Neuron, Vol. 21, 1399–1407, December, 1998, Copyright ©1998 by Cell Press

Neural Activity in the Primate Prefrontal Cortex during Associative Learning

Wael F. Asaad, Gregor Rainer, and Earl K. Miller*
Department of Brain and Cognitive Sciences and
The Center for Learning and Memory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Summary

The prefrontal (PF) cortex has been implicated in the remarkable ability of primates to form and rearrange arbitrary associations rapidly. This ability was studied in two monkeys, using a task that required them to learn to make specific saccades in response to particular cues and then repeatedly reverse these responses. We found that the activity of individual PF neurons represented both the cues and the associated responses, perhaps providing a neural substrate for their association. Furthermore, during learning, neural activity conveyed the direction of the animals' impending responses progressively earlier within each successive trial. The final level of activity just before the response, however, was unaffected by learning. These results suggest a role for the PF cortex in learning arbitrary cue–response associations, an ability critical for complex behavior.

Introduction

Many of our complex, learned behaviors depend on arbitrary stimulus–response associations. We learn that "green" means "go" and "red" means "stop," for example. In the laboratory, the ability to act voluntarily according to learned rules is studied using conditional visuomotor tasks (Petrides, 1985a, 1985b, 1986, 1990; Passingham, 1993). In such tasks, only a small number of responses can be performed. On any particular trial, the correct response depends on the identity of the cue used on that trial. One cue may require, for example, an eye or arm movement in a certain direction, whereas another cue would require one in another direction. No single cue or response invariably leads to reward. Instead, subjects must learn to map the cues to their associated actions.

Conditional visuomotor learning engages a wide variety of brain systems; visual stimuli must be identified, motor commands issued, and associations formed. Premotor areas involved in generating voluntary actions are important; damage severely compromises conditional visuomotor learning (Petrides, 1982; Halsband and Passingham, 1985). In fact, neurons in premotor cortex (area 6 and the supplementary eye fields [SEF]) show an evolution of activity that mirrors monkeys' acquisition of cue–response associations (Mitz et al., 1991; Chen and Wise, 1995a, 1995b, 1996). Medial temporal structures critical for long-term memories are also important: damage to the hippocampus and/or subjacent cortex (Murray and Wise, 1996) or to its principal output pathway, the fornix (Rupniak and Gaffan, 1987), impairs conditional visuomotor learning.

This type of learning also requires mechanisms that select and monitor behavioral events. Animals need to choose an action in response to each cue and keep track of these choices so that feedback about the consequences (e.g., reward) can be used to modify behavior. Here, the prefrontal (PF) cortex is likely to be involved. It has long been thought to play a role in the top-down control of complex behavior and is interconnected with virtually all sensory systems and with cortical and subcortical structures important for generating voluntary behavior (Schwartz and Goldman-Rakic, 1984; Pandya and Barnes, 1987; Pandya and Yeterian, 1990; Barbas and Pandya, 1991). There is some evidence that the PF cortex is involved in conditional visuomotor learning. Lesions of tissue around the arcuate sulcus cortex impair conditional visuomotor behavior, although damage included both PF area 8 and premotor area 6 (Petrides, 1982). Damage to the connections between the PF cortex and the inferior temporal (IT) cortex (a region important for object recognition) produces impairments on a variety of conditional tasks, including conditional visuomotor tasks (Gaffan and Harrison, 1988; Eacott and Gaffan, 1992; Parker and Gaffan, 1998). Neurons in the PF cortex have properties consistent with conditional visuomotor behavior. During such tasks, the activity of many PF neurons codes both cues and the behavioral responses instructed by them (Watanabe, 1986; Sakagami and Niki, 1994a, 1994b; Hasegawa et al., 1998).

To explore the role of the PF cortex in arbitrary cue–response learning, we studied the activity of lateral PF neurons during performance of a conditional visuomotor task. This task required the animals to associate a foveally presented cue object (500 ms) with a saccadic eye movement, either to the left or to the right (Figure 1). Two novel cue objects were used each day (i.e., for each recording session). They learned which response was required for each cue by trial and error. To separate neural events related to the cue and response, a delay (1000 ms) was imposed between the disappearance of the cue and the "go" signal. Rather than have monkeys learn a single cue–response association, we required them to learn to associate, on alternate blocks of 30 or more trials, each of two cue objects with each of two saccades. In other words, after having learned "object A, go right" and "object B, go left," the associations were reversed such that they now needed to learn "A, go left" and "B, go right." By pairing each cue with each possible response, we were able to disambiguate the effects of the cue and the behavioral response on neural activity and thus determine how PF neurons help represent these attributes.

*To whom correspondence should be addressed (e-mail: ekm@ai.mit.edu).

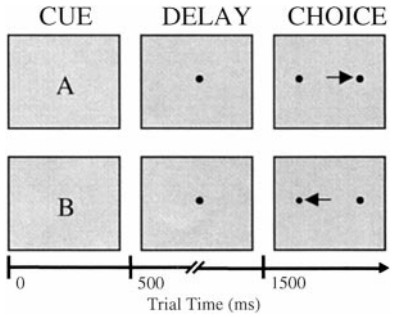

Figure 1. A Schematic Diagram of the Behavioral Task

Results

Behavior

Both monkeys learned each object–saccade direction pairing quickly. Figure 2a shows their average performance on each block of trials. Early in the sessions (blocks 1–3), the monkeys needed only 10–15 trials to reach seven out of seven correct (the behavioral criterion). Later in the session (blocks 4–6), their performance worsened somewhat (perhaps due to proactive interference, fatigue, and/or satiation), requiring 25–30 trials to reach criterion.

When the object–saccade pairings were reversed, performance was disrupted but quickly recovered. This is illustrated in Figure 2b. Trial zero indicates the first trial after the object–direction pairings reversed. At this point, the object that had required a rightward response for reward now required a leftward response and vice versa. Because the reversal was not explicitly signaled to the animal, they almost invariably made an incorrect choice on the first trial when the associations were reversed. Note that their performance immediately after reversals (average ~15% correct) was the converse of their performance just before reversals (~85% correct). Within about 6–12 trials, their overall performance reached 60%–70% correct (chance = 50%) and then gradually improved to above 80% correct by the end of the block.

Learning was also reflected in a change in the reaction times of the saccadic eye movements at the choice. Just before the reversals, reaction times averaged ~175 ms. Just after the reversals, saccadic reaction times increased to ~200 ms and then decreased in parallel with the change in animals' error rates. Indeed, the reaction times for the ten trials preceding a reversal were significantly shorter than those for the ten trials just after a reversal ($p < 0.001$, t test).

Neuronal Properties

General Properties

We recorded a total of 254 neurons from the left lateral PF cortices of two monkeys (146 from one and 108 from the other). Most of the neurons (249 of 254, or 98%) were responsive in at least one task period. This was determined by comparing each of the three task intervals (cue, delay, and presaccade; see Experimental Procedures) to baseline, or spontaneous, activity using paired t tests (evaluated at $p < 0.01$).

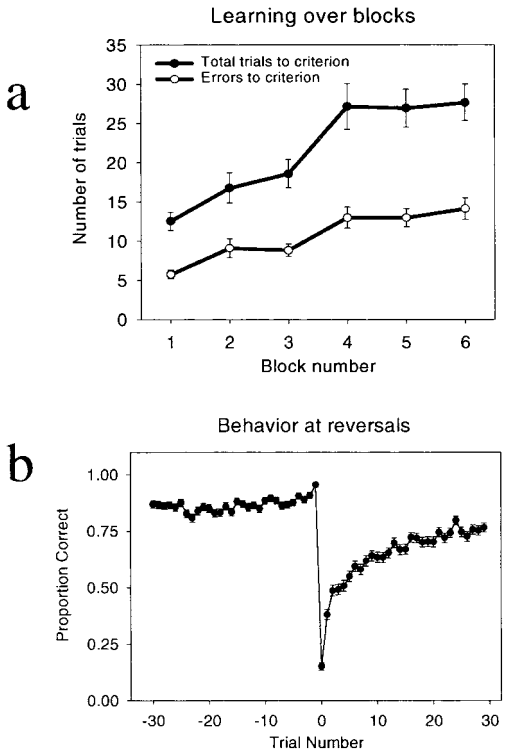

Figure 2. Cue–Saccade Learning over Blocks and at Reversals

(A) The number of trials required by the monkeys to learn the cue–saccade pairings. This figure plots the total number of trials (closed circles) or the total number of incorrect choices (open circles) taken by the two monkeys to reach the criterion of seven out of seven correct choices. Each reversal constituted a new block (x axis). The rate of learning by both measures decreased with additional reversals.

(B) Behavioral performance at the point of reversals. This panel plots percent correct performance (calculated using a moving window of three trials) as a function of trial number relative to a reversal (at trial zero). Performance preceding a reversal was generally around 85% correct, demonstrating that the animals learned the cue–saccade pairings to a high level by the end of each block. Performance does not return to prereversal levels in this graph because 30 trials represents a minimum number of trials completed. Blocks often lasted longer, depending on the animal's rate of learning in each particular block.

Learned-State Properties

First, we examined neuronal properties after the object and direction pairings were relatively well learned. For these analyses, we included only data extracted from "well-learned trials." Specifically, these were correct trials from groups of ten trials of at least 80% correct performance. We thus hoped to exclude, for the moment, any learning-related changes in activity in order to focus first on the steady-state properties, that is, the stable representation of cue–saccade associations. Table 1 summarizes the incidence of object and/or direction selectivity for each task interval based on ANOVAs (see Experimental Procedures), evaluated at $p < 0.01$. The majority of recorded cells (202 of 254, or 80%) displayed activity selective for the cue object, the saccade direction, or both, in at least one of the three task epochs.

Almost half of these neurons (96 of 202, or 48%)

Table 1. Number of Neurons with Task-Related Activity in the Different Epochs (See Experimental Procedures) of Each Trial

254 Total Cells	CUE	DELAY	PRESACCADIC	In Any Epoch
Number of Responsive Cells	212	207	203	249
Number of Selective Cells	131	155	108	202
	62% of 212	75% of 207	53% of 249	81% of 249
Object Selective	59	49	23	96
	45% of 131	32% of 155	21% of 108	48% of 202
Direction Selective	18	33	37	59
	14% of 131	21% of 155	34% of 108	29% of 202
Object and Direction (linear)	9	18	12	32
	7% of 131	12% of 155	11% of 108	16% of 202
Object and Direction (nonlinear)	45	55	36	88
	34% of 131	35% of 155	33% of 108	44% of 202

The number of selective cells is noted as a percentage of the total number of responsive cells in each epoch, and the number of cells displaying a particular type of selectivity is noted as a percentage of the number of selective cells.

showed activity selective for the cue object, but not the saccade direction, in at least one epoch. Examples of such cells are depicted in Figure 3a. They showed object-selective activity starting shortly after cue onset that was maintained throughout the trial. On the ANOVA, these cells showed a significant effect of OBJECT, but not of DIRECTION, and no significant interaction between the factors (evaluated at $p < 0.01$). Other cells (59 of 202, or 29%) were selective for the saccade direction, but not for the cue object in at least one epoch. They showed a significant effect of DIRECTION, but not OBJECT, and no interaction between the factors. An example is depicted in Figure 3b. It was relatively unselective during sample presentation. Then, during the delay, its activity was highly dependent on the direction of the forthcoming saccade but unaffected by the cue object.

Many cells (120 of 202, or 60%) showed activity that depended on both the cue object and saccade direction in at least one epoch. They showed a significant effect of OBJECT and DIRECTION and/or a significant interaction between the factors ($p < 0.01$). For some (32 of 202, or 16%), activity seemed to reflect a straightforward, linear addition of these attributes. For example, the cell illustrated in Figure 3c preferred rightward saccades over leftward saccades. Superimposed on this direction selectivity was a preference for object B over object A. Such cells showed a significant effect of OBJECT and DIRECTION with no interaction between the factors.

For the majority of object- and direction-selective cells (88 of 202, or 44%), the cue and saccade direction influenced activity in a nonlinear fashion. That is, their responses to specific object–saccade pairings could not be explained by simply summing their responses to the individual elements. For example, the cell illustrated in Figure 3d was nonselective during cue presentation. Then, in the delay, it showed the highest activity on trials in which monkeys saw object A and made a leftward saccade. Activity to all other combinations of cues and saccades elicited equally lower levels of activity. In other words, direction selectivity was apparent for only one of the two cue objects. These "nonlinear" cells showed

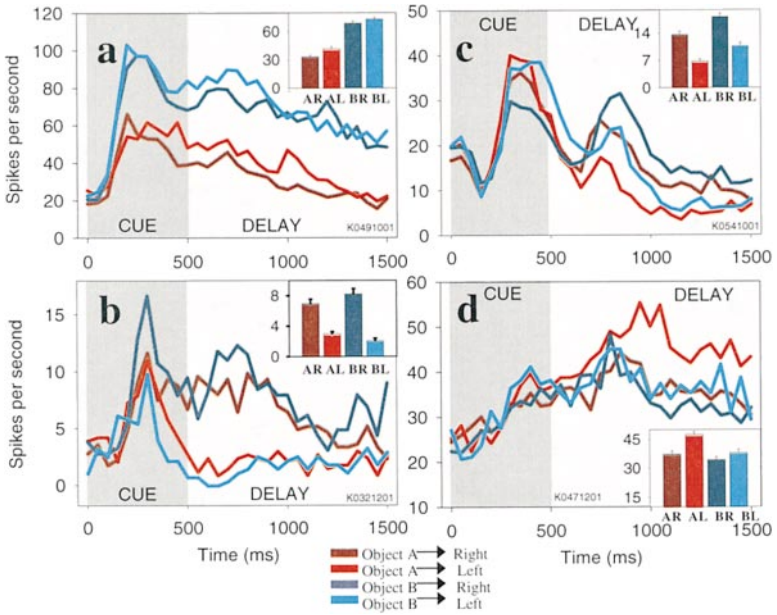

Figure 3. Examples of Neural Responses after Learning

The figure shows the responses of four neurons to each of the four possible cue–saccade combinations. The shaded area represents the time of cue presentation. The bar plot in the inset shows the mean level of delay activity with standard errors for each cue–saccade pairing (AR: object A, right; AL: object A, left; BR: object B, right; BL: object B, left). The colors used in the inset match those of the histograms and are keyed by the legend. Shown are an object-selective cell (A), a direction-selective cell (B), a linear object-and-direction-selective cell (C), and a nonlinear object-and-direction-selective cell (D). Bin width, 50 ms.

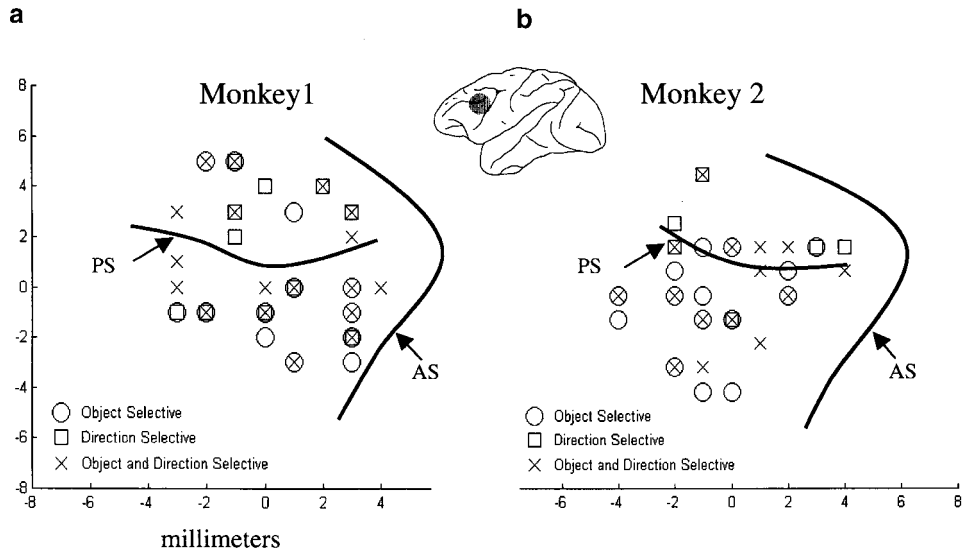

Figure 4. Recording Sites

Penetration sites where selective cells were found are plotted for each monkey. Different markers are used to represent locations where cells with the indicated types of selectivity were recorded. PS, principal sulcus; AS, arcuate sulcus. The inset shows, on a "generic" rhesus brain, the general area of recording.

a significant interaction between OBJECT and DIRECTION factors ($p < 0.01$).

Anatomical Locations of Cells

The recording sites for each of the two monkeys are shown in Figures 4a and 4b. Each symbol indicates a site where at least one neuron of the indicated selectivity was found. While we observed a slight bias for object-selective cells to appear ventrally in one animal, there was no strict segregation of cells with different types of responses. Cells with combined selectivity for cue identity and response direction were interspersed throughout the recording area.

Neural Activity and Learning

Direction selectivity was evident even early in each block of trials, before monkeys had learned the associations. This is depicted in Figure 5. This graph plots the average neural activity during the last 250 ms of the delay (which is similar to one of the epochs used by Chen and Wise [1995a, 1995b] in their studies of SEF neurons). For this analysis, we included the 64 cells whose delay activity showed a main effect of DIRECTION on the ANOVA ($p < 0.01$). The black circles show the average activity when a given object first cued the preferred direction and then reversed (at trial zero) to cue the nonpreferred direction. The white circles show the opposite transition. Note that the level of activity on even the first correctly executed trial after a reversal is indistinguishable from activity 30 correct trials later. That is, direction-selective activity abruptly "switched" to reflect the saccade the animal was about to perform. This was also evident in an examination of the animal's error trials; also plotted on the figure is the average activity during the first ten error trials in which the animals chose the wrong saccade (red circles). Note that this activity is similar to that before the reversal, when the saccade was correct. Thus, this activity was not reflecting the "output" of a learned association. Rather, it seemed simply to reflect the direction of whichever response was impending, regardless of whether or not the animal had learned it was "correct."

Learning was, however, reflected in an earlier appearance of direction-selective activity. After learning, direction selectivity appeared earlier in the trial than it had before learning. For these analyses, we aligned activity on the execution of the saccade to compensate for the change in saccade latency with learning. Collapsing across different blocks of trials, we computed a direction selectivity index (see Experimental Procedures) for each neuron for the first 30 correct trials of the block. The average values of the selectivity index are plotted in Figure 6a. Each box represents 25 ms of one trial. The bottom row, therefore, depicts the average selectivity index of 64 cells for the first correctly executed trial within each block. Subsequent trials are plotted upward. Robust direction selectivity (the yellow-white colors) tends to appear progressively earlier in each successive trial. In fact, the largest change seems to occur around trials 5–10, which is when the animals first showed evidence of learning the object–saccade associations (see Behavior).

We defined the time of appearance of directional selectivity as the earliest bin in which half of the maximal selectivity was reached. These times are plotted in Figure 6b. There was an inverse relationship between trial number and the time of appearance of directional selectivity, which we approximated with a simple sigmoid function ($r^2 = 0.97$, $p < 0.0001$). In the first five correct trials after a reversal, direction selectivity appears ~700 ms into the trial (almost 900 ms before the response). By the time the animals' performance begins to asymptote at trial 15, direction selectivity appears much earlier, at ~250 ms into the trial (1350 ms before the response).

Neural Correlates of Over-Learning

In addition to introducing two novel cue objects in each session, we used two highly familiar cue objects. These

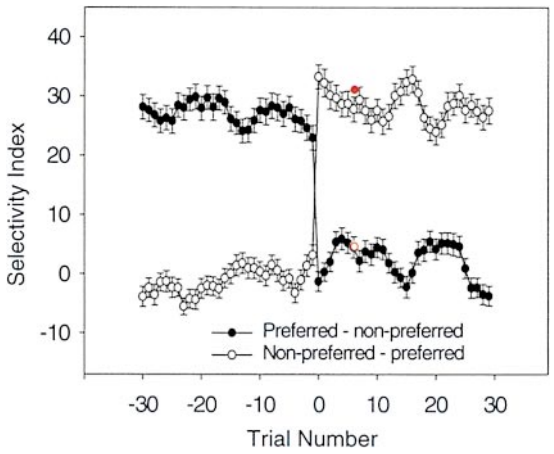

Figure 5. Directionally Selective Activity around the Time of Reversals

This figure plots the average selectivity index (see Experimental Procedures) for 64 direction-selective cells. The index is expressed as a percent change in activity to the preferred over the nonpreferred direction. The figure plots the average direction selectivity during the last 250 ms of the delay before and after reversals. The closed circles follow the responses to the object that, having indicated a saccade in the neurons' preferred directions, begins at trial zero to indicate a saccade in their nonpreferred directions. The open circles show the opposite, namely the activity following the object that had previously cued a saccade in their nonpreferred directions reversing, such that it subsequently required an eye movement in their preferred directions. The bars represent standard errors. Neural activity on even the first correctly executed trial after a reversal (trial zero) was determined entirely by the cells' directional preferences and not by the animals' level of learning. Activity to the incorrect choices (red circles) also reflected simply the direction of the intended response. The closed red circle represents activity preceding incorrect selection of the neurons' preferred directions, whereas the open red circle shows activity preceding incorrect selection of the nonpreferred direction.

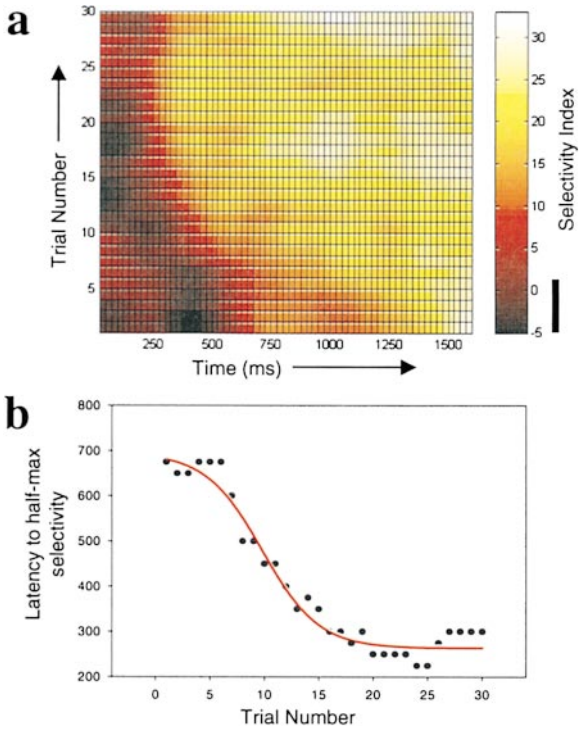

Figure 6. Direction-Selective Activity Appeared Earlier within Each Trial as the Animals Learned

(A) Change in latency of direction-selective activity with learning. The selectivity index for the same 64 cells as in Figure 5 is shown in this surface plot. Directional selectivity appeared earlier (further to the left) with increasing trial number (upward). Each individual box represents the average selectivity index for 25 ms of one trial. The data in this figure are smoothed with a sliding kernel of 5 × 5 bins. The trials are aligned on the initiation of the saccadic eye movement and include the cue and delay intervals. The black bar in the lower right corner illustrates the average standard error of the mean for all the data points. The data shown in the rightmost ten bins of each trial correspond to those used in Figure 5.

(B) The time at which half of the maximal selectivity was reached within each trial is plotted along with the fitting sigmoid function:

$$y = y_0 - \frac{a}{1 + e^{-\frac{(x - x_0)}{b}}}$$

where $y_0 = 698.6$, $x_0 = 9.8$, $a = 430.5$, and $b = 2$. The y axis represents the time of appearance of selectivity, while trial number (after a reversal) is represented along the x axis.

familiar objects were used throughout training and for every recording session. Each familiar object was always associated with the same saccade direction (i.e., they were never reversed). Figure 7 shows the average activity across the entire population of 254 PF neurons studied. This shows that the novel, reversing objects elicited more activity than the familiar, nonreversing objects. In fact, 215 of 254 neurons (85%) showed a significant difference in activity between the novel and familiar objects in at least one epoch ($p < 0.01$, t test, activity to familiar objects versus activity to novel objects).

We sought to determine if this difference in activity between novel and familiar objects could be developed "online," i.e., in the course of one recording session. We recorded the activity of an additional 30 neurons (beyond the 254 neurons reported above) from one monkey, during the presentation of two novel objects for which the initial object–direction associations were maintained throughout the recording session (1924 trials on average). We then compared the activity of these neurons on the first 50–100 correct trials to that from the last 50–100 correct trials of the session. No significant difference was observed ($p > 0.05$, t test on activity from early trials versus activity from late trials), suggesting that the difference reported above for the main population of neurons may depend upon repeated exposure and may require days or weeks to develop.

Discussion

To perform this task, the monkeys needed to identify cue stimuli, choose responses, and form associations between them. Because the presentation of the visual cue and the animal's response were separated in time, monkeys needed to bridge the gap. The role of PF neurons in sustaining representations of objects and/or cued locations is well established (Fuster, 1973; Fuster and Jervey, 1981; Kojima and Goldman-Rakic, 1982; Funahashi et al., 1989; di Pellegrino and Wise, 1991; Miller

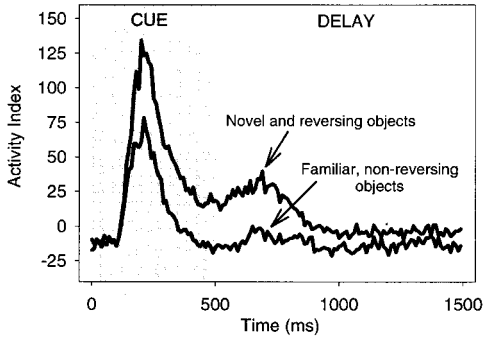

Figure 7. Neurons Were More Active to Novel Objects that Required Learning than to Familiar Ojects with Well-Learned Saccade Associations

This figure plots the normalized activity (see Experimental Procedures) of all 254 recorded neurons on trials in which novel, reversing objects were used (top line) versus trials in which highly familiar objects that never reversed directional association were used (bottom line). The shaded area represents the time of cue presentation. Bin width, 10 ms. Most individual cells (85%) discriminated between these two groups (see Results).

et al., 1996; Rao et al., 1997). It should be noted, however, that the ventral PF cortex at least is not critical for object memory per se; monkeys with ventral PF lesions can remember an object over a brief delay (Rushworth et al., 1997). This study shows that, for many PF neurons, during conditional visuomotor behavior, activity that conveyed the identity of the cue combined with activity related to the forthcoming saccade, a property also reported by Hasegawa et al. (1998). Activity reflecting conjunctions of cues and behavioral responses have also been reported during conditional visuomotor tasks that required monkeys to associate a visual and/or auditory cue with an immediate or delayed response (Watanabe, 1986; Sakagami and Niki, 1994a, 1994b). This ability to represent conjunctions of cues and responses seems fitting for a region at the apex of the "perception–action cycle" (Fuster, 1995). The direction-selective activity observed in our study could have been due to a "premotor" signal for the impending eye movement and/or a shift of visual attention preceding the eye movement. Both sensory and motor-related spatial signals have been reported in the PF cortex (Niki and Watanabe, 1976; Quintana and Fuster, 1992; Boussaoud and Wise, 1993a, 1993b; Funahashi et al., 1993). In either case, the spatial signals observed here could be used to plan the animal's behavioral response. Indeed, they reflect the forthcoming saccade prior to learning, as do many neurons in the premotor cortex (Mitz et al., 1991; Chen and Wise, 1995a), and do so irrespective of whether the behavioral response is correct or not, a phenomenon also observed by Niki and Watanabe (1976).

This ability to integrate information about cues and responses is critical for associative learning. Visual information about cue identity and the visuospatial information needed to direct actions are processed in largely separate systems in posterior neocortex (although there does seem to be "cross-talk" between these systems; e.g., see Sereno and Maunsell, 1998). Establishing cue–response associations, therefore, requires neural substrates where such information can converge. The present results show that this can be reflected at the single neuron level in the PF cortex. These results are consistent with earlier studies indicating that the PF cortex may play a role in integrating information about visual identity with spatial information (Fuster et al., 1982; Rao et al., 1997; Rainer et al., 1998a, 1998b). Taken together, the present results along with prior studies indicate that PF neurons can integrate diverse, behaviorally relevant information.

Many PF neurons showed nonlinear tuning for object–saccade conjunctions. That is, their activity for a given cue–response conjunction could not be predicted by linearly combining their activity for other combinations of the same elements. This property is also apparent in other studies of PF neurons during conditional visuomotor tasks (e.g., Watanabe, 1992). PF activity, therefore, seems ideally suited for representing specific cue–response conjunctions. Nonlinear tuning may be common in other areas involved in conditional learning. Hippocampal neurons, for instance, often show nonlinear responses to combinations of spatial and nonspatial, behaviorally relevant information (Young et al., 1994; Deadwyler et al., 1996). It is important to note that our distinction between "linear" and "nonlinear" is not intended to convey a categorical distinction between two discrete classes of cells. While cells were classified as "linear" or "nonlinear" according to a statistical criterion (a significant interaction on the ANOVA), they actually varied along a continuum. In fact, the observed incidence of nonlinear cells is probably a lower bound. It is possible, indeed likely, that had we used more than two objects and two saccade directions, more cells would have shown nonlinear effects.

Learning was reflected in PF activity by the earlier appearance of directionally selective activity within each trial. A similar phenomenon has been observed in the SEF. Neural activity after the monkey's response, but before reward, appeared progressively earlier as the monkeys acquired the task (see the "prereward" interval in Figures 5b, 8b, and 9a of Chen and Wise, 1995a). Also, many SEF neurons show an evolution of direction-selective activity as monkeys learn object–saccade associations (Chen and Wise, 1995a, 1995b, 1996). These findings suggest that the PF and premotor cortices may share similar learning-related mechanisms and may contribute to a distributed system for associative learning. Any more specific comparisons regarding these areas will require further study, however. The task used by Chen and Wise (four objects and four directions, typically no reversals) is different from our task (two objects and two directions, frequent reversals). Such comparisons are best made using a single task within the same monkeys (Chen and Wise, 1996).

Recent evidence suggests that dopaminergic (DA) signals arising from the midbrain, a region that provides the PF cortex with its major DA input, have properties that may be related to the progressively earlier appearance of direction selectivity. DA has the ability to modulate mnemonic activity in PF cortex (Williams and Goldman-Rakic, 1995). Additionally, midbrain DA neurons are initially activated by a reward. Then, as the animals

learn the predictive value of a conditioned stimulus, their responses shift earlier in time toward the presentation of that stimulus (Schultz et al., 1993). In our study, as the monkeys learned the behavioral requirements of the cues, these stimuli became increasingly reliable predictors of reward. Thus, according to the model proposed by Schultz et al. (1997), DA-related activity would be expected to shift in time toward the cue, possibly gating direction-selective delay activity at an earlier point in each succeeding trial as the animals learned. These may be the mechanisms that underlie the changes we observed in the latencies of direction-selective response-related activity.

Learning was also reflected in that most PF cells were activated more strongly by novel, reversing cues than by familiar, nonreversing cues. SEF neurons show similar effects (Chen and Wise, 1995a). Whether the difference in activity to novel and familiar stimuli is due to the novelty of the cues and associated responses, or the difficulty of the reversals, or some combination, is not yet clear. Responses in higher visual cortex have likewise been reported to decrease with increasing stimulus familiarity (Miller et al., 1991; Riches et al., 1991; Li et al., 1993) and increase with task difficulty or increased attention (Spitzer et al., 1988; Spitzer and Richmond, 1991). Human functional imaging studies have also reported that the level of PF activation can be modulated by over-training (Raichle et al., 1994; Berns et al., 1998). In humans, however, such changes occur sufficiently quickly that they can be viewed over the course of a single session. For the 30 cells we specifically tested for this effect, activity in response to novel, nonreversing cues did not decrease sufficiently in the course of a single recording session to be observed. If indeed these phenomena are related, the discrepancy in time course may be due to species or task differences.

In the early stages of learning a conditional association task, a response must be selected without any knowledge of its correctness or, in the case of reversals, despite previous knowledge to the contrary. Presumably, direction-selective PF activity, which was evident even before learning, reflected response selection. PF neurons that combine this activity with information about the preceding cue could play a role in guiding associative learning. Their activity could act in concert with reward signals that indicate a successful cue–response pairing to shape the neural circuitry that underlies the long-term associations between the cues and responses. Such an interpretation would be consistent with the large body of evidence attributing to the PF cortex a role particularly in novel or changing circumstances.

Experimental Procedures

Subjects

The subjects were two rhesus monkeys, *Macacca mulatta*, weighing 10 and 6 kg. Using previously described methods (Miller et al., 1993), monkeys were implanted with a scleral search coil (Robinson, 1963; Judge et al., 1980) to monitor eye movements, a head bolt to immobilize the head during recording, and recording chambers. Penetration sites were determined using magnetic resonance imaging (MRI). The recording chambers were positioned stereotaxically over the lateral PF cortex such that the principal sulcus and surrounding cortex, especially the ventrolateral PF cortex, were readily accessible. All surgeries were performed under aseptic conditions while the animals were anesthetized with isoflurane. The animals received postoperative antibiotics and analgesics and were always handled in accord with NIH guidelines and the recommendations of the MIT Animal Care and Use Committee.

Recording Technique

Monkeys were seated in primate chairs within sound-attenuating enclosures (Crist Instruments, Damascus, MD). Their heads were restrained and a juice-spout placed at their mouths for automated reward delivery. Recordings were made using arrays of four to eight dura-puncturing, tungsten microelectrodes (FHC Instruments, Bowdoin, ME). Electrodes were mounted on custom-made, independently adjustable miniature microdrives. These were introduced into the brain using a grid system (Crist Instruments) with 1 mm spacing between adjacent locations. 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 data collection began. This procedure was used to ensure an unbiased estimate of PF activity. In any given session, we were able to simultaneously record the activity of up to 18 individual neurons. Neural waveforms were digitized and stored for offline sorting into single-neuron records (DataWave, Longmont, CO).

Behavioral Task

The task required the animals to associate a foveally presented object with a saccade either to the right or left (Figure 1). The stimuli were small, complex objects about $2^\circ \times 2^\circ$ in size. For each recording session, two novel objects, never before seen by the animal, were selected from a large pool of objects. They were arbitrarily designated object "A" and object "B." The objects were presented on a computer screen positioned directly in front of the animal. Complex objects were used because they have been shown to elicit robust activity from lateral PF neurons (Miller et al., 1996). We made no attempt to determine which features of particular objects were responsible for the cells' responses; for this study, it was necessary only that different objects elicited selective activity from a number of PF neurons.

Each trial began with presentation of a small dot (a fixation point) that the monkey was required to foveate. It maintained gaze on the fixation point until the behavioral choice at the end of the trial. After one second of fixation, one of the objects was presented as a cue at the fovea for 500 ms, followed by a delay of 1000 ms. At the end of the delay, two dots were simultaneously presented as the central fixation dot was extinguished. One dot was 4° to the right and the other 4° to the left of fixation. Monkeys were required to saccade to either the right or left dot depending on the identity of the cue object. The correct response was immediately rewarded with a few drops of apple juice. An incorrect choice resulted in a 3 s "time out," followed by termination of the trial without delivery of reward. Incorrect trials were not immediately repeated, as is sometimes done in such tasks to aid learning.

The criterion for having learned a cue–direction pairing was seven consecutive correct choices for each object (the probability of achieving seven consecutive correct choices by chance is <0.01). Then, the monkeys' behavior was allowed to asymptote by allowing them to continue performing the learned cue–saccade pairings for an additional 20–50 trials. This allowed the animals to obtain enough apple juice to keep them motivated and also provided us with steady-state neural data. After this, the cue–response pairings were reversed; the cue that had required a rightward saccade now required a leftward saccade and vice versa. The monkeys were not explicitly cued that a reversal had occurred. Instead, they had to infer it from feedback about their performance. Again, the monkeys were required to learn the correct responses by trial and error. This learning–asymptote–reversal cycle continued for as long as the monkeys were willing to work. Each reversal was classified as a new block; the monkeys worked four to ten blocks (average = 6) in a day (i.e., for one recording session).

Data Analysis

We divided the trial into three contiguous, nonoverlapping epochs for analysis of neural activity. We defined "cue" activity as the neural

response from 100–700 ms after stimulus onset. The first 100 ms were excluded to compensate for the minimum latency of visual responses in PF cortex (100–150 ms). The length of this time window was selected to include any activity related to offset of the stimulus. The "delay" epoch began immediately after the cue period and lasted until 250 ms before the animal's response (usually, it lasted 700–750 ms). The ensuing activity within the final 250 ms before the response constituted the "presaccadic" period. Baseline (spontaneous) activity was assessed using the average activity during the 1 s period in the middle of the 3 s intertrial interval (ITI). These epochs were chosen for simplicity; the results reported here are insensitive to the exact time windows used. Except where noted (red circles in Figure 5), neural activity from only correct trials was examined.

Reversals allowed us to avoid confounding the effects of object and saccade direction on neural activity. Because, over any recording session, each object cued each saccade direction, we could dissociate the effect of each on neural activity. A two-way ANOVA was performed for each cell on activity from each trial epoch (evaluated at $p < 0.01$). One factor was the object used as a cue (OBJECT factor) and the other factor was the saccade direction associated with it (DIRECTION factor). To assess the neural effects of learning, we compared neural activity from correct trials early in each block, when the animals' overall performance was low, to correct trials later in the blocks, once the associations had become well learned. To compare direction selectivity across a population of neurons, we computed a selectivity index (Luck et al., 1997; Rainer et al., 1998b). It was computed by taking the absolute value of the difference in activity to the two saccade directions divided by their sum and then converting this value to a percent difference. Average activity across the entire population of PF neurons (in Figure 7) was normalized in a similar fashion. Neural activity is often not normally distributed. Thus, we also applied nonparametric statistics to neural activity (Mann–Whitney U test and the Kruskal–Wallis H test). These tests yielded virtually identical results to the parametric tests.

Neurophysiological experiments that compare activity across different blocks of trials must make efforts to ensure that any neural effects are not the result of artifacts of that design, such as slow-wave changes in neural activity over time. We made certain such artifacts did not influence our data in several ways. First, because animals performed four to ten blocks, all effects related to specific cue–response pairings were replicated two to five times throughout the recording session. We further controlled for any nonspecific effects by replicating all of the observed results after normalizing within-trial activity to spontaneous (baseline) activity. We employed two normalization methods. We subtracted the average baseline activity on each block from the within-trial activity to express neural responses as differences from baseline, and we divided within-trial activity by baseline activity so that within-trial activity would be expressed as a proportional change over baseline. Both techniques yielded identical results to the analyses based on raw data reported here.

Acknowledgments

The authors wish to thank Kenway Louie, Mayank Mehta, Susie Ro, Thanos Siapas, Richard Wehby, Marlene Wicherski, and Matthew Wilson for their comments on the manuscript. This work was supported by NINDS/NIH grant NS35145-02 and the Whitehall Foundation.

Received September 23, 1998; revised November 3, 1998.

References

Barbas, H., and Pandya, D. (1991). Patterns of connections of the prefrontal cortex in the rhesus monkey associated with cortical architecture. In Frontal Lobe Function and Dysfunction, H.S. Levin, H.M. Eisenberg, and A.L. Benton, eds. (New York: Oxford University Press), pp. 35–58.

Berns, G.S., Cohen, J.D., and Mintun, M.A. (1998). Brain regions responsive to novelty in the absence of awareness. Science *276*, 1272–1275.

Boussaoud, D., and Wise, S.P. (1993a). Primate frontal cortex: effects of stimulus and movement. Exp. Brain Res. *95*, 28–40.

Boussaoud, D., and Wise, S.P. (1993b). Primate frontal cortex: neuronal activity following attentional versus intentional cues. Exp. Brain Res. *95*, 15–27.

Chen, L.L., and Wise, S.P. (1995a). Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. J. Neurophysiol. *73*, 1101–1121.

Chen, L.L., and Wise, S.P. (1995b). Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. J. Neurosci. *73*, 1122–1134.

Chen, L.L., and Wise, S.P. (1996). Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations. J. Neurosci. *16*, 3067–3081.

Deadwyler, S.A., Bunn, T., and Hampson, R.E. (1996). Hippocampal ensemble activity during spatial delayed-nonmatch-to-sample performance in rats. J. Neurosci. *16*, 354–372.

di Pellegrino, G., and Wise, S.P. (1991). A neurophysiological comparison of three distinct regions of the primate frontal lobe. Brain *114*, 951–978.

Eacott, M.J., and Gaffan, D. (1992). Inferotemporal-frontal disconnection—the uncinate fascicle and visual associative learning in monkeys. Eur. J. Neurosci. *4*, 1320–1332.

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

Funahashi, S., Chafee, M.V., and Goldman-Rakic, P.S. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. Nature *365*, 753–756.

Fuster, J.M. (1973). Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. J. Neurophysiol. *36*, 61–78.

Fuster, J.M. (1995). Memory in the Cerebral Cortex (Cambridge, MA: MIT Press).

Fuster, J.M., and Jervey, J.P. (1981). Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. Science *212*, 952–955.

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

Gaffan, D., and Harrison, S. (1988). Inferotemporal-frontal disconnection and fornix transection in visuomotor conditional learning by monkeys. Behav. Brain Res. *31*, 149–163.

Halsband, U., and Passingham, R.E. (1985). Premotor cortex and the conditions for movement in monkeys. Behav. Brain Res. *18*, 269–276.

Hasegawa, R., Sawaguchi, T., and Kubota, K. (1998). Monkey prefrontal neuronal activity coding the forthcoming saccade in an oculomotor delayed matching-to-sample task. J. Neurophysiol. *79*, 322–333.

Judge, S.J., Wurtz, R.H., and Richmond, B.J. (1980). Vision during saccadic eye movements. I. Visual interactions in striate cortex. J. Neurophysiol. *43*, 1133–1155.

Kojima, S., and Goldman-Rakic, P.S. (1982). Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. Brain Res. *248*, 43–49.

Li, L., Miller, E.K., and Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. J. Neurophysiol. *69*, 1918–1929.

Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J. Neurophysiol. *77*, 24–42.

Miller, E.K., Gochin, P.M., and Gross, C.G. (1991). Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque. Vis. Neurosci. *7*, 357–362.

Miller, E.K., Li, L., and Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. J. Neurosci. *13*, 1460–1478.

Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J. Neurosci. *16*, 5154–5167.

Mitz, A.R., Godschalk, M., and Wise, S.P. (1991). Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. J. Neurosci. *11*, 1855–1872.

Murray, E.A., and Wise, S.P. (1996). Role of the hippocampus plus subjacent cortex but not amygdala in visuomotor conditional learning in rhesus monkeys. Behav. Neurosci. *110*, 1261–1270.

Niki, H., and Watanabe, M. (1976). Prefrontal unit activity and delayed response: relation to cue location versus direction of response. Brain Res. *105*, 79–88.

Pandya, D.N., and Barnes, C.L. (1987). Architecture and connections of the frontal lobe. In The Frontal Lobes Revisited, E. Perecman, ed. (New York: The IRBN Press), pp. 41–72.

Pandya, D.N., and Yeterian, E.H. (1990). Prefrontal cortex in relation to other cortical areas in rhesus monkey—architecture and connections. Prog. Brain Res. *85*, 63–94.

Parker, A., and Gaffan, D. (1998). Memory after frontal/temporal disconnection in monkeys: conditional and nonconditional tasks, unilateral and bilateral frontal lesions. Neuropsychologia *36*, 259–271.

Passingham, R. (1993). The Frontal Lobes and Voluntary Action (Oxford: Oxford University Press).

Petrides, M. (1982). Motor conditional associative-learning after selective prefrontal lesions in the monkey. Behav. Brain Res. *5*, 407–413.

Petrides, M. (1985a). Deficits in nonspatial conditional associative learning after periarcuate lesions in the monkey. Behav. Brain Res. *16*, 95–101.

Petrides, M. (1985b). Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man. Neuropsychologia *23*, 601–614.

Petrides, M. (1986). The effect of periarcuate lesions in the monkey on the performance of symmetrically and asymmetrically reinforced visual and auditory go, no-go tasks. J. Neurosci. *6*, 2054–2063.

Petrides, M. (1990). Nonspatial conditional learning impaired in patients with unilateral frontal but not unilateral temporal lobe excisions. Neuropsychologia *28*, 137–149.

Quintana, J., and Fuster, J.M. (1992). Mnemonic and predictive functions of cortical neurons in a memory task. Neuroreport *3*, 721–724.

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

Rainer, G., Asaad, W.F., and Miller, E.K. (1998a). Memory fields of neurons in the primate prefrontal cortex. Proc. Natl. Acad. Sci. USA, in press.

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

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

Riches, I.P., Wilson, F.A., and Brown, M.W. (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.

Robinson, D.A. (1963). A method of measuring eye movement using a scleral search coil in a magnetic field. IEEE Trans. Bio-Med. Eng. *BME-10*, 137–145.

Rupniak, N.M., and Gaffan, D. (1987). Monkey hippocampus and learning about spatially directed movements. J. Neurosci. *7*, 2331–2337.

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

Sakagami, M., and Niki, H. (1994a). Encoding of behavioral significance of visual stimuli by primate prefrontal neurons: relation to relevant task conditions. Exp. Brain Res. *97*, 423–436.

Sakagami, M., and Niki, H. (1994b). Spatial selectivity of go/no-go neurons in monkey prefrontal cortex. Exp. Brain Res. *100*, 165–169.

Schultz, W., Apicella, P., and Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. J. Neurosci. *13*, 900–913.

Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. Science *275*, 1593–1599.

Schwartz, M.L., and Goldman-Rakic, P.S. (1984). Callosal and intrahemispheric connectivity of the prefrontal association cortex in rhesus monkey: relation between intraparietal and principal sulcal cortex. J. Comp. Neurol. *226*, 403–420.

Sereno, A.B., and Maunsell, J.H.R. (1998). Shape selectivity in primate lateral intraparietal cortex. Nature *395*, 500–503.

Spitzer, H., and Richmond, B.J. (1991). Task difficulty: ignoring, attending to, and discriminating a visual stimulus yield progressively more activity in inferior temporal neurons. Exp. Brain Res. *83*, 340–348.

Spitzer, H., Desimone, R., and Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. Science *240*, 338–340.

Watanabe, M. (1986). Prefrontal unit activity during delayed conditional go/no-go discrimination in the monkey. II. Relation to go and no-go responses. Brain Res. *382*, 15–27.

Watanabe, M. (1992). Frontal units of the monkey coding the associative significance of visual and auditory stimuli. Exp. Brain Res. *89*, 233–247.

Williams, G.V., and Goldman-Rakic, P.S. (1995). Modulation of memory fields by dopamine D1 receptors in prefrontal cortex. Nature *376*, 572–575.

Young, B.J., Fox, G.D., and Eichenbaum, H. (1994). Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. J. Neurosci. *14*, 6553–6563.