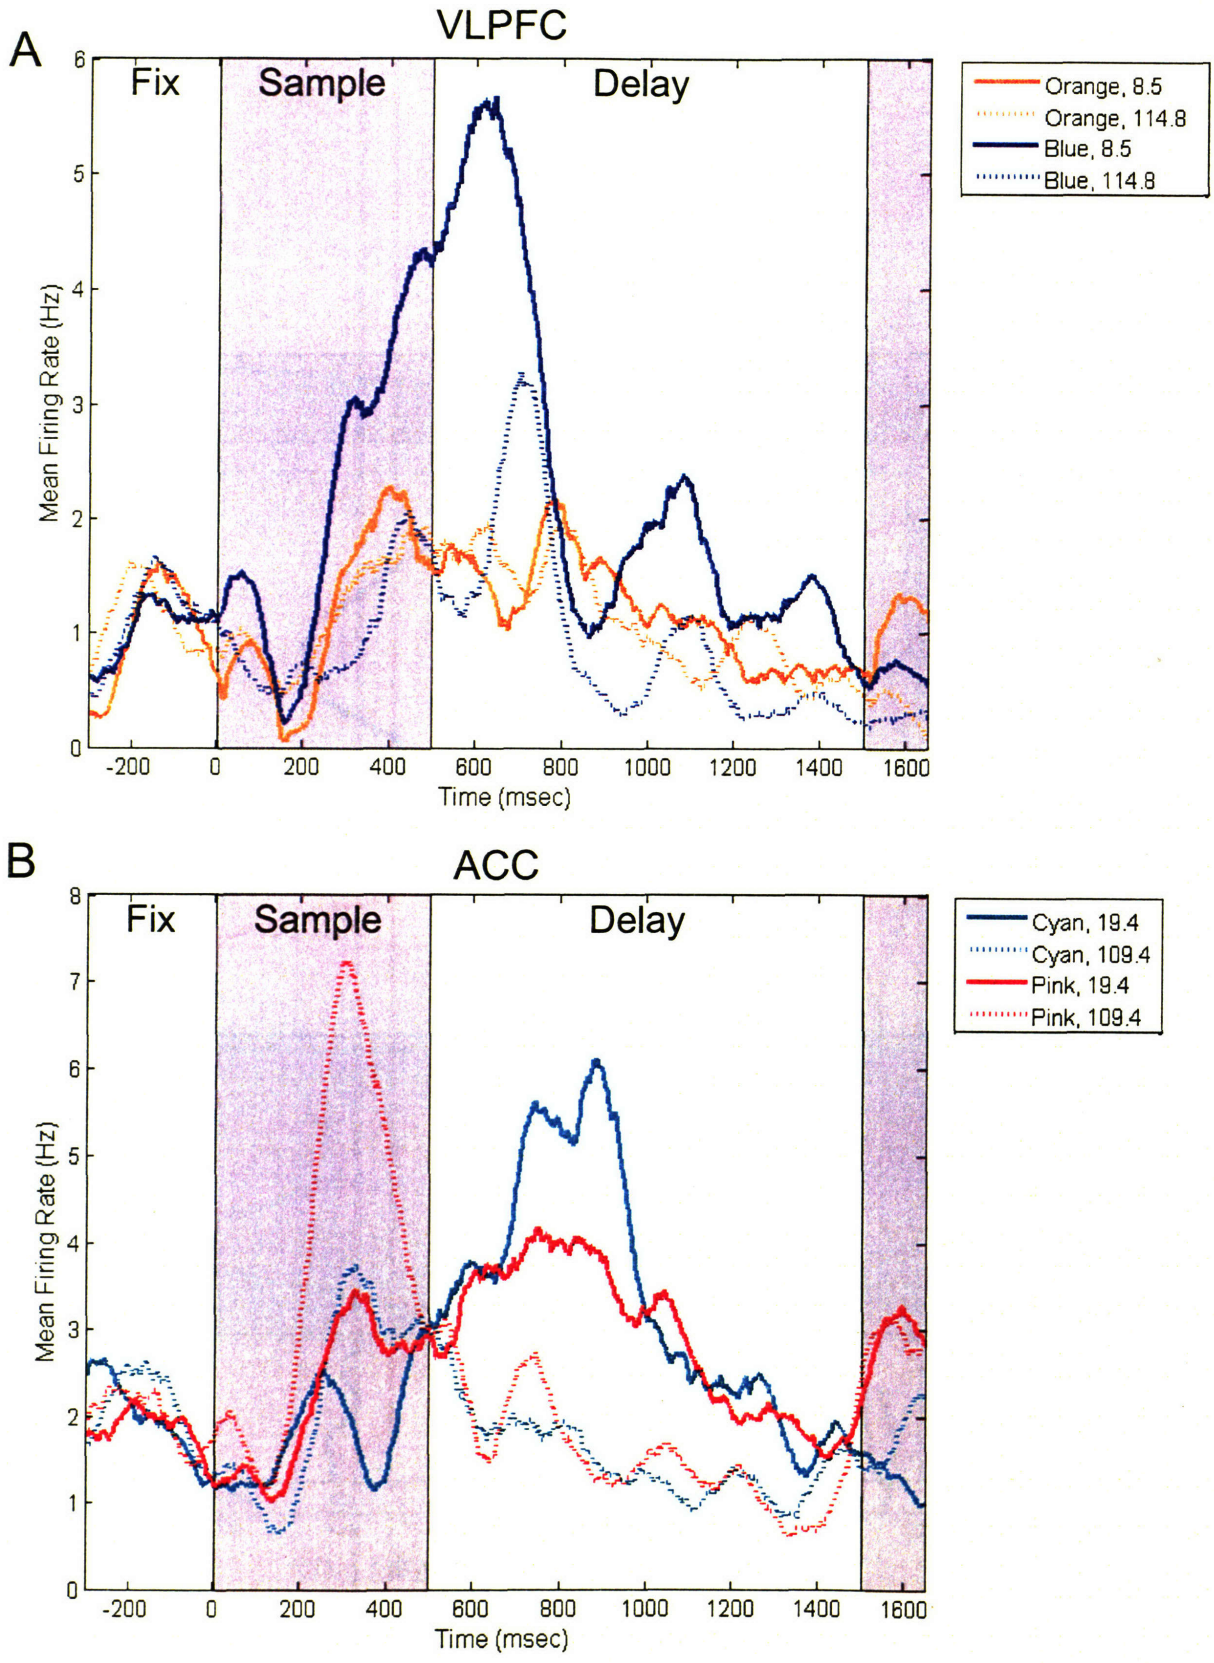

Fig. 13. Two cells that were selective for specific items during the sample period: **(A)**, VLPFC cell that preferred the blue sample object with an orientation of 8.5°; **(B)**, ACC cell that preferred the pink sample object with an orientation of 19.4°.

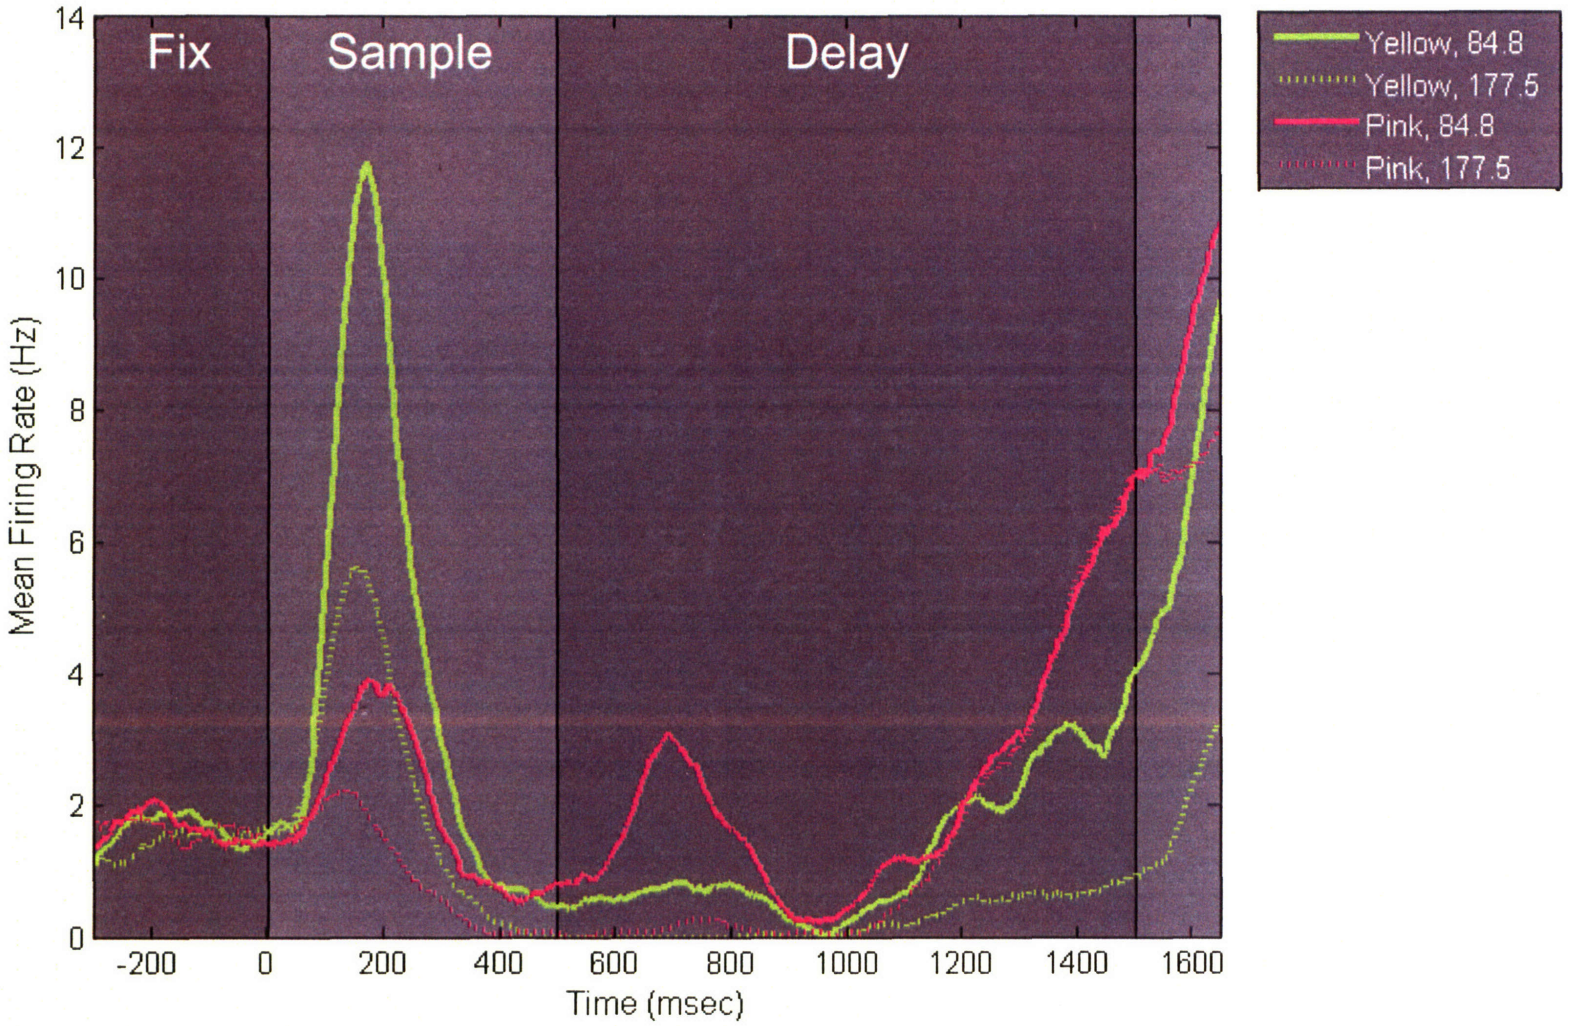

Fig. 14. Example of a VLPFC cell that was selective for both attributes during the sample period, but with no interaction between orientation-selectivity and color-selectivity. Data are shown using actual sample colors; histogram is displayed on a dark background in order to enhance contrast.

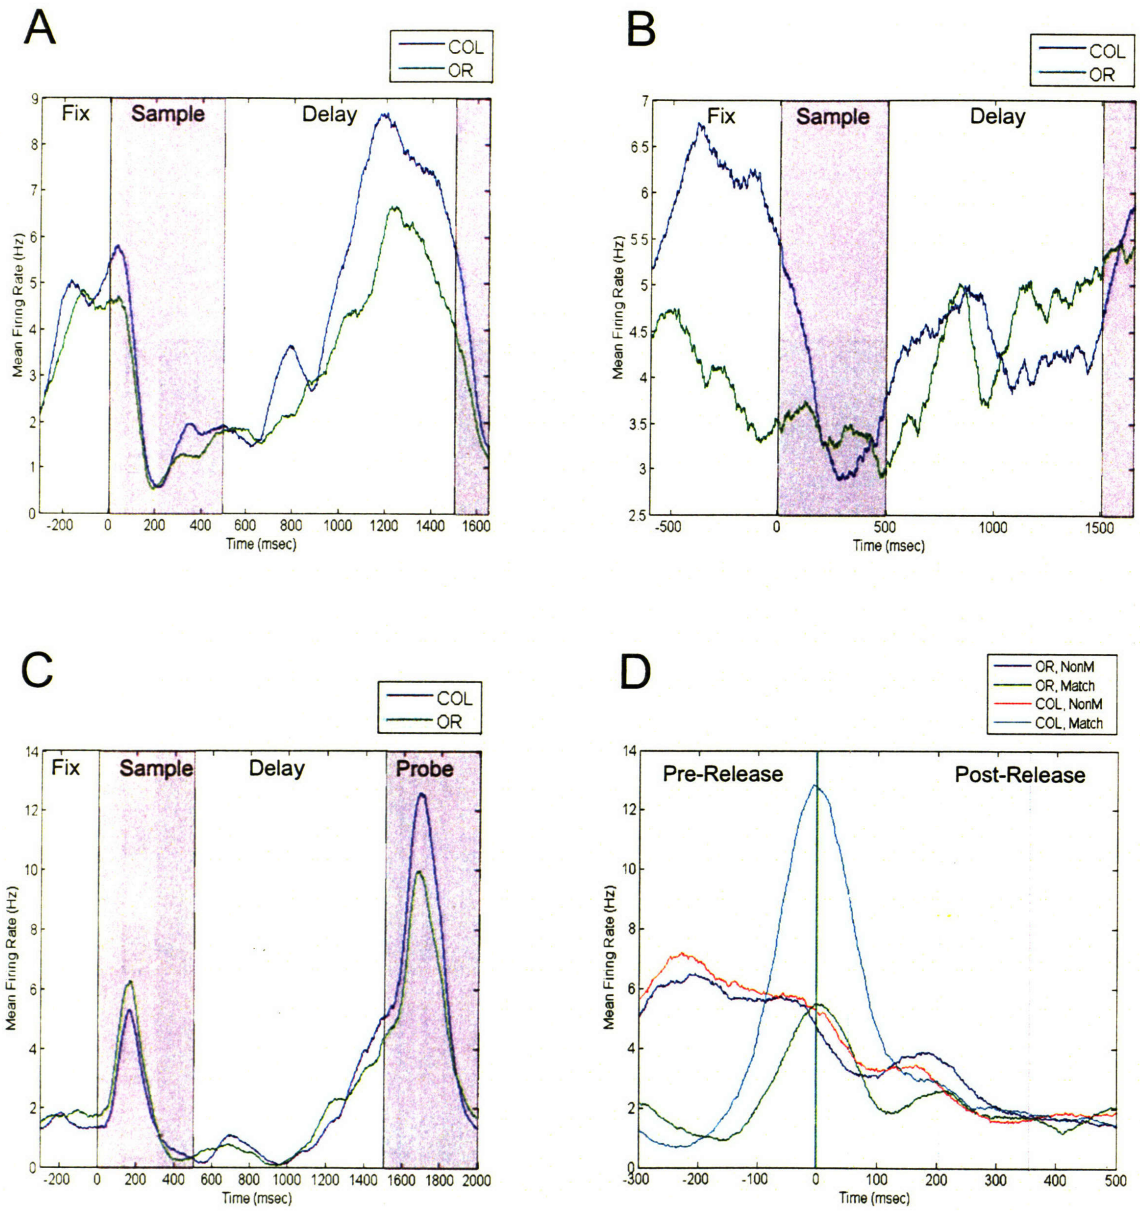

Fig. 15. Rule-selective cells: **(A),** VLPFC cell with a main effect of rule during the delay period; **(B),** ACC cell with a main effect of rule during the fixation period; **(C),** VLPFC cell with a main effect of rule during the early probe period; **(D),** ACC cell with a Rule x Decision interaction during the bar-release period; this cell was rule-selective during match trials only. All histograms aligned on sample onset except (D), which is aligned on the time of bar release, indicated by the vertical green line. The vertical gray lines in (D) indicate when the monkey received a drop of juice as a reward.

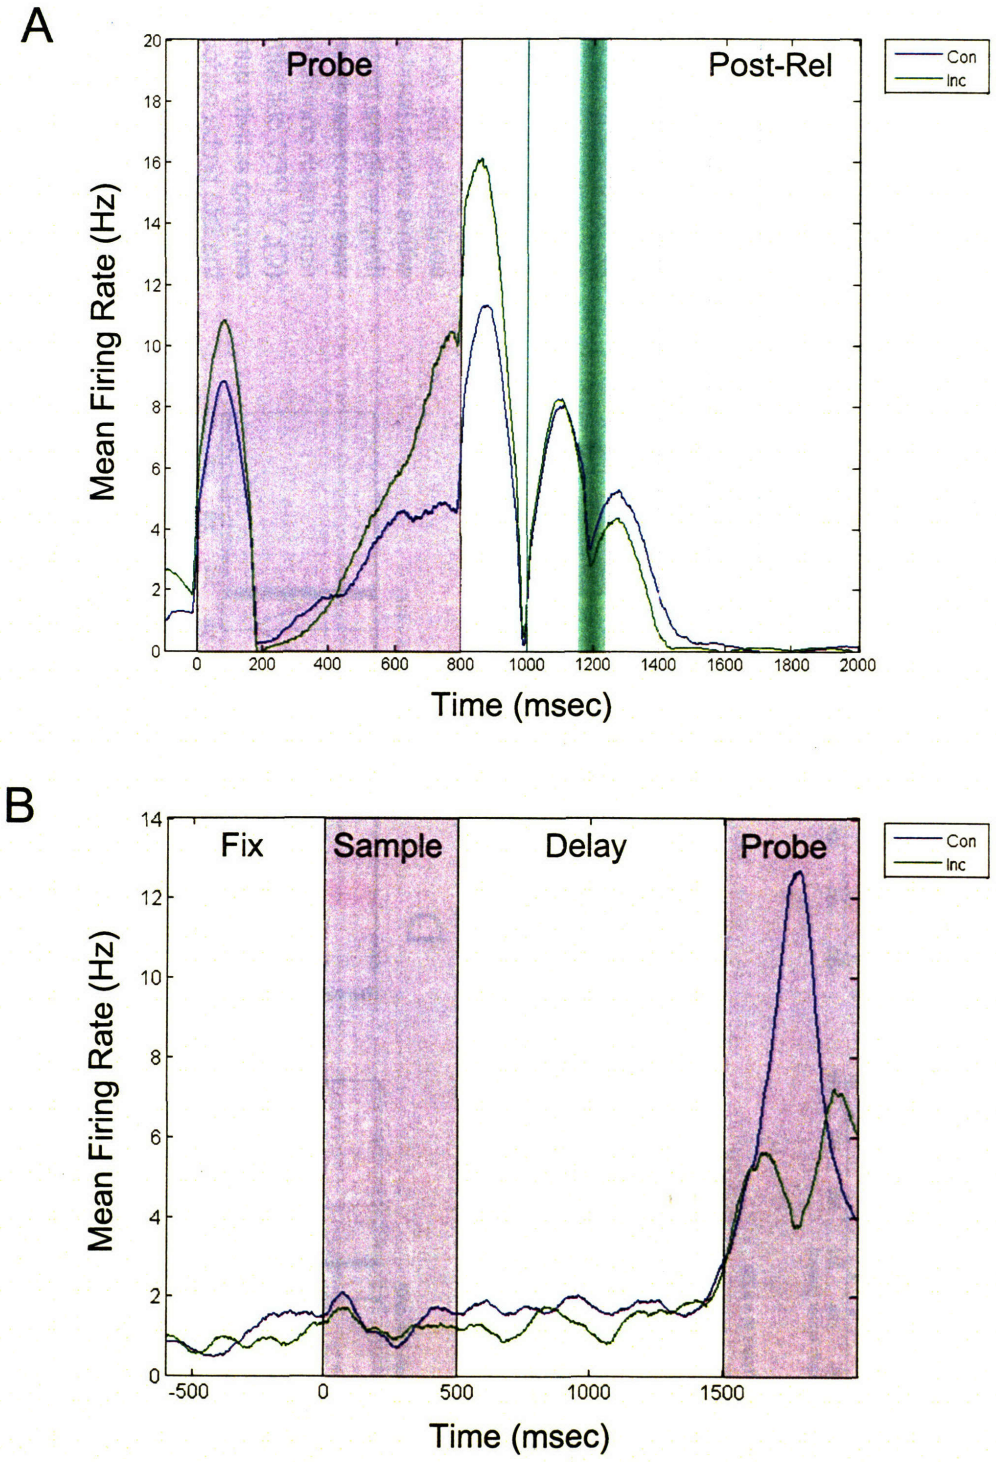

Fig. 16. Congruency-selective cells: **(A),** VLPFC cell that prefers incongruent conditions during the late probe period; **(B),** ACC cell that prefers congruent conditions during the early probe period.

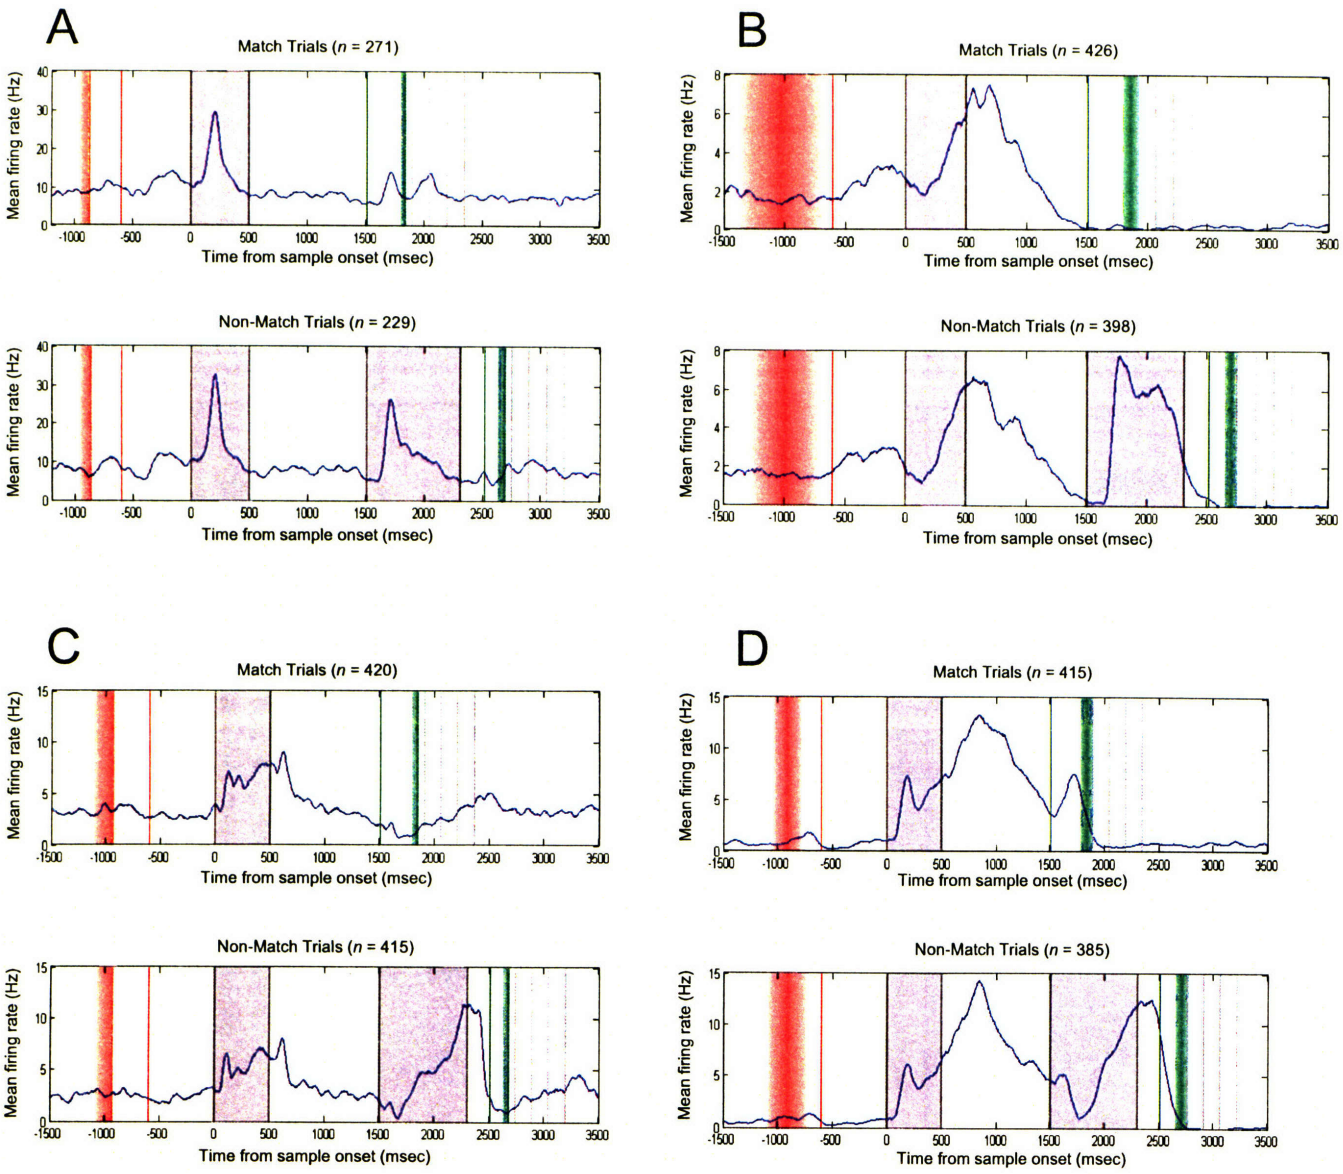

Fig. 17. Examples of cells showing significant effects of response type. In each pair of histograms, match trial data appear on top and non-match trial data appear on the bottom. Trials are aligned on sample onset time, and task events are denoted using the same conventions as for Fig. 4. **(A),** VLPFC cell with a stronger visual response to the probe object during non-match conditions; **(B),** ACC cell with a strong burst of firing while the probe was on-screen during non-match conditions; **(C),** VLPFC cell with ramping activity during the late probe period; **(D),** ACC cell showing a stronger visual response to the probe during match trials, and a strong ramp-up of firing during the late probe period of non-match trials.