Cerebral Cortex 2007;17:i41-i50
doi:10.1093/cercor/bhm070

The Representation of Multiple Objects in Prefrontal Neuronal Delay Activity

Melissa R. Warden and Earl K. Miller

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

The ability to retain multiple items in short-term memory is fundamental for intelligent behavior, yet little is known about its neural basis. To explore the mechanisms underlying this ability, we trained 2 monkeys to remember a sequence of 2 objects across a short delay. We then recorded the activity of neurons from the lateral prefrontal cortex during task performance and found that most neurons had activity that depended on the identity of both objects while a minority reflected just one object. Further, the activity driven by a particular combination of objects was not a simple addition of the activity elicited by individual objects. Instead, the representation of the first object was altered by the addition of the second object to memory, and the form of this change was not systematically predictable. These results indicate that multiple objects are not stored in separate groups of prefrontal neurons. Rather, they are represented by a single population of neurons in a complex fashion. We also found that the strength of the memory trace associated with each object decayed over time, leading to a relatively stronger representation of more recently seen objects. This is a potential mechanism for representing the temporal order of objects.

Keywords: memory, monkey, multi-item, order, sequence, temporal

Introduction

As an animal interacts with its environment, it is faced with a constant barrage of stimuli that it must sense, remember, and use. Many complex behaviors require that multiple pieces of information be held in short-term memory, and the sequencing of information is critical for planning. Yet in the laboratory, most studies employ tasks utilizing only the memory for a single item.

We sought to investigate how multiple objects might be held in the lateral prefrontal cortex (PFC), a brain region that has long been thought to be involved in short-term memory for objects, spatial locations, and other information. Damage or inactivation of the PFC in nonhuman primates causes performance deficits on tasks that impose a memory delay between a cue and a response based on that cue (Mishkin 1957; Gross and Weiskrantz 1962; Fuster and Alexander 1970; Goldman et al. 1971; Passingham 1975; Mishkin and Manning 1978). Correspondingly, neurophysiological experiments have established the existence of sustained memory delay activity in the PFC that reflects the identity of the remembered cue and/or forthcoming behavioral response (Fuster and Alexander 1971; Kubota and Niki 1971; Funahashi et al. 1989; Miller et al. 1996); this activity has been interpreted as the neural signature of short-term memory.

Basic questions about how this activity might underlie the holding of more than one object in memory remain unanswered. One possibility is that separate populations of neurons are responsible for the memory of each object. This would be analogous to an address in computer memory, where each object to be stored is placed in its own "box" (memory location). An alternative is that information about each object is stored in a single neuronal population. If so, how is information about multiple objects combined? Would the delay activity of a neuron representing a single object relate in a straightforward way to that neuron's multi-object activity? Addressing such questions can lead to insight into fundamental issues of neural coding.

We trained 2 monkeys to remember 2 sequentially presented objects and found that the majority of neurons had activity that reflected the identity of both objects in memory rather than the alternative possibility of a separate population of neurons for each object. The interaction between the representations of the 2 objects was complex and could not be modeled well as an addition of the activity driven by the individual objects. In fact, the second object appeared to have a profound impact on the way the first object was coded. We also found that the second object was represented more strongly than the first object, which is a possible neural representation of the order of the 2 objects.

Materials and Methods

Subjects and Surgery

The subjects were 2 rhesus monkeys (*Macaca mulatta*), one male and one female, weighing 6.0 and 6.5 kg. Eye movements were monitored and stored using an infrared eye-tracking system (ISCAN, Burlington, MA). Using previously described methods (Miller et al. 1993), monkeys were implanted with recording chambers and with a head bolt to immobilize the head during neuronal recordings. The location of the recording chambers and the location of recording penetrations were determined by structural magnetic resonance imaging (MRI) scans. Recording chambers were placed over the lateral PFC, centered over the principal sulcus, and anterior to the arcuate sulcus. 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 National Institutes of Health guidelines and the recommendations of the MIT Animal Care and Use Committee.

Bar-Release Sequence Task

Monkeys performed a 2-object sequence memory task (or delayed-match-to-sequence, Fig. 1) that required them to judge if 2 successively presented sequences of 2 natural objects were the same. The task was administered and behavior monitored by 2 computers running the "CORTEX" real-time control system (http://www.cortex.salk.edu). The trial began when the monkeys grasped a lever and fixated a small (0.15°) white spot at the center of a CRT screen. They were required to maintain fixation within a ±1.5° square window around the fixation spot for the entire trial. After the initial 1000 ms of fixation, an object was presented at the center of the screen for 500 ms. The object was then extinguished and was followed by a 1000-ms memory delay

(the 1-object delay). A second object was then presented for 500 ms and was also followed by a 1000-ms memory delay (the 2-object delay). The 2 delay periods were of identical duration to facilitate the comparison of these 2 epochs. The presentation of these 2 objects constituted the sample phase of the task because the monkeys were required to remember both of these objects throughout the duration of the trial. The sample phase was followed by the presentation of a temporally identical test sequence, again consisting of 2 objects presented on the screen for 500 ms each, separated by a 1000-ms delay. If the test sequence exactly matched the sample sequence, the monkeys were required to release the lever within 900 ms following the onset of the second test object in order to receive a juice reward. If the test sequence differed in any way from the original sample sequence (if either of the objects was different or their order was reversed), the monkey was required to continue holding the lever until a second test sequence was presented. This second test sequence was always a match and thus required a lever release. As a result, a sequence judgment was only required for the first test sequence; the second test sequence was used so that a behavioral response would be required on every trial. This ensured that the monkeys were always paying attention. Note that with this design, the behavioral response (lever release) is not uniquely associated with a sequence (it was used to signal "match" not a particular sequence), and further, the monkeys could not predict whether the first test sequence would require a response. Thus, any differential activity to the sample sequences could not be related to the behavioral response. Fifty percent of all trials were nonmatch trials, and 50% were match trials. A 1000-ms intertrial interval followed all trials.

For each recording session, 4 novel cue stimuli, never before seen by the animal, were chosen at random from a database of images (Corel, Ottawa, Canada). The stimuli were small complex objects about 2° × 2° in size. The objects were presented on a computer screen positioned directly in front of the animal. We made no attempt to determine which features of particular objects were responsible for the neurons' responses; for this study, it was necessary only that different objects elicited selective activity from a number of PFC neurons. Complex objects were used because they have been shown to elicit robust activity from lateral prefrontal neurons (Miller et al. 1996). Each of the 4 sample objects had the same chance of appearing as the first object (25%) and of appearing as the second object (25%). All combinations of 2 objects in sequence were used, including the 4 sequences composed of a single object shown twice (e.g., the sequence A-A), for a total of 16 sequences.

A crucial feature of our design was that object presentation was balanced. That is, across a set of trials in which a given sample object was the first or second cue, it was followed or preceded by each and every sample object with equal frequency. Thus, when we sort trials by the first object, any neural selectivity seen is directly attributable to that object, even after presentation of the second object because the influence of the object used as the other cue is factored out across trials.

Three types of nonmatching test sequences were used to ensure that the monkeys were remembering the sequence correctly (Fig. 2*A*). One type of nonmatch was that in which the first object changed and the second object remained the same. This nonmatch was used to ensure that the monkey remembered the first object—it would be impossible to correctly respond to this type of trial if the monkey only remembered the second object. The second type of nonmatch was a sequence in which the first object stayed the same but the second object changed. This was used to test the memory of the second object. The third type of nonmatch was that in which the same objects were used, but they were presented in the reverse order. This type of nonmatch was used to ensure that the monkeys were remembering the objects in the correct order. The monkeys performed well on all types of trials (Fig. 2*B*; first object 91% correct; second object 85% correct; order 95% correct; chance on all conditions was 50% when matched trials were included), indicating that they were remembering both objects and the order in which they were presented. A limitation of our task design is that it is, unfortunately, impossible to analyze the neural data from the error trials in any meaningful way. If the monkey made an error, it could be because of a faulty memory for the first item, the second item, the order of the items, or even combinations of these factors. In all cases, all we know is that the trial was incorrect, not why.

Recording Technique

Electrode penetration sites (Fig. 3*A*) were determined using MRI scans obtained prior to surgery. The recording chambers were positioned stereotaxically over the left lateral PFC of each animal such that the principal sulcus and lateral PFC were readily accessible (Fig. 3*A*). During the course of our experiments, we qualitatively determined that the cells in the ventral half of the recording chambers were more likely to respond selectively to objects and therefore concentrated our experimental sessions in this region.

Monkeys were seated in primate chairs within sound-attenuating enclosures (Crist Instruments, Damascus, MD). Their heads were restrained, and a juice spout was placed at their mouths for automated

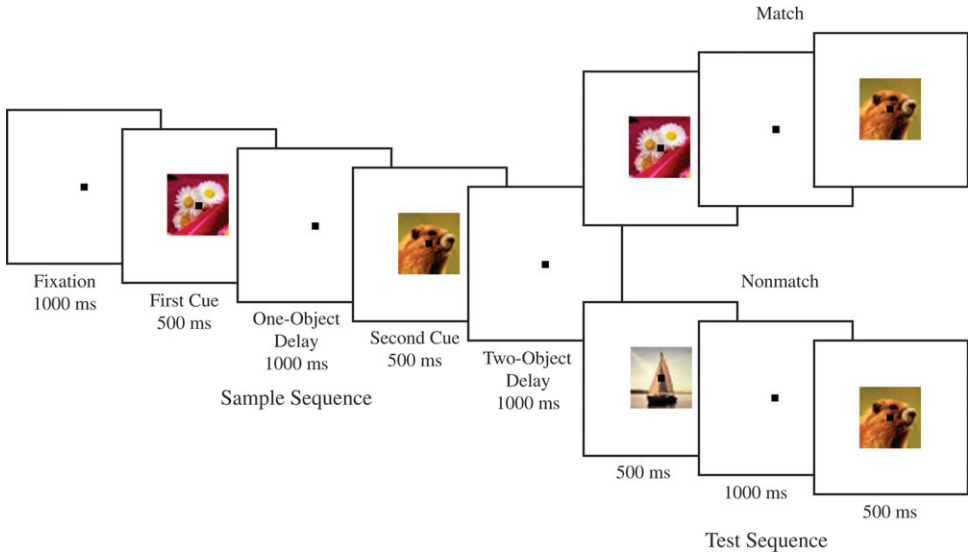

Figure 1. Behavioral task. The monkey was presented with a sequence of 2 objects, which consisted of one sample object, a 1-object delay period, a second sample object, and a 2-object delay period. This was followed by the presentation of a test sequence that had the same temporal structure as the first. If this test sequence matched the sample sequence, the monkey was rewarded for releasing a lever during the presentation of the second matching test object. If the test sequence was not an exact match, the monkey was required to continue grasping the lever until a match sequence appeared. A match sequence always appeared immediately following a nonmatch test sequence. See Materials and Methods for further information.

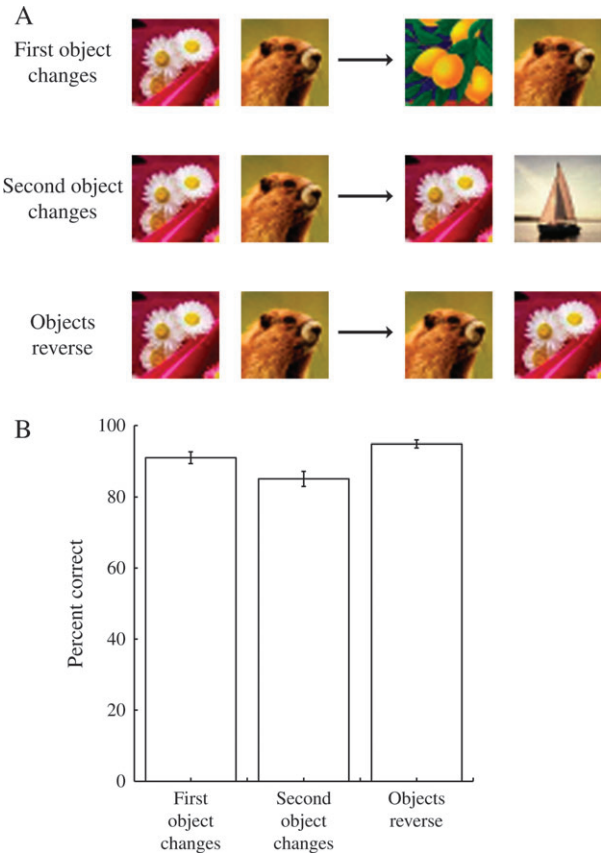

Figure 2. Test trial types. (*A*) Three types of nonmatching sequences were used to ensure that the monkey was correctly remembering the entire sequence. The first pair of objects in each row is the sample sequence, and the second pair is the test sequence. First row: memory for the first object. Second row: memory for the second object. Third row: memory for order. (*B*) Behavioral performance. The monkeys performed well on all 3 types of test sequence. The percent correct for each type of test sequence is shown; error bars represent the 95% confidence interval around the mean. The accuracy rate was 91% for the first condition (first object), 85% for the second condition (second object), and 95% for the third condition (order). Chance performance was 50% for each condition.

reward delivery. Recordings were made using arrays of 8 independently movable dura-puncturing tungsten microelectrodes (FHC Instruments, Bowdoinham, ME). The electrodes were advanced using custom-made screw-driven mini-microdrives (Nichols et al. 1998) mounted on a plastic grid (Crist Instruments) with 1-mm spacing between adjacent locations. Neuronal activity was amplified, filtered, and stored for off-line sorting into individual neuron records (Plexon Systems, Dallas, TX). We did not prescreen neurons for task-related activity such as visual responsiveness or stimulus selectivity. Rather, we randomly selected neurons for study by advancing 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 12 individual neurons (an average of 5.8 per recording session).

Analysis of Neural Data

Data were analyzed using custom-written routines in MATLAB (Mathworks, Natick, MA). We focused our analysis on 5 epochs, which we define here. The "fixation" period consisted of the 500 ms immediately preceding stimulus onset. The "first cue" period began 100 ms after the onset of the first object and had a duration of 400 ms. The first 100 ms were excluded to compensate for the minimum latency of visual responses in PFC. The "1-object delay" period started 200 ms after the offset of the first cue and had a duration of 800 ms. Likewise, the "second cue" period started 100 ms after the onset of the second object and had a duration of 400 ms, and the "2-object delay" period started 200 ms after the offset of the second object and had a duration of 800 ms. These epochs were chosen for simplicity. The results reported here were insensitive to the exact time windows used. All neural activity histograms were calculated with a resolution of 1 ms and then smoothed with a 50-ms boxcar window.

Analysis of Variance, Neural Histograms

To assess the effect of each of the 2 objects on neural activity, a 2-way ANOVA was performed for each neuron on the activity during each epoch. A significant effect of the first or second object meant that activity varied significantly with the identity of the first or second object during the analysis epoch. If the effect of one of the objects on neural activity depended on the identity of the other object, this would produce a significant interaction between objects, and neurons that showed such an interaction were included among the population of neurons selective for both objects. All ANOVAs were evaluated at $P < 0.05$.

Normalization

Normalized data (except for that used in the response surface analyses, see Materials and Methods below) were generated by dividing the firing rate obtained with a particular object during an epoch by the average firing rate during that epoch. This had the effect of transforming the mean firing rate during that epoch to 1 and was done in order to be able to compare epochs with very different firing rates.

Regression Analyses

For the 1-object regression analysis, we regressed the activity of each neuron during the 1-object delay period (D_1) against its activity during the first cue period (C_1) using the linear equation:

$$D_1 = \alpha + \beta C_1,$$

where β is the slope and α is the offset. If the cue and delay period activities were positively correlated, the regression would yield a positive slope (β). All the neurons used in this regression analysis were selective for the first object during both the first cue period and the 1-object delay period.

For the 2-object regression analysis, we regressed the 2-object delay activity (D_2) of each neuron against its activity during both the first cue period (C_1) and the second cue period (C_2) using the 2-factor linear model:

$$D_2 = \alpha + \beta_1 C_1 + \beta_2 C_2,$$

where β_1 and β_2 are the slopes and α is the offset. A positive β_1 indicates a positive correlation between the first cue and 2-object delay periods, whereas a positive β_2 indicates a positive correlation between the second cue and 2-object delay periods. In our regression model, the tilt of the plane and its offset were unconstrained to allow differential weighting of each object. All the neurons used in this regression analysis were selective for the first object during the first cue period, the second object during the second cue period, and both objects during the 2-object delay period. We used this population because the purpose of this analysis was to use the firing rates observed during the cue periods to predict the response during the 2-object delay period.

Response Surface Analysis

The axes used in these figures are 1) the response to the first object during the first cue period, 2) the response to the second object during the second cue period, and 3) the response to the combination of both objects during the 2-object delay period. The first object and second object firing rates were transformed to a zero to one scale by fixing the lowest firing rate at zero and the highest firing rate at one, with the 2 other rates linearly scaled. The 2-object delay response rates were similarly transformed. This normalization was used because we wanted a metric of selectivity that would have the same range for all neurons. The normalized 2-object delay response was then plotted as a function of both the first object and the second object responses. In the single neuron response surfaces, the values between the data points were linearly interpolated. The population response surface was an average of

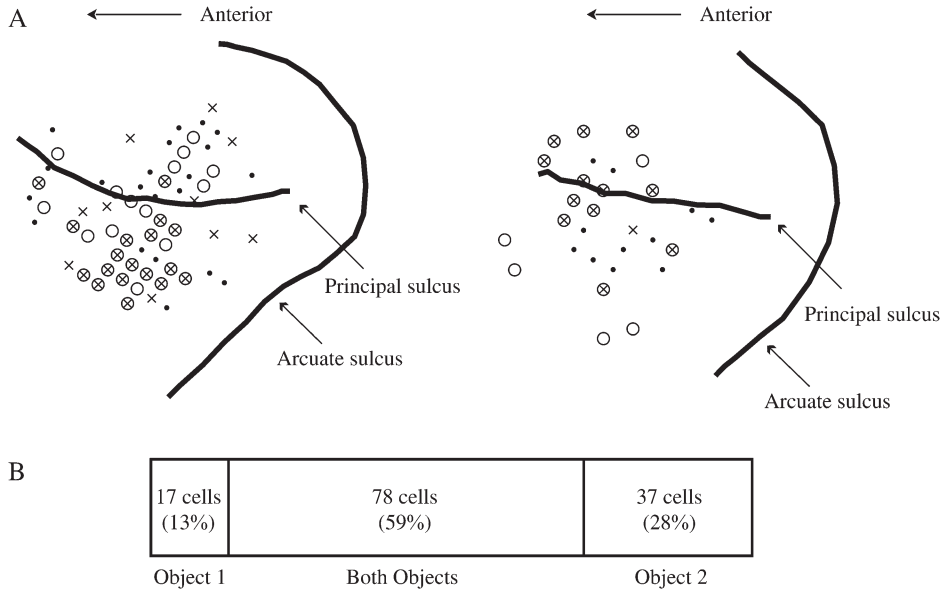

Figure 3. (*A*) Anatomical locations of recording sites and object-selective neurons in both monkeys. X and O, recording sites at which neurons selective for the first object or the second object during the 2-object delay period were found, respectively. Black dots, locations at which neurons were recorded, but no object-selective neurons were encountered. Multiple neurons were recorded at many locations. (*B*) Relative proportions of neurons selective during the 2-object delay period for only the first object, both objects, or only the second object. Area is to scale.

these surfaces over all neurons selective for both objects during the 2-object delay period.

Percent Variance Explained by Each Object

We calculated a sliding estimate of the percent variance explained by each object across the sample phase of the trial. We used a 200-ms sliding time window that moved forward every 20 ms, and an ANOVA was calculated for each window. Simple-effects ANOVAs were used in these analyses instead of 2-way ANOVAs because of the presence of a large amount of interaction between the first and second objects. The resultant sums of squares for each ANOVA were used to estimate the percentage of variance attributable to either the first or the second object for each neuron (Sokal and Rohlf 1995) as a function of time. All neurons were then averaged together, yielding a population estimate of the average percentage of variance explained by each object. All neurons contributed to the variance component figures shown in this paper, although repeating the analysis using only object-selective neurons did not alter the pattern of the results. The only effect of this modification was to increase the overall percentage of variance explained.

Results

Visual Responsiveness

A total of 222 lateral prefrontal neurons were recorded from the left hemispheres of 2 monkeys during performance of the 2-object sequence task (121 from *monkey A* and 101 from *monkey S*). Most of the neurons showed a significant change in activity relative to baseline activity during one or more of the task epochs (206/222 or 92.8%, 112 from *monkey A* and 94 from *monkey S*; 2-tailed *t*-test, evaluated at $P < 0.05$). In any single epoch, many neurons were responsive (128/222 or 57.7% during the first cue period; 150/222 or 67.6% during the 1-object delay period; 159/222 or 71.6% during the second cue period; and 142/222 or 64.0% during the 2-object delay period).

Object Selectivity

To identify neurons whose activity varied with either the first object or the second object, a 2-factor ANOVA (one factor for each object, evaluated at $P < 0.05$) was performed on each neuron during each trial epoch (see Materials and Methods). A majority of neurons (163/222, 73.4%) showed activity that varied significantly with the identity of at least one of the objects during at least one trial epoch. Table 1 shows the incidence of selectivity for each epoch; from one-third to one-half of all neurons showed selectivity during a given epoch.

Neural Activity when 2 Objects Are Held in Memory

Our first interest was to determine whether information about each of the 2 objects was maintained in a separate population of PFC neurons or if information about both objects was somehow combined on the single neuron level. We found examples of both, but the majority of object-selective neurons exhibited activity that depended on the identity of both objects.

During the 2-object delay period, over half of the recorded neurons (132/222, 59.5%) showed activity that varied significantly with one or both objects. There was no obvious topography for first object selectivity or second object selectivity across the recording sites (Fig. 3*A*); the signals appeared to be intermingled. The majority of these object-selective neurons (78/132, 59.1%) showed selectivity that depended on the identity of both objects ("2-object" neurons). Fewer neurons showed selectivity for only the first (17/132, 12.9%) or only the second (37/132, 28.0%) object during the 2-object delay period (Fig. 3*B*). Some 2-object neurons (30/78, 38.5%) required both objects to become selectively activated and did not show any selectivity until after the presentation of the second object. Other 2-object neurons (48/78, 61.5%) were also selectively activated during the first cue presentation, the 1-object delay period, or the second cue presentation.

The Relationship of 2-Object Activity to 1-Object Activity

Given that the activity of the majority of single PFC neurons showed activity that reflected a combination of 2 objects, we wondered whether this 2-object activity bore any simple relationship to the neural activity elicited by a single object.

Table 1

Percentage of neurons selective for the first or second object during the first cue, the 1-object delay, the second cue, or the 2-object delay

	First cue	1-object delay	Second cue	2-object delay	Any epoch
First object					
Neurons	76	82	81	95	144
Percentage of 222	34.2	36.9	36.5	42.8	64.9
Second object					
Neurons			96	115	135
Percentage of 222			43.2	51.8	60.8
Either object					
Neurons			107	132	163
Percentage of 222			48.2	59.5	73.4
Both objects					
Neurons			70	78	116
Percentage of 222			31.5	35.1	52.3
First X second objects					
Neurons			60	63	95
Percentage of 222			27.0	28.4	42.8

For example, was 2-object activity a simple addition of the activity driven by each single object?

At this point, it is important to note one of the crucial features of our experimental design: the design was completely balanced in that each possible first sample object was followed equally often by each possible second sample object. The converse was also true; each possible second sample object was preceded equally often by each possible first sample object. This allowed us to disambiguate the signals related to the first and second objects and to follow each signal independently throughout the course of the trial. This design ensures that if the second object simply erased the effects of the first object (as one might expect to find in a primary sensory area), the neuron would show no selectivity for the first object during the latter phase of the trial. However, if activity related to the first object was still carried by the neuron, this task design would allow us to extract that signal.

A single neuron is shown in Figure 4. When the trials are grouped according to the identity of the first sample object (Fig. 4*A*), activity during the first cue period is strongest for a particular first sample object (object "D"). This is maintained through the 1-object delay period immediately after the first cue presentation. In Figure 4*B*, the sample objects are arranged in descending preference as determined by first object selectivity during first cue presentation. A descending sample object preference curve during the 1-object delay indicates that this neuron maintains the same preferred sample object into this delay period.

But note this neuron's pattern of selectivity during the 2-object delay period when the trials are still grouped according to the identity of the first object (Fig. 4*A*,*C*). Because object presentation is balanced, the selective activity seen after presentation of the second object is attributable to the first object (see above). The neuron's activity still varies with the identity of the first object but in quite a different way than seen during first cue presentation. During the 2-object delay, the firing rates corresponding to the 4 different first sample objects have reversed order; now, the activity is lowest when object D was seen as the first object. Also note that when we sort the trials according to the identity of the second object (Fig. 4*D*), neural selectivity is strongest for object D, as it was when object D was the first object. Thus, it seems that adding a second object to memory alters neural selectivity to a previous (the first) object.

The same effects could be seen at the population level. Figure 5*A* shows a population average of the same analyses seen above (i.e., Fig. 4*B*). Activity was sorted by the identity of the first object and averaged over all the neurons selective for the first object during both the first cue and 1-object delay periods. The similarity of the curves shows that population neural selectivity during the first cue presentation continues through the 1-object delay period. However, the population selectivity for the first object changes after the second object is added (Fig. 5*B*, population average taken over all the neurons selective for the first object during both the first cue and 2-object delay periods). In fact, the average population activity curve during the 2-object delay period is flat when activity is sorted by the identity of the first object. This is not because selectivity is absent; it is because adding the second object changed selectivity for the first object in a nonsystematic fashion across the neuron population. Some neurons invert their preferences (as did the neuron of Fig. 4), some changed in a different fashion, and some maintained their preference. It is worthwhile noting here that the proportion of neurons selective for the first object is roughly comparable during the 1-object and 2-object delay periods (36.9% in the 1-object delay period and 42.8% in the 2-object delay period, Table 1). Only the neurons that were selective for the first object during the 2-item delay period went into this analysis, so neural selectivity for the first object has not disappeared—it has just changed form.

We performed a regression analysis to further examine the correspondence between first object-related activity during cue presentation and the 1-item delay (Fig. 6*A*), primarily in order to determine how this relationship changed with the presentation of the second object. We fit a linear model to each neuron using its activity during the first cue period as the regressor and the activity during the 1-object delay period as the response (see Materials and Methods for details). A positive β (slope) indicates that the response of the neuron during the 1-object delay period varies directly with its response during first cue presentation. All βs were tested for a significant difference from zero (2-tailed *t*-test, $P < 0.05$). Most neurons in the population had βs greater than zero, and the distribution of βs was significantly greater than zero (1-tailed *t*-test, $P < 0.05$). These results lend support to the conclusion that object preferences during the 1-object delay period are very similar to object preferences during cue presentation.

To attempt to gain a better understanding of how information about the 2 objects are combined in neural activity, we tried to predict the level of activity in the 2-object delay period from the activity elicited by single objects using a similar approach. To do this, we fit a 2-factor linear model to the data (see Materials and Methods for details) using the activity driven by the first object and the second object during their presentations as the regressors. This is perhaps the most intuitively plausible model; it would be a simple linear combination of activity to the 2 objects. The population of cells used for this analysis was significantly selective for the first object during first cue presentation, the second object during second cue presentation, and both objects during the 2-item delay period.

We found that the distribution of β_1s, reflecting the correlation between the first cue period selectivity and the 2-object delay period selectivity for the first object, was not significantly different from zero (Fig. 6*B*, 1-tailed *t*-test, $P > 0.05$). There are many neurons with slopes that are significantly different from zero, as would be expected, given that all the neurons in this

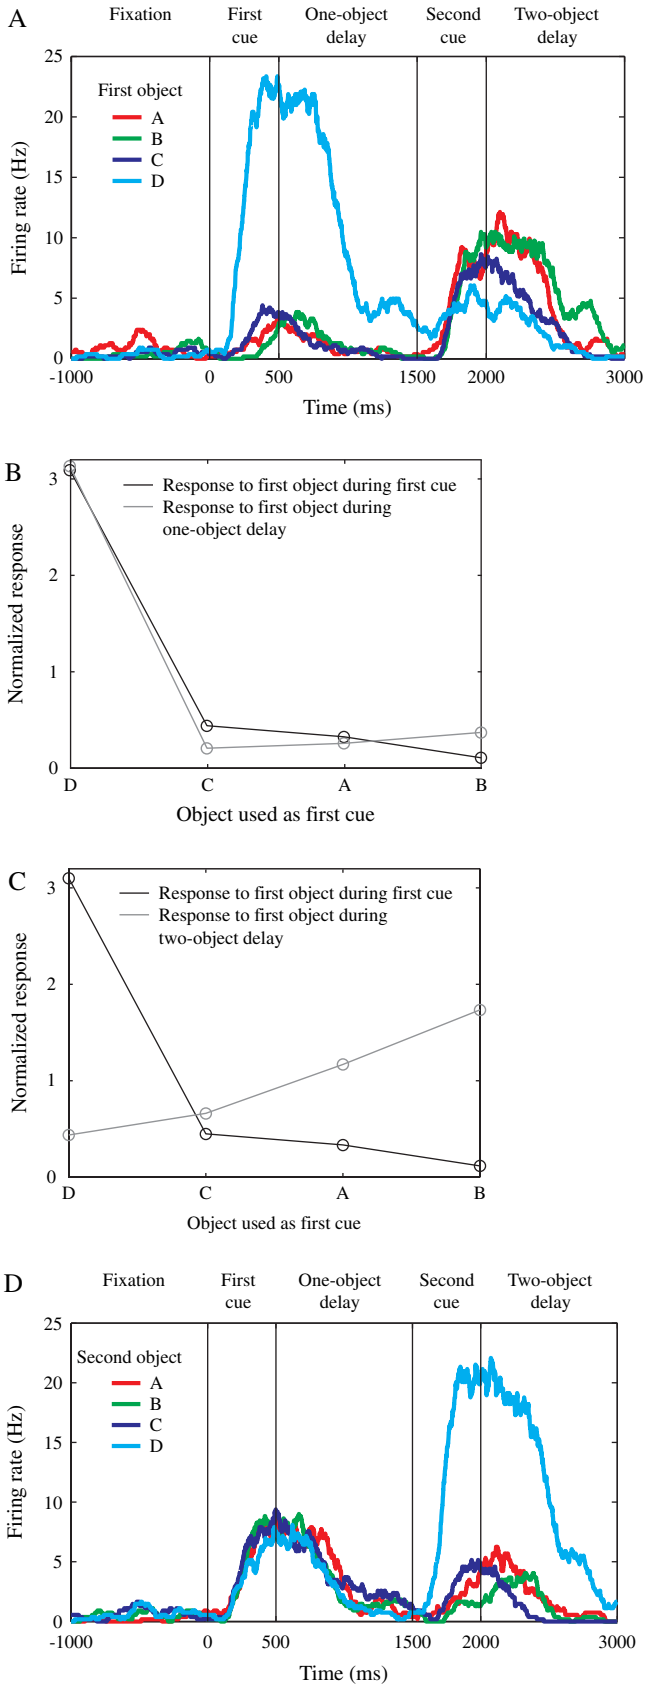

Figure 4. (*A*) Activity of a single prefrontal neuron, trials grouped according to which object appeared as the first object. This cell shows selective activity for the first object during both the 1-object and 2-object delay periods, but the preferred object has changed between these 2 epochs. (*B*) Normalized response of this neuron to the first object during both the first cue period and the 1-object delay period. The similarity of the 2 curves shows that this neuron maintains its object preference during the 1-object delay. (*C*) Normalized response of this neuron to the first object during both the first cue period and the 2-object delay period. The curves do not look the same, indicating that this neuron has changed its preferred first object. (*D*) Activity of the same neuron, now grouped according to which object appeared as the second object.

figure have object selectivity for the first object during the 2-object delay period. However, there are just as many negative slopes as positive slopes, indicating that these neurons have changed their object preferences as a result of the addition of the second object to memory. The neurons with nonsignificant slopes are also still selective for the first object, as assayed by ANOVA, but they have changed their first object coding in a nonlinear fashion. The centering of the mean on zero also indicates that while the memory of the first object is preserved after the addition of a second object to memory, it has changed form in a nonsystematic fashion.

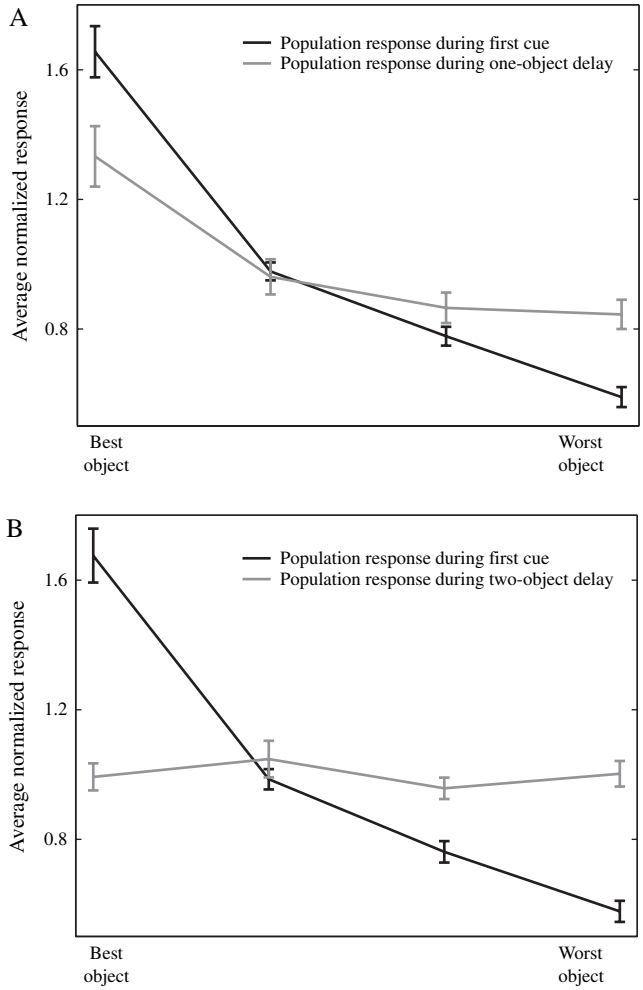

Figure 5. (*A*) The same analysis shown in Figure 4*B*, averaged over all neurons selective for the first object during both the first cue and 1-object delay periods. Population object preference is maintained into this delay period. (*B*) The same analysis shown in Figure 4*C*, averaged over all neurons selective for the first object during both the first cue and 2-object delay periods. Many neurons have changed their preferred first object and, when averaged together, produce a flat response curve.

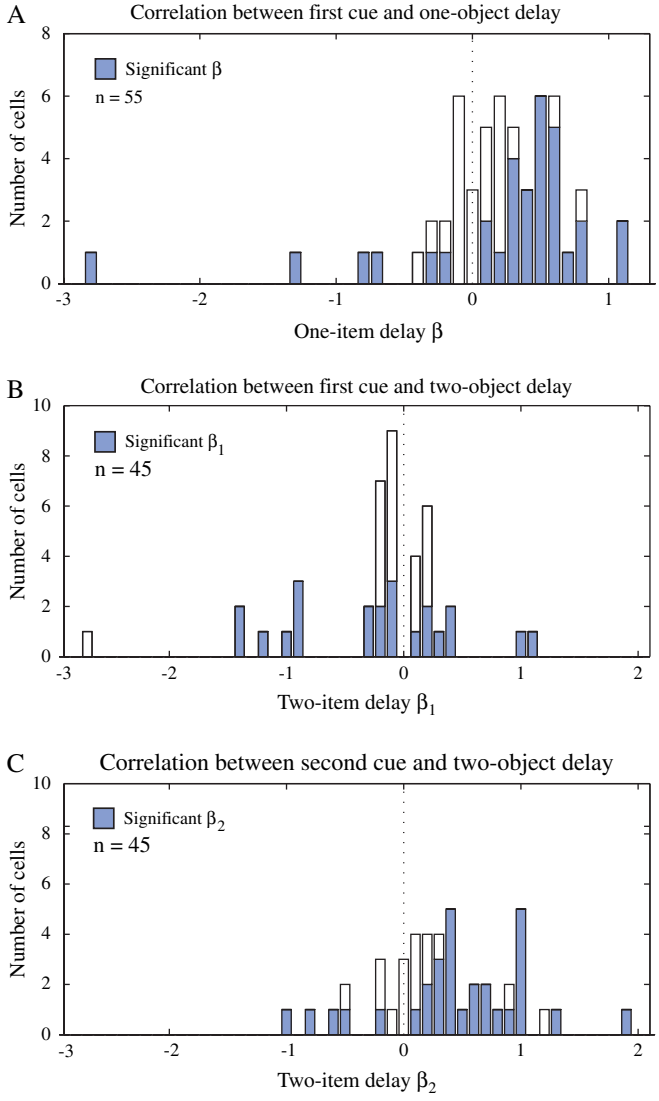

Figure 6. (*A*) Delay activity regressed against cue activity. The histogram shows the distribution of βs (regression slopes) across the population of neurons selective for the first object during both the first cue and 1-object delay periods. Colored bars are neurons with significant β ($P < 0.05$). The mean of the distribution was significantly greater than zero ($P < 0.05$), indicating that neurons maintain their object preferences into the 1-object delay period. (*B*) The distribution of the values of β_1 obtained when every selective neuron (selective for the first object during both the first cue and 2-object delay periods and selective for the second object during both the second cue and 2-object delay periods) was fit to a 2-factor model. The mean of the distribution was not significantly greater than zero ($P > 0.05$). These cells continue to encode the first object but have changed their preferred first object. (*C*) The distribution of the values of β_2 using the neural population as defined in part B. The mean of the distribution was significantly greater than zero ($P < 0.05$), indicating that these neurons are maintaining their preferred second object into the 2-object delay period.

The mean of the distribution of β_2s, on the other hand, is significantly greater than zero (Fig. 6*C*, 1-tailed *t*-test, $P < 0.05$). There are many neurons with significant slopes, and most of these slopes are positive. This demonstrates that the preferred second object during the second cue presentation is similar to the preferred second object during the 2-object delay period. This is consistent with the single neuron examples and population averages presented above: neural activity during presentation of a single object simply carries forward to the immediately following delay but addition of a second object to memory changes the relationship between the first object and neural activity in a nonsystematic fashion across different neurons.

To get a better idea of how adding the second object changes neural activity, it is informative to look at actual response surfaces obtained from single neurons without fitting them to a model. Examples of response surfaces for 2 single neurons are shown in Figure 7*A,B*. We first normalized the activity of each neuron in order to easily compare the response surfaces of different neurons (see Materials and Methods for definitions and normalization). We then plotted the 2-object delay activity of the neuron as a function of its response to the 2 objects. Figure 7*A* shows one of the few neurons that had a maximal response when the best first object and the best second object were used as the sequence of objects. The neuron shown in Figure 7*B* is more representative of the population in the sense that neural activity to the 2 objects did not combine in a straightforward fashion. It shows the highest activity in the 2-object delay when the worst first object is used in combination with the best second object.

We computed the population average response surface by averaging all the individual response surfaces together (Fig. 7*C*). Only neurons selective for both objects during the 2-object delay were used in this analysis. Note the large net positive slope in the second object direction and no net slope in the first object direction. This means, as suggested by the analyses above, that the first object does not have a systematic effect on each neuron's activity after the second object has been presented. A wide variety of first object-related responses are seen, and on average they produce a flat curve. Also, the activity related to the second object during the 2-object delay is quite faithful to what we would find if it was presented in isolation. These results indicate that the newest object in memory is represented as if it were the only object in memory, whereas the representation of the older object has changed significantly.

One interesting question is whether the representation of sequences with repeated items (such sequences as A-A or B-B) is somehow special. We did not find evidence for such an effect in our data, and this can be seen from the average response surface plot in Figure 7. If these sequences were systematically associated with a higher or lower firing rate, the diagonal of this figure would be higher or lower than the surrounding sequences.

Order-Dependent Effects

Until this point, we have discussed the representation of the 2 objects in memory but have not explicitly considered how their order might be represented. The behavior of the monkeys clearly indicates that they do remember the order in which the objects were presented, and this must somehow be represented in neural activity. One way in which order might be represented would be via neurons specifically selective for individual sequences. Consistent with previous reports (Averbeck et al. 2003; Ninokura et al. 2003), we do indeed find these neurons in our data (not shown). The criteria for categorizing a neuron as a "sequence-selective neuron" are somewhat arbitrary, which makes determining the fraction of the population involved difficult. However, one definition of sequence selectivity would be that a neuron is selective for both the first and the second object during the 2-object delay period and that the neuron requires the presentation of both of these objects to show

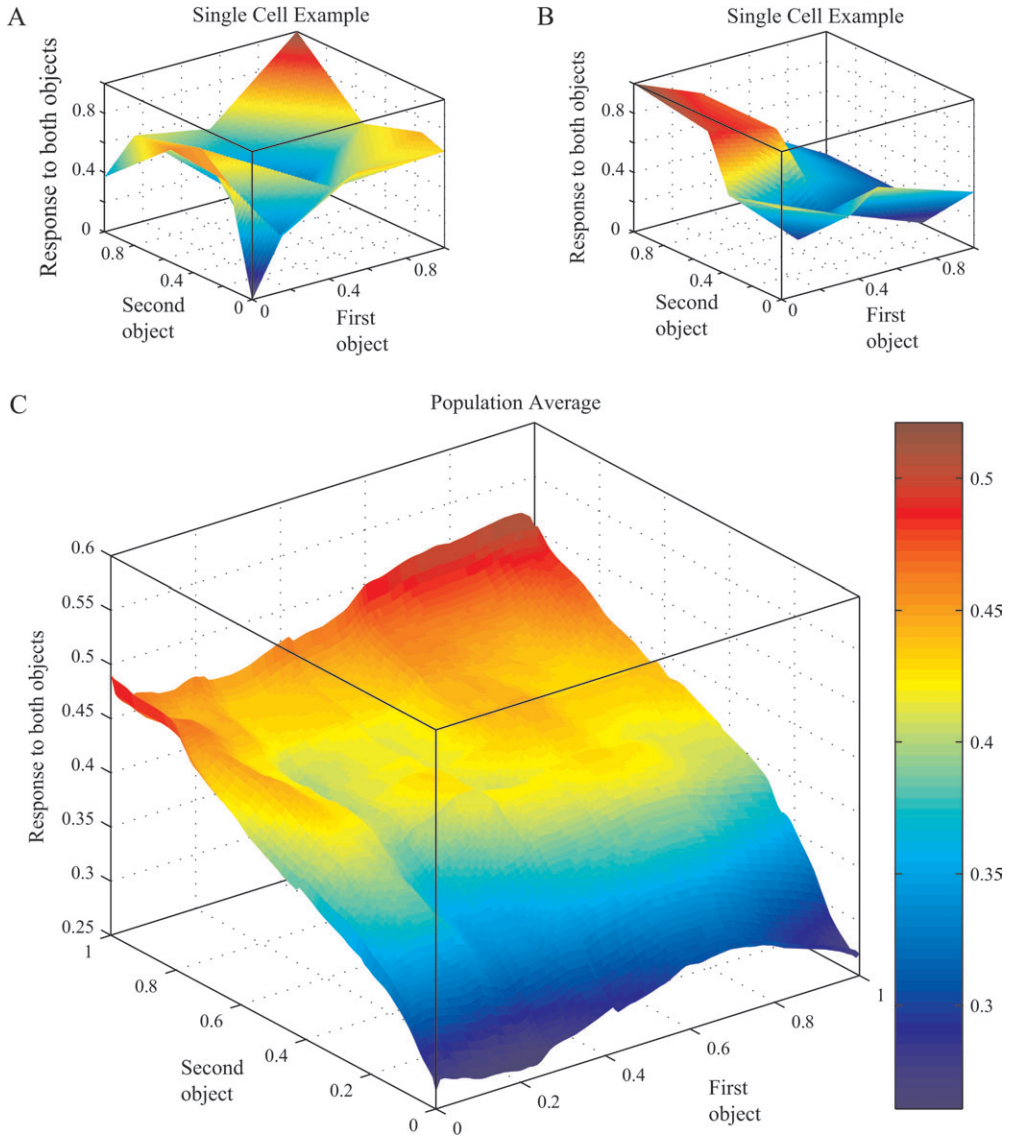

Figure 7. (*A*) An actual response surface from a single neuron is shown. The first object axis is the neuron's response to the first object during the first cue period, and the second object axis is the neuron's response to the second object during the second cue period. The height of the response surface is the neuron's response during the 2-object delay period. This cell maintained its preference for both objects into the 2-object delay period. (*B*) Another response surface from a different cell. This cell changed its preferred first object but maintained its preferred second object into the 2-object delay period. (*C*) The average response surface of all neurons selective for both objects during the 2-object delay period. The population response surface is flat in the first object direction and positively tilted in the second object direction. This indicates that the preferred first object has changed, but the preferred second object has been maintained in the 2-object delay period.

selectivity. By this definition, 30/222 (13.5%) of our neurons are sequence selective.

Another possible way of representing order information would be through the use of order-dependent memory trace strengths. It is possible that the strength of a memory decreases over time or with the addition of a second object, which would mean that more recently seen objects would have a greater effect on neural activity. We did in fact find this effect, as seen in Figure 8. We used the percentage of the variance explained by each object as a measure of the amount of information in the neural population about each object and computed this quantity in a window slid across the length of the trial (see Materials and Methods for details). When the relative strengths of both objects were examined during the 2-object delay period, we found that the second object had a stronger representation in neural activity than the first object.

Discussion

We have found that when a sequence of 2 objects must be remembered across a brief memory delay, both objects are reflected in the activity of single prefrontal neurons. This finding is consistent with the hypothesis that information about multiple objects is combined in a single population of neurons and less compatible with an alternative model in which separate memories are stored in separate neural populations, analogous to addresses in computer memory. Further, we found that there was not a straightforward relationship between the delay period activity corresponding to a single object and the delay period activity reflecting that object after a second object was added to memory. The addition of the second object to memory changed activity related to the first object in a nonsystematic fashion across neurons and in a fashion that was unpredictable from

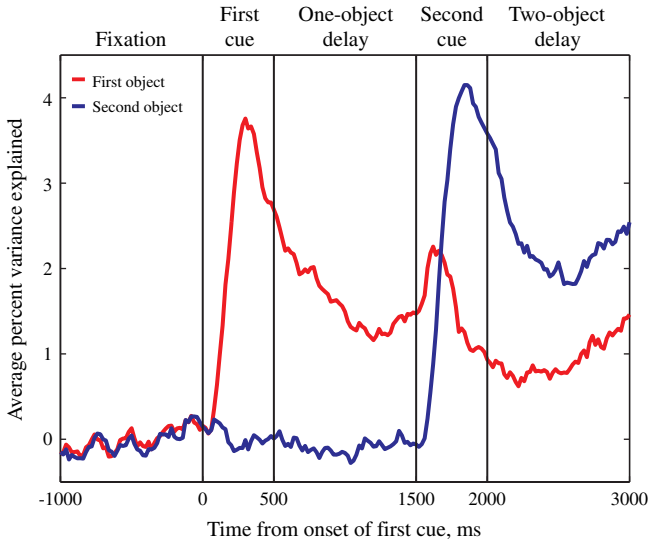

Figure 8. Relative object strengths. Red curve: percent variance explained by the first object, averaged across the entire population of neurons. Blue curve: percent variance explained by the second object, averaged across the population of neurons. During this task, the population of neurons encoded the second object more strongly than the first object during the 2-object delay period.

each neuron's single-object activity. Finally, in addition to single neurons selective for specific sequences of objects, the most recently seen object had a stronger representation in neural activity than the first object. This suggests another possible manner in which sequence information could be maintained.

Relationship to Prior Neurophysiological Studies

There have been several studies demonstrating prefrontal neural activity selective for specific sequences of objects, spatial locations, and movement sequences (Barone and Joseph 1989; Shima and Tanji 2000; Ninokura et al. 2003). In fact, microstimulation of the frontal cortical areas (especially the supplementary eye fields) can disrupt monkeys' ability to perform a remembered saccade sequence but leave memory for the individual saccade locations intact (Histed and Miller 2006). Consistent with our results, other neurophysiological studies have found that single PFC neurons often showed a unique level of activity for specific sequences. Our results extend these studies by addressing the question of how the memory trace for a single object is modified when a new object is loaded into memory. The activity driven by the first object was not simply added to the activity driven by the second object. Rather, the new memory changed the representation of the older object in memory. Also, our results extend previous knowledge by examining the relative strengths of each object in memory. We demonstrated that the signal corresponding to a given object decays over time and that newer objects in memory are represented more strongly than older objects.

One interesting question is whether the lateral PFC is uniquely responsible for maintaining a multi-item memory buffer. An associated question is whether or not different subregions of the PFC contribute in different ways to multi-item memory storage. Our data show no obvious difference between the dorsolateral and ventrolateral PFC with respect to the simultaneous storage of 2 objects by single neurons, which is consistent with a previous study on sequence memory (Ninokura et al. 2003). However, we did find object-selective cells with a higher probability in the ventrolateral PFC and therefore concentrated our recording in this area. As a natural result of this concentration, we found a greater number of 2-object neurons in this area. One other area of great interest for future investigation would be the inferior temporal cortex, which is known to exhibit object-specific delay period activity (Miyashita and Chang 1988).

Comparison to Computational Models

There have been several attempts to model the short-term memory for multiple items as the simultaneous activation of multiple delay circuits. A number of groups have approached this problem, and although the specific systems under study are quite different, the solutions that they propose are similar. Amit and colleagues have created a network of integrate-and-fire neurons that is capable of simultaneous delay activity for up to 6 objects in short-term memory (Amit et al. 2003; Yakovlev et al. 2004). Each object is represented by a distinct population of neurons, each with a "perfectly sharp" tuning curve; in other words, if a neuron responds selectively to a particular object, it has a complete lack of response to all other objects. Because this model assumes that a population of neurons corresponds to each object, multiple objects are represented by the activation of multiple, corresponding populations of neurons. This model is perfectly additive, that is, memories for different objects neither directly interfere with each other nor does adding objects to memory affect the way a previous object is stored.

Another approach has been the creation of a neural network for the simultaneous representation of several spatial targets (Tanaka 2002a, 2002b). This is also an integrate-and-fire network, and it relies on a topographically organized spatial map with hills of activity representing the memory of particular spatial locations. Although in this model the neurons are not perfectly sharply tuned, as they are in the Amit model, the same effect is realized because the model represents multiple spatial memories with nonoverlapping hills of activity. This type of network does a very good job at storing multiple memories when they are distant on the spatial map but breaks down when the memories are close enough in space to interfere with each other. Multiple memories can only be stored with high accuracy by completely separate subpopulations within these networks; when memory traces begin to overlap, fidelity is compromised.

These 2 models are similar in that they store memories for multiple objects using distinct populations of neurons. They are also similar in that new memories are stored in the network in an additive way. Our results suggest that the PFC instead combines the memory representations of multiple objects in a single population of neurons; in most cases, the activity of a single neuron was determined by the identity of both objects in short-term memory. Prefrontal neurons rarely had a strong response to one object and a baseline response to all other objects; in most cases, a spectrum of firing rates is observed for a large number of objects. This is difficult to reconcile with models that require completely separate populations of neurons for each object/spatial location. Amit and colleagues (2003) (Curti et al. 2004; Mongillo et al. 2005; Romani et al. 2006) have begun to address these issues through the creation of a more realistic spiking network model that incorporates neurons that respond selectively for more than one object.

The sequence-selective cells that we and others have found (Averbeck et al. 2003; Ninokura et al. 2003) have been predicted by theoretical studies (Grossberg 1978a, 1978b). In particular,

these studies proposed the existence of sequence-selective cells (or "list chunk" cells) that would be activated by one or more items of a sequence that is stored in working memory.

Summary

In conclusion, we have shown that the primate lateral PFC exhibits signals related to the maintenance of multiple objects in short-term memory. We have found that a single population of neurons is capable of encoding 2 objects and that the signal related to a newer object is not simply overlaid on the signal related to an older object. Instead, the memory trace for the newer object appears to change the older signal. In addition, the strength of each object in memory appears to decay as time progresses, with the result that newer objects in memory are represented more strongly than older objects. This difference in strength may be used to encode the temporal order of the objects. It remains to be shown how each memory can be reliably read out and reconstructed based on such a population code.

Notes

We are grateful to E. Antzoulatos, S. Brincat, T. Buschman, J. Cromer, A. Diogo, M. Histed, R. Meyer, M. Puig, J. Roy, and M. Siegel for discussions and comments and to K. MacCully for technical assistance. This work was supported by the National Science Foundation CELEST Science of Learning Center. *Conflict of Interest*: None declared.

Address correspondence to Earl K. Miller, Building 46-6241, 77 Massachusetts Avenue, Cambridge, MA 02139, USA. Email: ekmiller@-mit.edu.

References

Amit DJ, Bernacchia A, Yakovlev V. 2003. Multiple-object working memory—a model for behavioral performance. Cereb Cortex. 13:435–443.

Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP. 2003. Neural activity in prefrontal cortex during copying geometrical shapes. I. Single cells encode shape, sequence, and metric parameters. Exp Brain Res. 150:127–141.

Barone P, Joseph JP. 1989. Prefrontal cortex and spatial sequencing in macaque monkey. Exp Brain Res. 78:447–464.

Curti E, Mongillo G, La Camera G, Amit DJ. 2004. Mean field and capacity in realistic networks of spiking neurons storing sparsely coded random memories. Neural Comput. 16:2597–2637.

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, Alexander GE. 1970. Delayed response deficit by cryogenic depression of frontal cortex. Brain Res. 20:85–90.

Fuster JM, Alexander GE. 1971. Neuron activity related to short-term memory. Science. 173:652–654.

Goldman PS, Rosvold HE, Vest B, Galkin TW. 1971. Analysis of the delayed-alternation deficit produced by dorsolateral prefrontal lesions in the rhesus monkey. J Comp Physiol Psychol. 77:212–220.

Gross CG, Weiskrantz L. 1962. Evidence for dissociation of impairment on auditory discrimination and delayed response following lateral frontal lesions in monkeys. Exp Neurol. 5:453–476.

Grossberg S. 1978a. A theory of human memory: self-organization and performance of sensory-motor codes, maps, and plans. In: Rosen R, Snell F, editors. Progress in theoretical biology. New York: Academic Press. p. 233–374.

Grossberg S. 1978b. Behavioral-contrast in short-term memory—serial binary memory models or parallel continuous memory models. J Math Psychol. 17:199–219.

Histed MH, Miller EK. 2006. Microstimulation of frontal cortex can reorder a remembered spatial sequence. PLoS Biol. 4:e134.

Kubota K, Niki H. 1971. Prefrontal cortical unit activity and delayed alternation performance in monkeys. J Neurophysiol. 34:337–347.

Miller EK, Erickson CA, Desimone R. 1996. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J Neurosci. 16:5154–5167.

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.

Mishkin M. 1957. Effects of small frontal lesions on delayed alternation in monkeys. J Neurophysiol. 20:615–622.

Mishkin M, Manning FJ. 1978. Non-spatial memory after selective prefrontal lesions in monkeys. Brain Res. 143:313–323.

Miyashita Y, Chang HS. 1988. Neuronal correlate of pictorial short-term memory in the primate temporal cortex. Nature. 331:68–70.

Mongillo G, Curti E, Romani S, Amit DJ. 2005. Learning in realistic networks of spiking neurons and spike-driven plastic synapses. Eur J Neurosci. 21:3143–3160.

Nichols AM, Ruffner TW, Sommer MA, Wurtz RH. 1998. A screw microdrive for adjustable chronic unit recording in monkeys. J Neurosci Methods. 81:185–188.

Ninokura Y, Mushiake H, Tanji J. 2003. Representation of the temporal order of visual objects in the primate lateral prefrontal cortex. J Neurophysiol. 89:2868–2873.

Passingham RE. 1975. Delayed matching in rhesus-monkeys with selective prefrontal lesions. Brain Res. 85:190–191.

Romani S, Amit DJ, Mongillo G. 2006. Mean-field analysis of selective persistent activity in presence of short-term synaptic depression. J Comput Neurosci. 20:201–217.

Shima K, Tanji J. 2000. Neuronal activity in the supplementary and presupplementary motor areas for temporal organization of multiple movements. J Neurophysiol. 84:2148–2160.

Sokal RR, Rohlf FJ. 1995. Biometry. 3rd edition. New York: W.H. Freeman and Company.

Tanaka S. 2002a. Dopamine controls fundamental cognitive operations of multi-target spatial working memory. Neural Netw. 15:573–582.

Tanaka S. 2002b. Multi-directional representation of spatial working memory in a model prefrontal cortical circuit. Neurocomputing. 44–46:1001–1008.

Yakovlev V, Bernacchia A, Orlov T, Hochstein S, Amit D. 2005. Multi-item working memory—a behavioral study. Cereb Cortex. 15:602–615.