

Prospective Coding for Objects in Primate Prefrontal Cortex

Gregor Rainer, S. Chenchal Rao, and Earl K. Miller

Department of Brain and Cognitive Sciences, Center for Learning and Memory, and RIKEN-MIT Neuroscience Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

We examined neural activity in prefrontal (PF) cortex of monkeys performing a delayed paired associate task. Monkeys were cued with a sample object. Then, after a delay, a test object was presented. If the test object was the object associated with the sample during training (i.e., its target), they had to release a lever. Monkeys could bridge the delay by remembering the sample (a sensory-related code) and/or thinking ahead to the expected target (a prospective code). Examination of the monkeys' behavior suggested that they were relying on a prospective code. During and shortly after sample presentation, neural activity in the lateral PF cortex primarily reflected the

sample. Toward the end of the delay, however, PF activity began to reflect the anticipated target, which indicated a prospective code. These results provide further confirmation that PF cortex does not simply buffer incoming visual inputs, but instead selectively processes information relevant to current behavioral demands, even when this information must be recalled from long-term memory.

Key words: prefrontal cortex; pair association; working memory; associative memory; physiology; vision; long-term memory; recall; monkey

When a delay is imposed between two behaviorally related events, animals can bridge this gap by remembering the event before the delay or anticipating the event after the delay. Consider a visual delayed paired associate (DPA) task in which there are two sample objects (e.g., "S1" and "S2") and two choice objects ("C1" and "C2"). If S1 is the sample, the animal must choose C1 after the delay, and if S2 is the sample, the animal must choose C2. Animals could solve this task by maintaining over the delay a memory of the samples (S1 or S2) or a memory of the anticipated choice stimuli (C1 or C2). The former (memories of events in the recent past) is called retrospective processing, and the latter (anticipation of future events) is called prospective processing (Roitblat, 1980; Honig and Thompson, 1982).

Prospective processing allows animals to optimize behavior by preparing in advance for critical events. They can, for example, act more quickly and direct attention to catch a fleeting event that might otherwise be missed. Given such advantages, it is not surprising that behavioral evidence for prospective processing has been found in several animal species, including monkeys and pigeons (Gaffan, 1977; Roitblat, 1980; Honig and Thompson, 1982; Colombo and Graziano, 1994). Studies exploring prospective processing have used associative tasks, such as the DPA task described above. These tasks thus allow perceptions and memories of the sample to be distinguished from expectation of events after the delay. For example, Colombo and Graziano (1994) trained monkeys on an auditory–visual matching task in which the samples were tones and the matches were objects associated with the tones. They found that visual, but not auditory, interference in the delay disrupted performance, which suggested that

monkeys had generated a prospective code of the anticipated object.

Neurophysiological evidence for prospective processing has been found in extrastriate visual areas thought to be involved in visual memory and perception. During the delay of visual DPA tasks, the activity of some neurons in inferior temporal (IT) cortex reflects the object expected after the delay (Sakai and Miyashita, 1991). During the delay of cross-modal DPA tasks, many visually responsive neurons in area V4 show activity after a haptic sample that seems to reflect the visual stimulus associated with it (Haenny et al., 1988; Maunsell et al., 1991).

Another region likely to be involved in prospective processing is the prefrontal (PF) cortex. It is critical for a wide range of visually guided behaviors (Goldman-Rakic, 1987; Passingham, 1993; Petrides, 1994; Fuster, 1995; Miller, 1999). Furthermore, the lateral PF cortex is interconnected with extrastriate visual cortex, including areas in which prospective processing is evident (the IT cortex and area V4). Although PF neuronal activity has been shown to reflect prospective coding for anticipated rewards (Watanabe, 1996) and actions (Quintana and Fuster, 1992; Asaad et al., 1998), its role in prospective processing of visual information has not been tested. To explore the role of the lateral PF cortex in prospective coding for objects, we trained monkeys to perform a visual DPA task.

MATERIALS AND METHODS

Subjects. Two adult rhesus monkeys (*Macaca mulatta*), one female and one male, weighing 6 and 8 kg, respectively, were used in this study. Using previously described methods, monkeys were implanted with a scleral search coil to monitor eye movements, a head bolt to immobilize the head during recording, and recording chambers (Miller et al., 1993). Penetration sites were determined using magnetic resonance imaging. The recording chambers were positioned stereotactically over the lateral prefrontal cortex such that the principal sulcus and surrounding cortex were readily accessible. All surgeries were performed under sterile conditions while the animals were anesthetized with isoflurane. The animals received postoperative antibiotics and analgesics and were always handled in accord with National Institutes of Health guidelines and

Received Feb. 19, 1999; revised April 15, 1999; accepted April 19, 1999.

This work was supported by a grant from the Whitehall Foundation. We thank K. Anderson, W. Asaad, M. Metha, A. Siapas, R. Wehby, M. Wicherski, and M. Wilson for valuable comments on this manuscript and M. Histed for expert help.

Correspondence should be addressed to Dr. Earl K. Miller, Building E25, Room 236, Massachusetts Institute of Technology, Cambridge, MA 02139.

Copyright © 1999 Society for Neuroscience 0270-6474/99/195493-13\$05.00/0

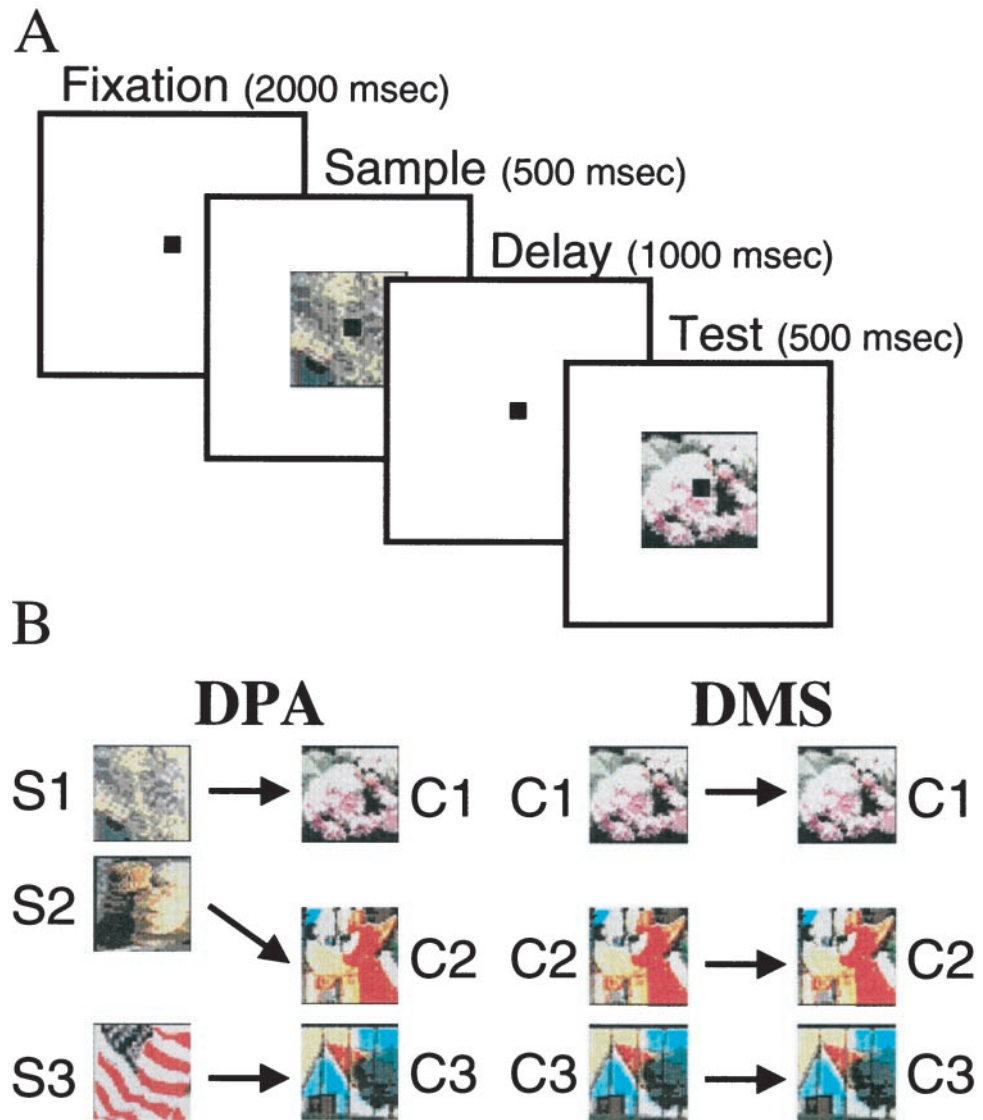


Figure 1. *A*, Schematic diagram of the DPA task. After fixating for 2 sec, a sample object was presented at the center of gaze for 500 msec. After a subsequent 1000 msec delay, a test object was presented, which could be either a target or nontarget. The sequence of events was identical for the DMS task. *B*, Sample–target associations for both tasks. For each task, the test stimulus could either be a target (match) or nontarget (non-match). Stimuli for both tasks were chosen such that two of the three were similar to each other. Similarity is depicted as vertical proximity, such that S1 and S2, and C2 and C3 were similar and each were different from S3 and C1, respectively.

the recommendations of the Massachusetts Institute of Technology Animal Care and Use Committee.

Recording techniques. Neural activity was isolated using arrays of two to eight independently movable tungsten microelectrodes (FHC Instruments, Bowdoin, ME). The electrodes were advanced using custom-made screw mini-microdrives mounted on a plastic grid (Crist Instruments, Damascus, MD). Neural activity was amplified, filtered and stored for off-line cluster separation (DataWave Technologies, Longmont, CO). We did not prescreen neurons for task-related responses. Instead, we advanced each electrode until the activity of one or more neurons was well isolated and then began data collection. This procedure was used to ensure an unbiased estimate of prefrontal activity. In any given session, we were able to simultaneously record the activity of up to 18 individual neurons. On average, we collected data from approximately six cells per recording session.

Behavioral task. Monkeys performed alternate blocks of a DPA task and a delayed match-to-sample (DMS). Each block consisted of 50 trials, and blocks were presented in alternation for an average of three repetitions of each type of block (a total six blocks). Both DPA and DMS trials used the following protocol, shown schematically in Figure 1*A*. Each trial began when the monkey grasped a metal bar and initiated fixation on a small (0.3°) spot of light (fixation spot). Their gaze needed to remain within 1.5° of the fixation spot throughout the trial. After 2000 msec of fixation, a sample object was presented for 500 msec at the center of gaze. After a 1000 msec delay, a test object was presented for 500 msec. If the test object was the correct choice object (the target), animals were

required to release the metal bar within 500 msec to obtain a juice reward. The fixation point was superimposed on the sample and test stimuli so that the monkeys could maintain central gaze during stimulus presentation.

For both tasks, a test object could either be a target or a nontarget. On DMS trials, it was a target if it was identical to (i.e., matched) the sample object. Three stimuli were used, each as a sample–target and as a nontarget (nonmatching) test object on different trials. On DPA trials, a test object was the target if it was the object that had been associated with the sample object during several months of training. If the test object was a nontarget, monkeys were required to hold the bar for the duration of test object presentation and for an additional 1000 msec delay (data not shown in Fig. 1*A*), which was always terminated by the presentation of a target. Monkeys did not have to hold object-specific information in memory during this second delay; it was included in both the DMS and DPA task to ensure that monkeys made a behavioral response on every trial. That way, we could be sure that the monkey was actively engaged by the task on every correctly performed trial. The test object was always extinguished as soon as the monkey released the metal bar. Rewards were withheld on incorrect responses, and a 2 sec time-out period was imposed. In both tasks, the same behavioral response (lever release) was required on each trial, and it always led to the same reward (apple juice). In contrast, the samples and expected targets were varied randomly from trial to trial.

Stimuli. All objects were multicolored pictures of real world objects, each $\sim 2 \times 2^\circ$ in size. Their average luminance was 15.3 cd/m^2 (SD of 0.7

cd/m²) and had an eight-bit color depth. They were presented at the center of gaze on a uniform background with a luminance of 0.8 cd/m². We used three pairs of associated objects in the DPA task. The samples will be referred to as S1, S2, and S3 and the choices as C1, C2, and C3. Monkeys needed to choose C1 if S1 was the sample and so on (i.e., S1→C1, S2→C2, S3→C3). The three choice stimuli were also used as nontarget test stimuli. In the DMS task, the three choice stimuli were used as both samples and matches (i.e., C1→C1, C2→C2, C3→C3). They also appeared as nonmatching test stimuli on different trials. Note that, across the DMS and the DPA tasks, monkeys always chose the same targets (C1, C2, C3). The tasks only differed in which samples preceded these targets. Thus, the monkeys did not need explicit cueing as to which task (DPA or DMS) they were performing. This type of design in which a larger group of samples is used to instruct choice of a smaller set of targets has been shown previously to foster prospective processing (Zentall et al., 1993).

The objects were chosen such that two of the samples (S1 and S2) and two of the targets (C2 and C3) were similar. The two similar samples (S1 and S2) were associated with two dissimilar choice stimuli (C1 and C2), and the two dissimilar samples (S2 and S3) were associated with two similar choices (C2 and C3). This relationship is depicted in Figure 1*B*. Similarity was assessed in pilot studies with human observers. The stimuli that to humans looked physically similar also seemed so to the monkeys; their pattern of errors in the DMS task reflected this similarity (see Results and Fig. 3*B*). The “confusion matrix” design used in the DPA task has been used in previous studies to explore whether animals adopt a prospective strategy or a retrospective (sensory-related) strategy (Gaffan, 1977; Roitblat, 1982). The pattern of errors is thought to reflect which stimulus (the sample or target) the animal actively holds in working memory. If errors reflect the similarity of the targets (C2 and C3), this would suggest a prospective code of the targets. In contrast, if errors reflect the similarity of the samples (S1 and S2), this would suggest that animals were maintaining a sensory-related code of the samples. In the former (prospective) strategy, recall of the targets from long-term memory would take place before presentation of the test object. In the latter (retrospective) strategy, recall would be initiated by presentation of the test object, when monkeys need to determine whether it is the correct target.

We used complex objects because previous studies have shown they readily elicit stimulus-selective activity from PF neurons (Miller et al., 1996). Complex stimuli, however, do not lend themselves to an objective measure of similarity. Thus, to control for the possibility that the similar samples were more similar to one another than the similar targets were to one another (or vice versa), we trained one monkey to perform the task with sample and target objects reversed. That is, the objects used as C1, C2, and C3 for one monkey were used as S1, S2, and S3 for the other and vice versa. This counter-balances the effects of similarity of specific stimuli across the two monkeys. It also allows us to control for any misclassification of stimuli with regard to their similarity. Any misclassification would result in a different pattern of results across the two monkeys. As will be seen below, this was not the case. Similar results were seen in both monkeys.

Data analysis. We calculated the firing rates of PF neurons in four epochs. These epochs were selected by inspecting the average activity of all cells in the entire population, as shown in Figure 2*B*. Neural activity during sample presentation was calculated over the entire 500 msec period of sample presentation, offset by 100 msec to compensate for the latency of PF neurons. For analysis of activity during the delay, we used the last 700 msec of the 1000 msec delay. As in previous studies, we excluded the first part of the delay to exclude any effects related to the offset of the sample. Analysis of responses to the test object (test epoch) was restricted to a 150 msec period starting 100 msec after test object presentation (a period before the bar release) to exclude any effects related to the behavioral response. Baseline or spontaneous activity was assessed over 800 msec before sample onset. We used data from correctly executed trials only. There was an average of 50 correct trials per sample condition. The monkeys did not accumulate enough error trials to allow analysis of neural activity during errors.

For some of the comparisons of neural selectivity, we computed a measure of the difference of activity between the tasks normalized by the variance:

$$d = |F_1 - F_2| / \sqrt{\sigma_1^2 + \sigma_2^2}$$

where F_i is the average firing rate, and σ_i^2 is the variance of neural activity to the i -th sample object. The vertical bars denote the absolute

value operator. The normalized difference d allows the direct comparison of activity differences for neurons with widely different baseline firing rates. If neural activity reflected prospective processing, then, across the two tasks, neural responses to samples predicting the same targets should be similar. For example, the responses to sample S1 (DPA task) should be similar to responses to sample C1 (DMS task), because in both cases the monkey needs to select target C1. We assessed this by computing two numbers: a normalized difference to associated objects (NDA) and a normalized difference to nonassociated objects (NDN). The NDA was the average difference in activity between DMS and DPA trials that predicted the same targets. We computed the absolute values of the mean differences in activity between trials in which S1 and C1 were samples (both predict target C1), between trials in which S2 and C2 (target C2) were samples, and between trials in which S3 and C3 (target C3) were samples and then calculated their average. In contrast, the NDN was the mean of the differences in activity between all possible pairs of samples that predict different targets (e.g., S1 vs S2, S1 vs S3, S1 vs C2, S1 vs C3, etc). Thus, the NDA is a measure of how similar activity is on trials in which the same target is expected, whereas the NDN is a measure of the general selectivity of the neuron, computed by comparing trials in which different targets were expected.

RESULTS

Behavioral data

Pattern of errors

Two monkeys completed a total of 52 recording sessions, with a mean number of 342 correct trials per session. For behavioral analyses, we computed the average error rate for each session. The monkeys were well trained on both the DPA and DMS tasks; their average performance across the sessions was >95% correct. They made few false negative responses; on trials in which the test object was a target, average performance across all sessions was 98% on both tasks. Most errors were false alarms: that is, incorrect responses to a nontarget object as if it were a target. False alarm rates for both monkeys from the DPA task are shown in Figure 3*A*. The pattern of errors was consistent with previous reports of prospective errors (Gaffan, 1977; Roitblat, 1982; see caveats in Discussion). That is, errors were based on the similarity of the target objects and not on the similarity of the sample objects. For example, after sample S3, monkeys made many errors responding to test C2, which was physically similar to C3 (Fig. 3*A*, *second horizontal bar from top*). On the other hand, few sensory-related errors, those confusing similar samples, were made. For example, there were few false alarms to test C2 when it followed sample S1, although S1 was physically similar to S2, the sample associated with C2 (Fig. 3*A*, *third horizontal bar from top*). Using the error rate for each session as an observation, we found that the difference between prospective and sensory-related errors was significant for both monkeys (t test; $p < 0.0001$).

An examination of performance during the DMS task confirmed that the physical similarity of the target objects was indeed contributing to the pattern of errors observed during the DPA task. False alarm rates for both monkeys from the DMS task are shown in Figure 3*B*. More errors were made when the test object was similar to the target than when it was dissimilar. Overall, the DPA task was more difficult than the DMS task. On average, monkeys made more errors on the DPA task than the DMS task (t test; $p < 0.00001$), making false alarm errors on 19% and 11% of trials, respectively.

Reaction time

When the test object was a target, monkeys needed to release the bar during the 500 msec of the presentation of this target. Histograms of these reaction times are shown in Figure 4. This figure is based on 9665 correctly executed trials from both monkeys.

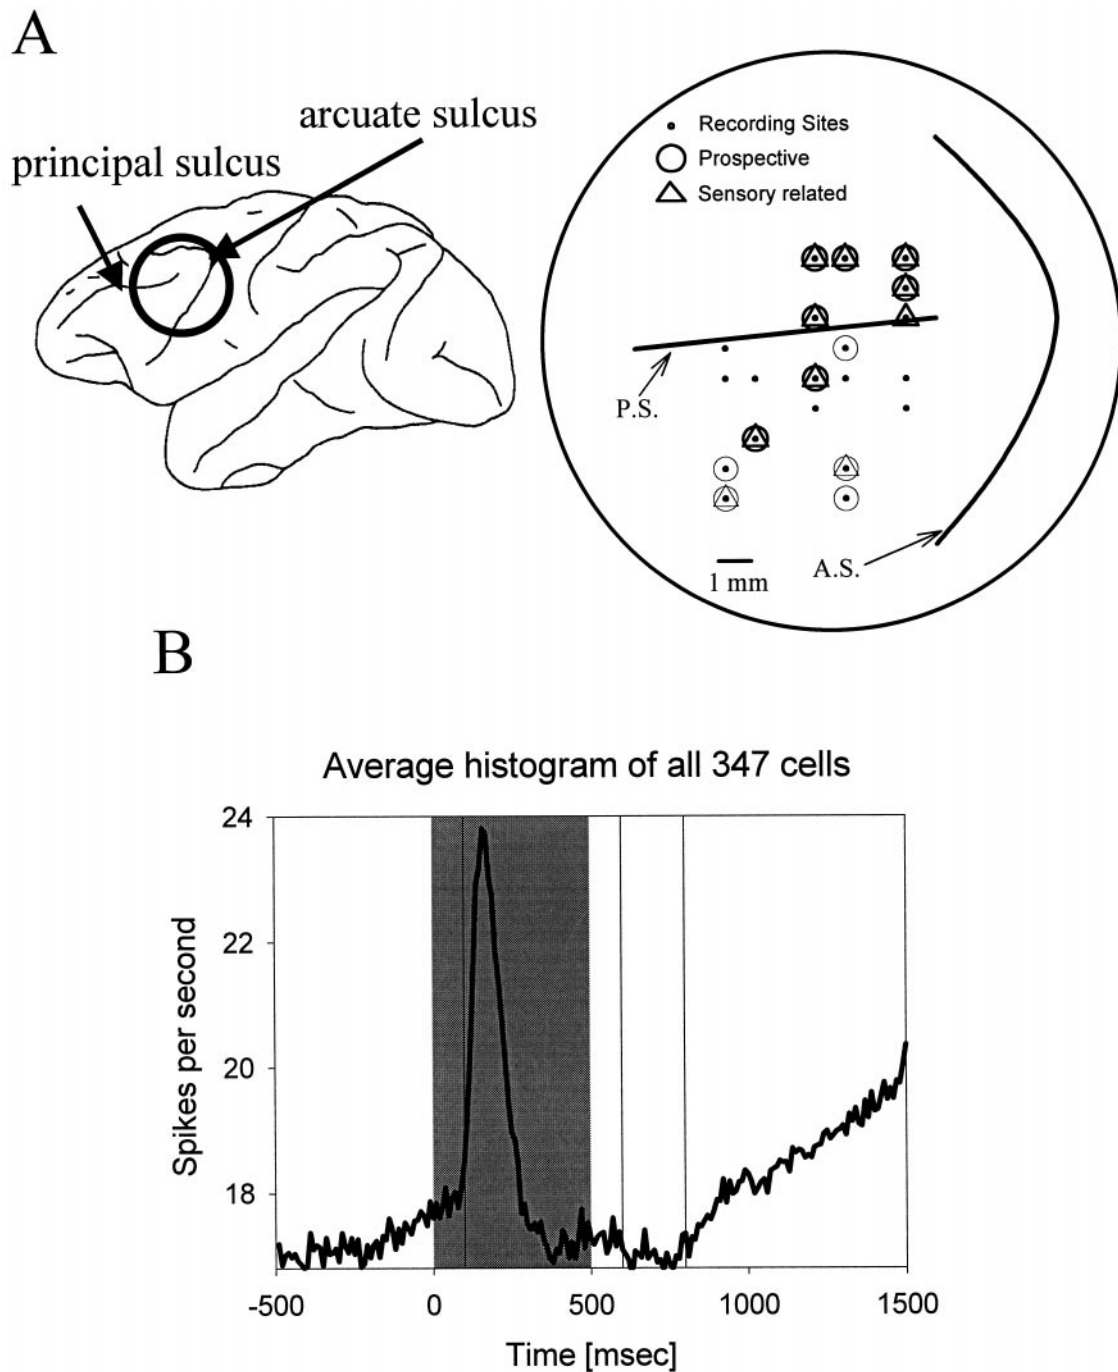


Figure 2. *A*, Location of recording sites. The recording sites were located around the principal sulcus and on the inferior convexity of the prefrontal cortex. Recording sites in the two monkeys were intermixed and nonoverlapping. The circles and triangles depict where prospective and sensory-related cells were found. Recording sites from monkey A are shown as thick symbols, and those from monkey B are shown as thin symbols. *B*, Average histogram for all recorded cells. The curve represents the mean firing rate (bin width, 10 msec). The shaded gray area represents the period of sample object presentation. The black vertical lines delineate the sample and delay epochs described in Materials and Methods.

First, note that reaction times were shorter to target C1 than to C2 and C3. This presumably reflects the fact that C1 was dissimilar from C2 and C3, and thus C1 could be discriminated more rapidly. More importantly, note that, for the same targets, reaction times were similar across the tasks. There was no significant difference between reaction times in the DPA (solid lines) and DMS (dashed lines) (t test; $p > 0.1$). In fact, reaction times were identical to target C1 on DMS and DPA trials, although they were

shorter to this target than the other targets. The reaction time was 318 ± 27 (mean \pm SE) msec for C1, 342 ± 32 msec for C2, and 338 ± 29 msec for C3. The absence of a reaction time difference between DPA and DMS tasks is suggestive of prospective processing. If monkeys were simply maintaining a memory of the sample objects during the delay, one might expect the target decision to take longer in the DPA task because the recall of the target would have to take place during the presentation of the test

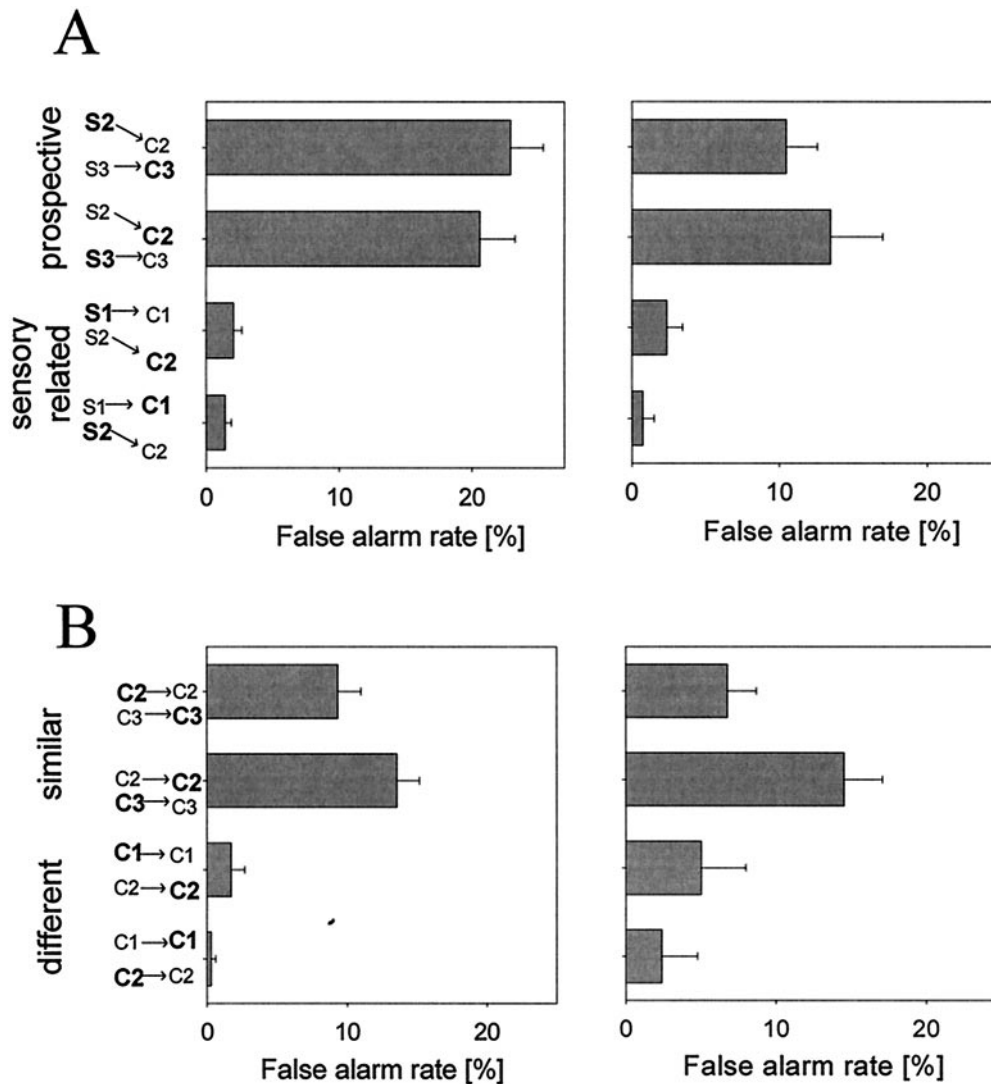


Figure 3. Average false alarm rates pooled across both monkeys for DPA (*A*) and DMS (*B*) tasks. Each bar gives the false alarm rate for the condition described on the left. Shown in bold are the sample and the test object chosen in error for that trial. The arrow points to the correct target for that sample. Thus, the top bar in *A* shows the error rate when the sample was S2 and the monkey chose C3 instead of the correct target S2. The error bars show the SEM. Performance of monkey A is shown on the left, and monkey B is shown on the right.

object. The fact that reaction times were identical across the two tasks suggests that the recall step had already been completed before the test object presentation; that is, the monkeys were using prospective processing.

Neural data: general properties

A total of 347 neurons (252 from monkey A, 95 from monkey B) were recorded in the left lateral prefrontal cortex of two monkeys. Neurons were not screened for responsiveness and thus represent an unbiased sample of prefrontal activity. Figure 2*B* illustrates the average activity across the entire population of 347 PF cells. Note that there is a marked increase in activity over the delay as the time of test object presentation approaches. Such “climbing activity” has been suggested to reflect prospective coding (Quintana and Fuster, 1992; Miller et al., 1996). In principle, climbing activity could reflect any end-of-trial event, the test object, the reward, or the behavioral response. In this task, however, the behavioral response and reward were constant across trials, whereas the expected target object was varied. As will be demonstrated below, many PF neurons showed activity near the end

of the delay that varied with the anticipated test object. Of course, expectation of the reward or of the behavioral response could have contributed to any nonselective delay activity. In this experiment, modest climbing activity was also evident before sample presentation (Fig. 2*B*). Although the monkey could not predict which sample would appear, it could predict when it would appear. This activity may reflect that expectation.

We assessed responsiveness by comparing activity within a trial (sample and delay intervals) to baseline, or spontaneous, activity using paired sample *t* tests (evaluated at $p < 0.01$). Based on this criterion, a total of 181 of 347 cells (52%) were responsive. To assess selectivity, we performed ANOVAs (evaluated at $p < 0.01$) on the sample and delay periods separately for all neurons. Both tasks were treated together, so that there were six levels in the ANOVA corresponding to each of the possible sample objects in DMS and DPA. We tested for selectivity for samples rather than targets because it would reveal cells that showed any selective activity, whether for the samples or targets. According to this criterion, 146 cells showed activity that varied significantly with




Figure 4. Histograms of bar release reaction times to the test object presentation pooled across both monkeys. Solid lines and dashed lines correspond to DPA and DMS tasks, respectively. The bin width was 10 msec.

the sample during the sample period, 149 during the delay, and 87 during both periods. Because the main focus of this report will be on changes in neural activity as the animal perceives the sample and then waits for the test object presentation, most analyses will focus on the 87 cells (58 from monkey A, 29 from monkey B) that were selectively modulated during both epochs.

Figure 2*A* shows the locations where these neurons were recorded. Object-selective cells were found both near and around the principal sulcus (area 46) and on the inferior convexity (area 12/45). Both regions receive inputs from the inferior temporal cortex, a region known to analyze object features (Barbas, 1988; Ungerleider et al., 1989). Area 46 also receives inputs from the posterior parietal cortex (Goldman-Rakic and Schwartz, 1982), a region that also contains object-selective neurons (Sereni and Maunsell, 1998). Previous studies have shown that these prefrontal regions contain object-selective neurons (Watanabe, 1981; Fuster et al., 1982; Wilson et al., 1993; Rao et al., 1997; Asaad et al., 1998; Rainer et al., 1998a,b; O Scalaidhe et al., 1998).

Prospective coding

Single cell analyses

We defined prospective activity as that reflecting a forthcoming target object. Because the target must be recalled from memory, prospective activity is mnemonic in nature. Sensory-related activity is defined as that reflecting the samples and thus could either be a result of a visual response to the sample and/or its maintained memory. In other words, sensory-related activity reflects a current or recent sensory event (the sample), whereas prospective activity reflects the expectation of a sensory event (the target). Both types of activity were evident in the PF cortex. The cells illustrated in Figure 5*A* had activity consistent with prospective coding. The cell on the *left* was nonselective during sample presentation (ANOVA; $p = 0.232$). Then, during the second half of the delay, it showed high levels of activity (over twice as high as baseline firing rate) when the monkey needed to choose target C1, irrespective of whether the preceding sample had been C1 itself (DMS) or its associate S1 (DPA). Activity was lower for all other DMS and DPA conditions. Thus, the activity of

this cell seemed to reflect the target. This activity did not appear to be related to other end-of-trial events, such as the reward or behavioral response. The reward and behavioral response were the same for all trials, whereas the activity of this neuron varied significantly with which object was the target. On DPA and DMS trials, the activity was significantly greater when the target was C1 than during other DPA and DMS conditions (t test; $p < 0.0001$). On the *right* of Figure 5*A* is another prospective cell. On DPA trials when S1 was the sample and C1 the target, activity in the second half of the delay climbed to the same level of activity evident throughout the delay of DMS trials when C1 was both the sample and target.

Whereas the cells in Figure 5*A* showed activity that reflected the expected target, the cells shown in Figure 5*B* had activity consistent with an object-selective sensory response to the sample. During the DPA task, the cell on the *right* of Figure 5*B* showed a visual response to sample object S3 that lasted into the early delay. In contrast, on DMS trials when C3 was the sample (and target), there was little or no activity. This suggests that the high level of activity on DMS trials reflected a visual response to sample S3 and not an expectation of target C3. Similarly, the cell on the *left* of Figure 5*B* showed a strong response early in the delay of DPA trials when S1 was the sample (and C1 the target) but almost no activity on DMS trials when C1 was the sample and target.

In the DMS and the DPA tasks, the targets were the same; the tasks only differed in which stimuli were used as samples. A cell with prospective activity, then, should show a similar pattern of selectivity on DMS and DPA trials that required the same target to be chosen. For example, say that during the DPA task, a neuron shows the greatest activity on trials in which S1 was the sample and C1 the target. If this activity reflected a prospective code, it would reflect C1 rather than S1. Thus, this neuron should also show a high level of activity on DMS trials in which C1 was both the sample and the target. A cell with sensory-related activity, in contrast, would show no correspondence between the tasks because they use different samples. To classify whether single cells had activity patterns consistent with prospective or sensory-related coding, we used a two-way ANOVA. One factor was the target for a given trial (TARGET factor, C1, C2, or C3), whereas the other factor was which task the animal was performing (TASK factor, DMS or DPA). If cells with target-selective activity were involved in sensory processing (i.e., activity was related to the samples rather than to the targets), there should be an interaction between TARGET and TASK factors because the sample that preceded a given target would depend on which task the animals were performing. On the other hand, if cells with target-selective activity were involved in prospective coding (i.e., activity was related to the targets rather than to the samples), there should be no interaction between TASK and TARGET factors. That is, the task would not influence the pattern of selectivity of the cell because both tasks required that the same targets be chosen. Thus, prospective activity was defined as having a significant effect of TARGET ($p < 0.01$) and no significant interaction between TARGET and TASK ($p > 0.1$). Sensory-related activity was defined as exhibiting both effects of TARGET and a significant interaction between TARGET and TASK ($p < 0.01$).

Figure 6 shows how many cells were classified as prospective or sensory-related in each of the epochs using these criteria. Figure 6*A* shows data from the entire population of 181 responsive cells.

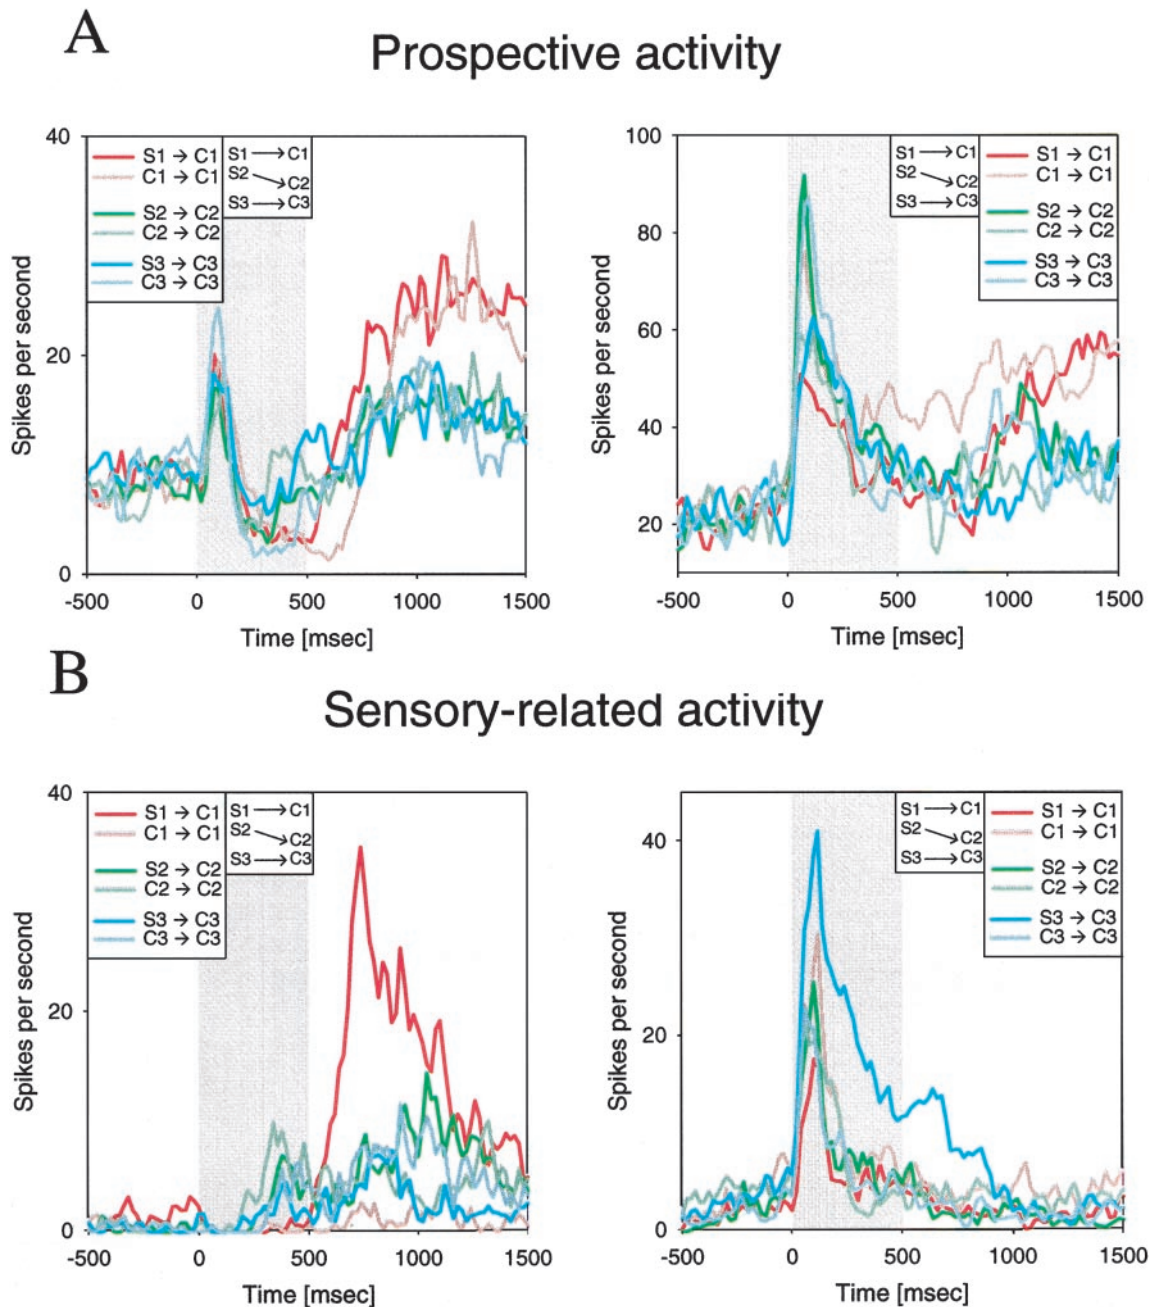


Figure 5. Single cell examples of prospective (*A*) and sensory-related (*B*) memory. The cells on the *right* were recorded from monkey A, and the cells on the *left* were recorded from monkey B. Each colored line represents the firing rate of a single neuron after the sample object shown in the legend. The gray shaded area represents the period of sample object presentation, which is followed by 1000 msec of delay shown to the *right* of the gray area. Histograms were binned at 20 msec and smoothed with a 20 msec Gaussian.

Figure 6*B* is based on the 87 cells that showed selective activity in both epochs. Both figures show similar trends. There was a decrease in the number of sensory-related cells and an increase in the number of prospective cells from the sample to the delay (χ^2 , $p < 0.01$). This trend was also apparent in cells from both monkeys considered separately. This suggests that activity related to processing of information about the samples dominated during sample presentation, but during the delay, there was an increase in prospective processing related to the anticipated targets. Figure 2*A* shows the locations in which we found cells that showed either sensory-related or prospective delay activity. These cells were intermixed throughout the recording sites.

It is important to note that the classification of activity into “sensory-related” or “prospective” categories is somewhat artificial. In this experiment, as in most others, neuronal selectivity varied along a continuum. In addition, cells selective for one of the sample objects but none of the target objects must necessarily be classified as sensory-related, although they do not argue against prospective coding for objects that were not included as targets. The remaining analyses, then, will focus on average activity across a population of cells, without categorizing activity as sensory-related or prospective. As we will see below, the trend for neural activity to “switch” from sensory-related to prospective coding can be observed at the population level.




Figure 6. Histograms of cell counts classified as prospective and sensory-related in the population of 181 responsive cells (**A**) and in the 87 cells that showed selective activity in both the sample and the delay intervals (**B**). Classification is based on an ANOVA (see Results).

Population analyses

In the DMS task, two of the samples, C2 and C3, were physically similar and were different from C1. We first wanted to establish whether this physical similarity, which manifested itself in the monkeys' behavior, was also reflected in PF activity. To assess the similarity of neural activity, we computed for each cell average differences in normalized firing rate (see Materials and Methods) between trials for each of the three possible combinations of sample objects (C1–C2, C1–C3, and C2–C3). If physical similarity was reflected in neural activity, these values would be small for the physically similar stimuli (C2–C3, or the SIMILAR comparison) and larger for physically dissimilar stimuli (C1–C2 and C1–C3, or the DIFFERENT comparisons). Figure 7*A* shows that this is indeed the case. It illustrates the average differences for the 87 neurons that exhibited selectivity in both intervals. The difference values are smaller to similar samples (*red line*) than to dissimilar samples (*black and blue lines*). We computed a two-way repeated measures ANOVA that used comparison and time bin as factors. Differences between the two DIFFERENT comparisons (*black and blue lines*) did not reach significance at any time point (repeated measures ANOVA; $p > 0.1$). However, the differences between the SIMILAR and each of the DIFFERENT comparisons (*black and red lines*, *blue and red lines*) were significant for the time bins that included 1100 to 1500 msec (repeated measures ANOVA with *post hoc* contrasts; $p < 0.01$). The fact that the three curves do not cross indicates that selectivity is similar throughout the DMS trial. In other words, it indicates that, on average, stimulus preference was maintained across all three epochs. We conclude that physically similar sample objects tend to elicit similar neural activity in PF cortex and that this difference is apparent in both visual response and mnemonic activity (sample and delay epochs, respectively).

Having established the effect of physical similarity in the DMS task, we now turn to the DPA task. The key concept in the DPA task is that physically similar sample objects (S1 and S2) instruct choice of dissimilar targets (C1 and C2) and that dissimilar samples (S2 and S3) instruct similar targets (C2 and C3). Thus, there are two critical comparisons. One is between the activity on trials in which two similar sample objects instructed two dissimilar target objects (S1→C1 vs S2→C2). This will be referred to as the SIMILAR → DIFFERENT comparison. The other comparison is between activity during trials in which two dissimilar sample objects instructed two similar target objects (S2→C2 vs S3→C3; i.e., the DIFFERENT → SIMILAR comparison).

If activity reflected sensory-related processing, then the comparisons should reflect the similarity of the samples and not the targets. That is, the DIFFERENT → SIMILAR (Fig 7*B*, *blue line*) comparison should yield larger difference values than the SIMILAR → DIFFERENT comparison (*red line*). If prospective processing predominated, then the reverse should occur; the comparisons should reflect the similarity of the targets. That is, the SIMILAR → DIFFERENT comparison (*red line*) should yield larger difference values than the DIFFERENT → SIMILAR comparison (*blue line*). Figure 7*B* shows the plots of the average values for the 87 neurons that showed selectivity in the sample and delay intervals. During sample presentation and the early portion of the delay, activity reflected the samples. The DIFFERENT → SIMILAR values (*blue line*) are higher than the SIMILAR → DIFFERENT values (*red line*). Near the end of the delay, however, this situation reverses. Then, activity reflected the targets. The SIMILAR → DIFFERENT values (*red line*) are larger than the DIFFERENT → SIMILAR values (*blue line*). The remaining comparison is between the activity on trials in which dissimilar samples instruct dissimilar targets (S1→C1 vs S3→C3, the DIFFERENT → DIFFERENT comparison). Here, we expect difference values to remain high throughout the trial. As can be seen in Figure 7*B*, they do (*black line*).

The time bin 100 msec after sample offset (centered on 600 msec in Fig. 7*B*) showed the most consistent sensory-related activity. The SIMILAR → DIFFERENT values (*red line*) were significantly lower than both the DIFFERENT → SIMILAR values (*blue line*) and the DIFFERENT → DIFFERENT values ($p < 0.01$). There was no significant difference between the latter two values ($p = 0.24$). In contrast, the last two time bins of the delay (centered on 1400 and 1500 msec in Fig. 7*B*) showed evidence of prospective activity. The DIFFERENT → SIMILAR values (*blue line*) were significantly less than the SIMILAR → DIFFERENT values (*red line*) and the DIFFERENT → DIFFERENT values (*black line*) (both $p < 0.01$). At these times, there was no significant difference between the latter two values (*blue vs black lines*; $p > 0.1$). Thus, in the DPA task, sensory-related activity was strongest just after sample offset, and prospective activity was strongest near the end of the delay. This is consistent with the notion that prospective coding increases as the time for the expected target approaches.

To investigate prospective activity in more detail, we tested whether samples that predicted the same targets (e.g., S1 and C1, which both predict target C1) elicited similar activity relative to samples that predicted different targets (e.g., S1 and C2). For both the sample and delay epochs, we computed two numbers for each cell: NDA and NDN (see Materials and Methods). The NDA value reflects similarity of activity on DPA and DMS trials that used associated objects as samples (e.g., S1 and C1); that is, when the same target was expected (e.g., C1). The smaller the differ-

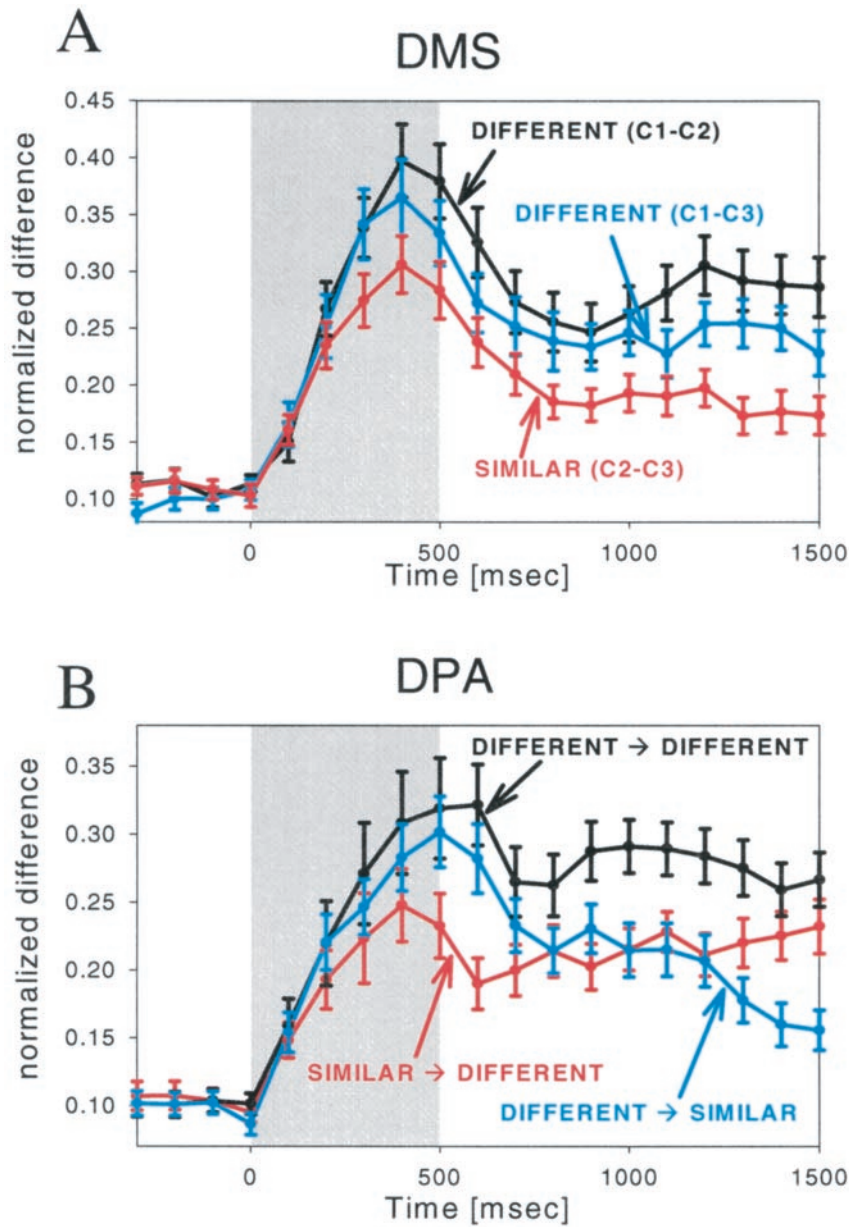


Figure 7. Time course of differences in firing rate between different combinations of conditions in the DMS (A) and DPA (B) task. Each plot corresponds to the average normalized difference in firing rate between conditions described in the legend (see Materials and Methods). Differences were calculated using a bin width of 100 msec. The average differences and SEs are shown for the population of 87 cells that showed selective activity in the sample and delay epochs. The gray shaded area represents the time of sample object presentation.

ence, the more similar the activity. Of course, if a cell was not strongly stimulus-selective, NDA would also be low because activity would tend to be similar on all trials. Thus, it is important to also have a measure of selectivity of each neuron to other (nonassociated) objects. The NDN provides this. It reflects the average difference in activity between trials that used nonassociated samples: that is, those that predicted different targets. The greater the NDN value, the more selective the activity. The lower NDA value relative to the NDN value, then, the more similar was the activity of the cell when the same targets were expected relative to when different targets were expected. In other words, the lower the ratio of NDA to NDN, the more prospective was the activity of the cell.

Figure 8, A and B, shows the NDA plotted against the NDN values in the sample and delay epochs for each of the 87 cells that showed selectivity in both epochs. The black symbols represent the cells that met the single neuron criteria for prospective coding in that epoch (the two-way ANOVA described above). Neurons

from monkey A are depicted as circles, and neurons from monkey B are depicted as triangles. The insets show histograms of the distribution of the NDA minus NDN values for each cell (NDA – NDN). More similarity in activity on trials that predicted the same targets (i.e., prospective coding) should yield lower NDA values relative to NDN values and thus more data points falling in the top left of the scatterplot and more values in the distribution falling to the left of zero.

We again found that prospective activity was more prevalent in the delay than in the sample interval. During the sample epoch, only 11 of 87 cells (13%) (Fig. 6B) were classified as prospective, and the distribution of NDA – NDN values was not significantly different from zero (*t* test; *p* = 0.70). Many of the data points obtained from the sample epoch cluster around the diagonal, indicating that selectivity to associated and nonassociated samples was approximately similar. During the delay, however, more cells (26 of 87, or 30%) were classified as prospective, and more data points fall to the left of the diagonal. This indicates that, for

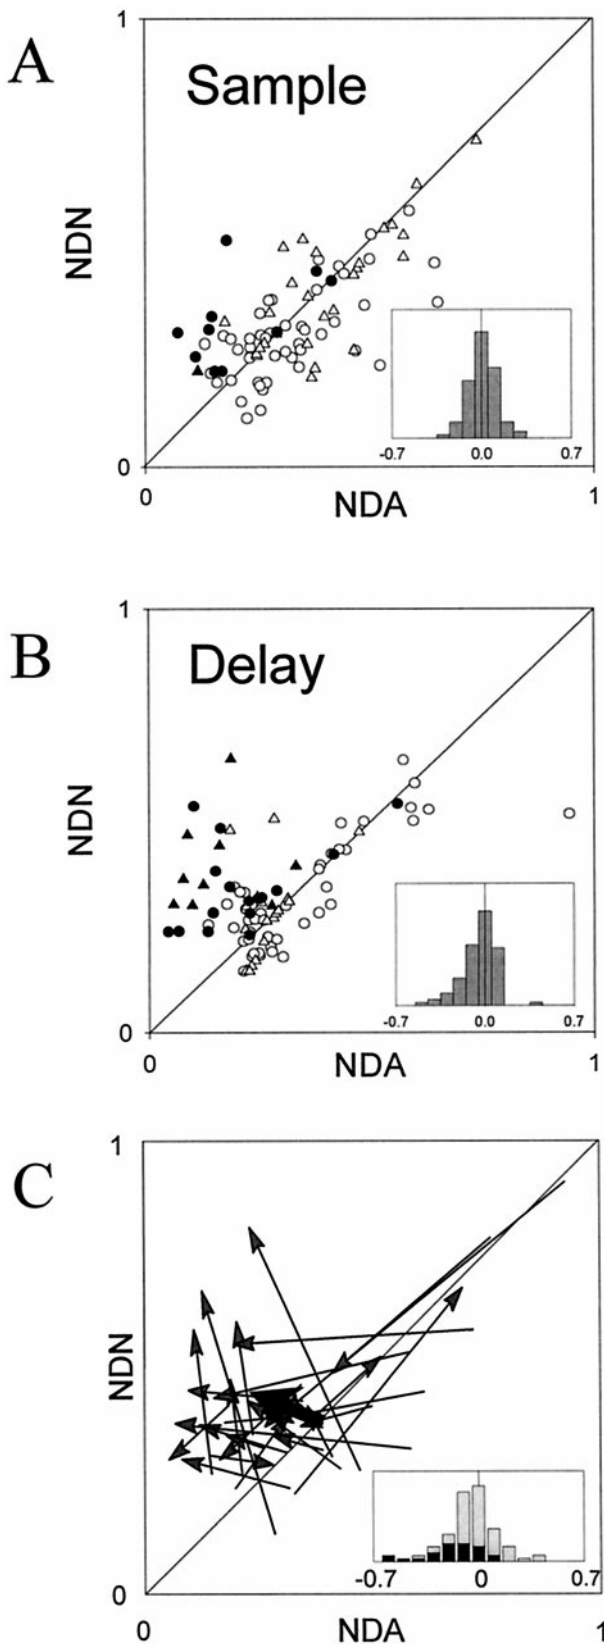


Figure 8. Scatterplots showing the normalized difference of each cell to associated objects (NDA) plotted against the average normalized difference to nonassociated associated objects (NDN) for the sample (*A*) and delay periods (*B*). The insets in *A* and *B* show the distribution of NDA – NDN values. Neurons from monkeys A and B are shown as circles and

triangles, respectively. The *black symbols* represent cells classified as prospective using an independent method (two-way ANOVA; see Results). *C*, Vectors connecting change in corresponding NDA and NDN values from the sample to the delay interval, i.e., the change in the data point position for each cell from *A* to *B*. This is shown for the cells classified as prospective during the delay period (*gray arrows*). The *black arrow* represents the average vector. The inset shows the distribution of $\delta\text{NDN} - \delta\text{NDA}$ values (see Results) for the population of 87 cells with selective activity in both intervals (*gray bars*) and for the cells classified as prospective in the delay (*black bars*).

these cells, activity after associated samples (that predicted the same targets) was more similar than activity to nonassociated samples (that predicted different targets). This is further illustrated by the distribution of the NDA – NDN values, which was significantly shifted to the *left* of zero (*t* test; $p = 0.0014$). If the cells were simply becoming less selective, the data points in the scatterplot would have fallen along the diagonal closer to the origin, and the distribution of NDA – NDN would center at zero. For both epochs, the NDA – NDN distributions computed with neurons from the two monkeys were not significantly different (*t* test; $p > 0.1$).

Figure 8*C* illustrates, for the 26 cells that met the statistical criterion for prospective coding in the delay, the change in their relative NDA and NDN values from the sample to the delay epoch. That is, it shows the change in position of a cell's data point from Figure 8*A* (sample epoch) to Figure 8*B* (delay epoch). A leftward shift indicates a decrease in NDA, i.e., an increase in similarity of activity to samples predicting the same targets. Almost all showed a shift toward the *left* or *top left* of the scatterplot, toward values indicating prospective coding. Note that there is relatively little decrease in NDN (a shift downward). This means that, even as activity to associated samples becomes more similar in the delay, activity to nonassociated samples remains different or (in the case of upward shifts) becomes even more different. In fact, the average vector (*large black arrow*) illustrates that, on average, NDA values decrease as NDN values increase from the sample to the delay. That is, as responses to associated samples become more similar, responses to nonassociated samples become more dissimilar. Figure 8*C*, inset, shows a histogram of $\delta\text{NDN} - \delta\text{NDA}$: that is, the difference between the NDA – NDN distributions from the sample and delay epochs. Note that the distribution is shifted to the left, especially for the 26 cells that were significantly prospective in the delay (*black bars*), indicating an increase in prospective coding from the sample to the delay epochs. The shift was significantly different from zero for the 26 prospective cells (*t* test; $p < 0.0001$) and approached significance for the entire population of 87 cells (*t* test; $p = 0.09$).

To examine the time course of prospective processing in more detail, we defined a prospective index (PI) as the NDN – NDA values. The greater the value of the PI, the more similar was the activity of the cell when the same targets were expected relative to trials in which different targets were expected: that is, the more strongly prospective was the activity. Figure 9 shows the average value of the PI over time for the entire population of 87 cells. Again, prospective activity dominated near the end of the delay. The PI values were significantly different from zero for the last three time bins of the delay (*t* test; $p < 0.01$). This indicates that, as the end of the delay and presentation of the expected targets neared, there was an increase in prospective activity.

←

triangles, respectively. The *black symbols* represent cells classified as prospective using an independent method (two-way ANOVA; see Results). *C*, Vectors connecting change in corresponding NDA and NDN values from the sample to the delay interval, i.e., the change in the data point position for each cell from *A* to *B*. This is shown for the cells classified as prospective during the delay period (*gray arrows*). The *black arrow* represents the average vector. The inset shows the distribution of $\delta\text{NDN} - \delta\text{NDA}$ values (see Results) for the population of 87 cells with selective activity in both intervals (*gray bars*) and for the cells classified as prospective in the delay (*black bars*).




Figure 9. PI as a function of time. Each bar shows the value of the PI calculated over a 400 msec bin. The PI value for each cell was calculated by first subtracting the NDA from the NDN value at each time bin. These values were then averaged across the 87 cells with selective activity in both the sample and delay epochs. High PI values indicate more similarity in activity when the same target object is expected relative to trials when different targets were expected. That is, the higher the PI value, the more prospective the activity. The gray shaded area corresponds to the period of sample object presentation. The asterisks indicate a significant difference between NDN and NDA values, evaluated at $p < 0.01$ using a t test.

Match effects

Previous studies have shown that, during DMS tasks, many neurons in the PF cortex and in extrastriate visual cortex respond differently to a test stimulus depending on whether or not it matches the sample (Gross et al., 1979; Miller et al., 1991, 1993, 1996; Riches et al., 1991; Lueschow et al., 1994; Miller and Desimone, 1994; Constantinidis and Steinmetz, 1996). Some cells show stronger responses to matches relative to nonmatches (match enhancement), and others show the opposite effect (match suppression) (Miller and Desimone, 1994; Miller et al., 1996). In this study, we found that these effects also occurred for symbolic matches (targets in the DPA task). For each stimulus and each task, we compared activity when that stimulus appeared as a target (“match”) and as a nontarget (“nonmatch”). For each stimulus and each neuron, we compared activity to the stimulus when it was a match with activity when it was a nonmatch using t tests. During the DMS task, 137 of the comparisons (of 543 possible comparisons) yielded a significant difference in activity to matches and nonmatches (t test; $p < 0.05$). Of these, 85 showed match suppression, and 52 stimuli showed match enhancement. On the DPA task, a similar number of comparisons (145) yielded significant differences in activity to matches and nonmatches. For this task, however, there was a significantly greater incidence of match enhancement (84 comparisons) than match suppression (61 comparisons; χ^2 ; $p < 0.01$). Bar graphs showing average firing rates to matching and nonmatching stimuli are shown in Figure 10. The incidence of match–nonmatch effects was similar in both monkeys.

Match enhancement was stronger on DPA trials than DMS trials. On DPA trials, there was an average 30% increase in

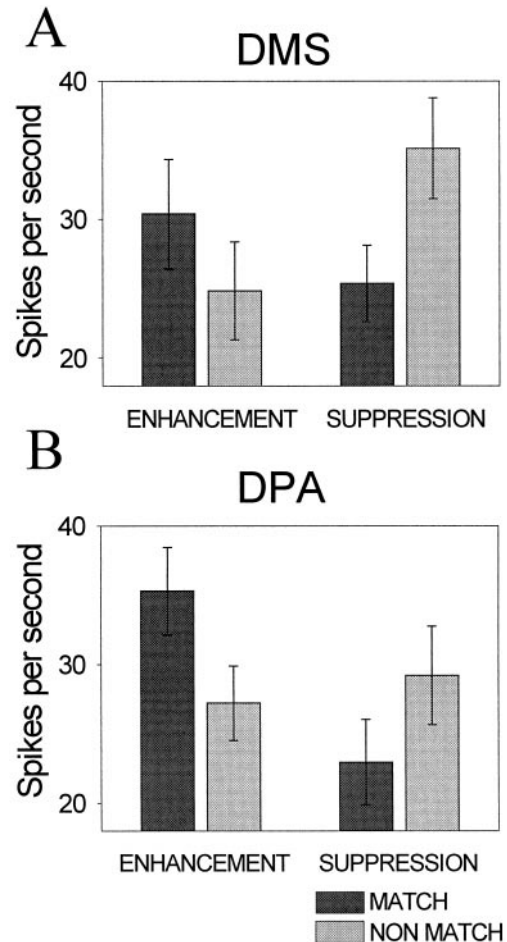


Figure 10. Match–nonmatch effects in the DMS (A) and DPA (B) tasks. Each pair of bars shows the average firing rate to matching (targets) and nonmatching (nontargets) test objects for stimuli that elicited significant enhancement or suppression of responses to matches relative to nonmatches. The error bars show the SD.

responses to matches relative to nonmatches (t test; $p = 0.05$), whereas on DMS trials, the average enhancement was weaker (a 23% increase) and did not reach significance across the population (t test; $p = 0.29$). In contrast, the opposite was true for match suppression; it was stronger on DMS trials than DPA trials. On DMS trials, there was an average 38% decrease in responses to matches relative to nonmatches (t test; $p = 0.03$), whereas on DPA trials, there was a 28% decrease in responses to matches (t test; $p = 0.18$). Thus, in terms of both incidence and strength of effects, match enhancement predominated on DPA trials, and match suppression predominated on DMS trials. A given neuron, however, did not necessarily show consistent effects across the DMS and DPA tasks. There were only 25 incidents (of a possible 137 comparisons) in which a given stimulus elicited significant match effects from a given neuron during both the DPA and DMS task. For only approximately half (13 of 25) were the effects consistent across the tasks (both suppression or both enhancement). Thus, the DPA and DMS tasks seem to elicit suppression and enhancement effects from somewhat different populations of neurons.

DISCUSSION

Many neurophysiological studies of visual memory have used standard, or identity, matching-to-sample tasks in which monkeys

are cued with a sample object and then, after a brief delay, must select that object from a display. It is widely assumed that behavior is guided by a matching process in which sensory inputs are compared with the maintained memory of the sample. In the real world, however, sought-after objects are rarely available for inspection shortly before we search for them. When we look for our missing keys, for example, we must first recall what they look like. Then, incoming visual inputs can be compared against this recalled image. The ability to recall stored information in anticipation of its use is referred to as prospective memory.

The results of this study show for the first time that the activity of PF neurons can reflect prospective coding for expected objects. During the DPA task, average PF activity first reflected the sample objects. Then, beginning before test object presentation, this activity began to reflect the target object associated with the sample. Behavioral results were also consistent with prospective coding. The monkeys' reaction times to a given target object were similar regardless of whether the monkey was cued with the object itself (DMS task) or its paired associate (DPA task). This suggests that, in both cases, the monkey had recalled the target from long-term memory before the test object appeared. The monkeys' pattern of errors was based on the similarity of targets rather than the similarity of the samples. This also suggests that the monkeys' behavior was guided by a prospective code of the target rather than a sensory-related code of the sample. Other studies (Gaffan, 1977; Roitblat, 1982) have used the confusion matrix design used in this study and found a similar pattern of errors. In principle, interpretations involving a confusion matrix depend on increasing the delay length so that errors can be attributable to increased mnemonic demands (Gaffan, 1977; Roitblat, 1982) rather than different behavioral demands between test stimulus and sample stimulus presentation. Hence, some caution is required in interpreting these errors in our study. However, the pattern of errors we observed is likely to reflect prospective coding given their similarity to those of previous studies that have used increasing delays and also given our other behavioral and neural results. Finally, we also found evidence of sensory-related activity that reflected the samples seen at the start of the trial. This has been found in previous studies (Fuster et al., 1982; Funahashi et al., 1989; di Pellegrino and Wise, 1991; Wilson et al., 1993; Miller et al., 1996; Rao et al., 1997; Asaad et al., 1998; Rainer et al., 1998a,b).

Further support for prospective coding comes from the observation that match enhancement and suppression was apparent to both actual matches (DMS task) and symbolic matches (DPA task), suggesting that a similar comparison process was used to evaluate the test object in both the DMS and DPA tasks. Miller and Desimone (1994) suggested that suppression effects reflected an automatic mechanism sensitive to repetition, whereas enhancement reflected a volitional mechanism related to matching *per se*. Results from this study are consistent with this interpretation. Suppression effects were more common in the DMS task than in the DPA task, further supporting the idea that suppression is a passive process related simply to stimulus repetition. Enhancement effects, in contrast, predominated in the DPA task, which needed to be solved by an active, symbolic, matching process involving the recall of target information from long-term memory. Suppression and enhancement effects related to identity matching are well established in the PF cortex and extrastriate visual cortex (Gross et al., 1979; Miller et al., 1991, 1993, 1996; Riches et al., 1991; Lueschow et al., 1994; Miller and Desimone, 1994; Constantinidis and Steinmetz, 1996).

A possibly related form of prospective processing in the PF cortex was described by Watanabe (1996). He found that the activity of many PF neurons reflects an expected reward (e.g., raisins, cabbage, potato). He suggested that this might be because of visual, gustatory, or olfactory images of the reward. The prospective activity we observed reflected an expected object and were thus visual in nature. This suggests that the expectancy signals observed by Watanabe may be, at least in part, a result of visual prospective coding for the expected rewards. Prospective coding is likely to occur for a wide variety of stimuli, however. Activity reflecting anticipated odors has been recently reported in the orbitofrontal cortex of rats, for example (Lipton et al., 1998).

Evidence for prospective processing of objects has also been reported in the IT cortex. Miyashita and colleagues found prospective coding for objects in tasks similar to the DPA task used here (Sakai and Miyashita, 1991; Naya et al., 1996). Erickson et al. (1998) also found prospective coding for objects in the perirhinal cortex, although most of the neurons studied had properties consistent with sensory-related coding. In addition, Sakai and Miyashita (1991) reported that neurons in the IT cortex tended to have similar visual responses to pairs of objects that were associated in training ("pair-coding neurons"). They interpret this finding as a mechanism for storing the associations between the objects. Interestingly, we did not find a similar effect in PF cortex in our task. Rather, in the PF cortex, visual similarity of stimuli appeared to be the sole determinant of activity during the sample presentation. This difference may reflect an important functional dissociation between IT and PF cortex, perhaps indicating the greater role of the IT cortex in long-term storage of the associations. The PF cortex, in contrast, may be primarily involved in the recall and maintenance of memories. This role is further suggested by the observation that humans with PF damage have severe impairments in source memory, the ability to recall the circumstances in which a certain item was learned (Janowsky et al., 1989). Functional imaging studies have also implicated the human PF cortex during recall (Buckner et al., 1996).

The recent study by Hasegawa et al. (1998) is particularly relevant. They found that, during a DPA task when both sample and target objects were presented in the same visual hemifield, monkeys' performance was unaffected by section of both the posterior and anterior corpus callosum. However, when sample and target were presented in opposite hemifields, performance was intact for the posterior section but devastated after section of the anterior corpus callosum. The authors argue that this implies a role for the PF cortex in recalling the target from storage in the temporal cortex, a conclusion supported by the results of this study. Further support comes from Gutnikov and colleagues (1997), who found that severing connections between the PF and temporal cortex disrupts monkeys' ability to perform object-object association tasks.

Prospective memory processing is critical for planning complex behavior. It allows animals to anticipate important events and modify their behavior or plans accordingly. The results of this study and others indicate that the PF cortex is involved in the prospective processing of a wide variety of information, including expected visual stimuli, rewards, and actions. This property seems fitting for a region at the apex of the perception-action cycle (Fuster, 1995) and thought to be involved in planning and organizing goal-directed behavior. Our results provide further confirmation that PF cortex does not simply buffer incoming visual inputs but instead selectively processes information relevant to current behavioral demands. Previously, we have demonstrated

that PF neurons selectively represented only the information from a cluttered display that was relevant for task performance (Rainer et al., 1998b). Here, we show that PF neurons selectively process relevant information, even when it needs to be recalled from long-term memory.

REFERENCES

- Asaad WF, Rainer G, Miller EK (1998) Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21:1399–1407.
- Barbas H (1988) Anatomic organization of basoventral and mediodorsal visual recipient prefrontal regions in the rhesus monkey. *J Comp Neurol* 276:313–342.
- Buckner RL, Raichle ME, Miezin FM, Petersen SE (1996) Functional anatomic studies of memory retrieval for auditory words and visual pictures. *J Neurosci* 16:6219–6235.
- Colombo M, Graziano MS (1994) Effects of auditory and visual interference on auditory-visual delayed matching to sample in monkeys (*Macaca fascicularis*). *Behav Neurosci* 108:636–639.
- Constantinidis C, Steinmetz MA (1996) Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *J Neurophysiol* 76:1352–1355.
- di Pellegrino G, Wise SP (1991) A neurophysiological comparison of three distinct regions of the primate frontal lobe. *Brain* 114:951–978.
- Erickson CA, Changizi BK, Desimone R (1998) Experience-dependent alterations in neuronal activity in perirhinal (PRh) cortex of rhesus monkeys. *Soc Neurosci Abstr* 23:1508.
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331–349.
- Fuster JM (1995) *Memory in the cerebral cortex*. Cambridge, MA: MIT.
- Fuster JM, Bauer RH, Jervey JP (1982) Cellular discharge in the dorsolateral prefrontal cortex of the monkey in cognitive tasks. *Exp Neurol* 77:679–694.
- Gaffan D (1977) Response coding in recall of colours by monkeys. *Q J Exp Psychol* 29:597–605.
- Goldman-Rakic PS (1987) Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In: *Handbook of physiology: the nervous system* (Plum F, ed), pp 373–417. Bethesda, MD: American Physiological Society.
- Goldman-Rakic PS, Schwartz ML (1982) Interdigitation of contralateral and ipsilateral columnar projections to frontal association cortex in primates. *Science* 216:755–757.
- Gross CG, Bender DB, Gerstein GL (1979) Activity of inferior temporal neurons in behaving monkeys. *Neuropsychologia* 17:215–229.
- Gutnikov SA, Ma YY, Gaffan D (1997) Temporo-frontal disconnection impairs visual-visual paired association learning but not configural learning in *Macaca* monkeys. *Eur J Neurosci* 9:1524–1529.
- Haenny PE, Maunsell JHR, Schiller PH (1988) State dependent activity in monkey visual cortex. II. Retinal and extraretinal factors in V4. *Exp Brain Res* 69:245–259.
- Hasegawa I, Fukushima T, Ihara T, Miyashita Y (1998) Callosal window between prefrontal cortices: cognitive interaction to retrieve long-term memory. *Science* 281:814–818.
- Honig W, Thompson RKR (1982) Retrospective and prospective processing in animal working memory. In: *The psychology of learning and motivation: advances in research and theory* (Bower GH, ed), pp 239–283. New York: Academic.
- Janowsky JS, Shimamura AP, Squire LR (1989) Source memory impairment in patients with frontal lobe lesions. *Neuropsychologia* 27:1043–1056.
- Lipton PA, Alvarez P, Eichenbaum H (1998) Interactions between orbitofrontal cortex and the hippocampal system in the coding of prospective responses to olfactory and location cues. *Soc Neurosci Abstr* 24:1424.
- Lueschow A, Miller EK, Desimone R (1994) Inferior temporal mechanisms for invariant object recognition. *Cereb Cortex* 5:523–531.
- Maunsell JHR, Sclar G, Nealey TA, DePriest DD (1991) Extraretinal representations in area V4 in the macaque monkey. *Vis Neurosci* 7:561–573.
- Miller EK (1999) The Prefrontal cortex: complex neural properties for complex behavior. *Neuron* 22:15–17.
- Miller EK, Desimone R (1994) Parallel neuronal mechanisms for short-term memory. *Science* 263:520–522.
- Miller EK, Li L, Desimone R (1991) A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* 254:1377–1379.
- Miller EK, Li L, Desimone R (1993) Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13:1460–1478.
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154–5167.
- Naya Y, Sakai K, Miyashita Y (1996) Activity of primate inferotemporal neurons related to a sought target in pair-association task. *Proc Natl Acad Sci USA* 93:2664–2669.
- O Scalaidhe SP, Wilson FA, Goldman-Rakic PS (1997) Areal segregation of face-processing neurons in prefrontal cortex. *Science* 278:1135–1138.
- Passingham R (1993) *The frontal lobes and voluntary action*. Oxford: Oxford UP.
- Petrides M (1994) Frontal lobes and working memory: evidence from investigations of the effects of cortical excisions in nonhuman primates. In: *Handbook of neuropsychology*, Vol 9 (Boller F, Grafman J, eds), pp 59–82. Amsterdam: Elsevier.
- Quintana J, Fuster JM (1992) Mnemonic and predictive functions of cortical neurons in a memory task. *NeuroReport* 3:721–724.
- Rainer G, Asaad WF, Miller EK (1998a) Memory fields of neurons in the primate prefrontal cortex. *Proc Natl Acad Sci USA* 95:15008–15013.
- Rainer G, Asaad WF, Miller EK (1998b) Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393:577–579.
- Rao SC, Rainer G, Miller EK (1997) Integration of what and where in the primate prefrontal cortex. *Science* 276:821–824.
- Riches IP, Wilson FA, Brown MW (1991) The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *J Neurosci* 11:1763–1779.
- Roitblat HL (1980) Codes and coding processes in pigeon short-term memory. *Anim Learn Behav* 8:341–351.
- Roitblat HL (1982) The meaning of representation in animal memory. *Behav Brain Sci* 5:353–406.
- Sakai K, Miyashita Y (1991) Neural organization for the long-term memory of paired associates. *Nature* 354:152–155.
- Sereno AB, Maunsell JHR (1998) Shape selectivity in primate lateral intraparietal cortex. *Nature* 395:500–503.
- Ungerleider LG, Gaffan D, Pelak VS (1989) Projections from inferior temporal cortex to prefrontal cortex via the uncinate fascicle in rhesus monkeys. *Exp Brain Res* 76:473–484.
- Watanabe M (1981) Prefrontal unit activity during delayed conditional discriminations in the monkey. *Brain Res* 225:51–65.
- Watanabe M (1996) Reward expectancy in primate prefrontal neurons. *Nature* 382:629–632.
- Wilson FAW, O Scalaidhe SP, Goldman-Rakic PS (1993) Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science* 260:1955–1958.
- Zentall TR, Sherburne LM, Steirn JN (1993) Common coding and stimulus class formation in pigeons. In: *Animal cognition* (Zentall TR, ed), pp 217–236. Hillsdale, NJ: Erlbaum.