

Timing and Neural Encoding of Somatosensory Parametric Working Memory in Macaque Prefrontal Cortex

Carlos D. Brody¹, Adrián Hernández², Antonio Zainos² and Ranulfo Romo²

¹Cold Spring Harbor Laboratory, 1 Bungtown Road, Cold Spring Harbor, NY 11724, USA and ²Instituto de Fisiología Celular, Universidad Nacional Autónoma de México, Apartado Postal 70-253, C.P. 04510 México, D.F. México

We trained monkeys to compare the frequencies of two mechanical vibrations applied sequentially to the tip of a finger and to report which of the two stimuli had the higher frequency. This task requires remembering the first frequency during the delay period between the two stimuli. Recordings were made from neurons in the inferior convexity of the prefrontal cortex (PFC) while the monkeys performed the task. We report neurons that fire persistently during the delay period, with a firing rate that is a monotonic function of the frequency of the first stimulus. Separately from, and in addition to, their correlation with the first stimulus, the delay period firing rates of these neurons were correlated with the behavior of the monkey, in a manner consistent with their interpretation as the neural substrate of working memory during the task. Most neurons had firing rates that varied systematically with time during the delay period. We suggest that this time-dependent activity may encode time itself and may be an intrinsic part of active memory maintenance mechanisms.

Introduction

Neural activity that depends on a sensory stimulus and that is persistent, outlasting the presence of the stimulus that initially drove it, can be used to convey information about the stimulus after it is gone. In this sense, the activity can act as a representation of the memory of the stimulus. Persistent stimulus-dependent activity has been found in many systems (for examples, see Fuster, 1973; Gnadt and Andersen, 1988; Funahashi *et al.*, 1989; Miller *et al.*, 1993; Taube, 1998; Prut and Fetz, 1999; Aksay *et al.*, 2000).

One way in which persistent stimulus-dependent activity can represent information is in what has been termed a monotonic firing rate code. Romo *et al.* (1999) trained macaques to perform a task that required remembering a scalar parameter that fell within a finite range (see Fig. 1). In this task, subjects must remember the frequency of a vibrotactile mechanical stimulus during a delay period several seconds long. Romo *et al.* found neurons in prefrontal cortex (PFC) that fired persistently during the delay period, with firing rates that either continually increased as the frequency of the stimulus held in memory increased (positive monotonic neurons; see Fig. 2), or continually decreased as the stimulus frequency increased (negative monotonic neurons). If the monotonic functions relating stimulus and firing rate of each neuron are smooth and broad, as in the example of Figure 2, then all cues in the range of cues used in the experiment drive all participating neurons to appreciable firing rates. Thus, the key factor in determining the identity of the stimulus held in memory is not which neurons are firing above their spontaneous rate, but how much the neurons are firing. Here we will describe the smoothness properties of the monotonic code found by Romo *et al.* (1999).

To decode a smooth monotonic code, one may invert the stimulus \rightarrow firing rate function (Fig. 2c) in order to obtain a firing rate \rightarrow stimulus map. However, Romo *et al.* (1999) found that for many neurons, this map changed as a function of time (see right three panels of Fig. 2b; note changing *y*-axis scales). One could assign meaning to a firing rate (e.g. what stimulus produces 11 spikes/s?) if one knew how much time had elapsed since the memory cue (e.g. if we know that we are in the third second of the delay period, we know that 11 spikes/s corresponds to the lowest memorized stimulus values). Starting with some of the first studies of neural responses during short-term memory (Fuster, 1973; see also Chafee and Goldman-Rakic, 1998), it was reported that many neurons had firing rates that varied systematically with time during the delay period. These neurons constitute a potential timing signal that could allow decoding of time-varying smooth monotonic maps. Here we describe how the overwhelming majority of PFC neurons recorded during the vibrotactile discrimination task had firing rates that varied systematically with time during the delay period. Current models of persistent activity as a basis for short-term memory typically assume that the memory is held in a pattern of firing rates that is essentially constant throughout the delay period (Amit and Brunel, 1997; Zipser *et al.*, 1993; Compte *et al.*, 2000; Wang, 2001; but see Miller *et al.*, 2003). We suggest that the striking prominence of time-dependence in experimentally observed delay period responses argues against models that ignore such time-dependence.

The monotonic encoding we focus on here is not specific to the vibrotactile discrimination task. The neural activity necessary to drive eye muscles so as to keep the eyes still while away from their natural equilibrium position can be thought of as representing a memory of eye position (Lopez-Barneo *et al.*, 1982; Cannon *et al.*, 1983; Seung, 1996). Studies in goldfish have revealed the presence of neurons with firing rates that are persistent during eye fixations and that depend monotonically on eye position (Aksay *et al.*, 2000, 2001). Although very different in many respects, the vibrotactile task (which is trained) and the eye position task (which is part of natural behavior) are similar in that both of them require remembering a scalar parameter that falls within a finite range. In both tasks, there is a well-defined ordinal relationship between possible sensory stimuli to be held in memory (if stimulus A < stimulus B and B < C, then A < C). We speculate that other tasks that satisfy these conditions would also give rise to a monotonic delay period encoding (Romo *et al.*, 1999).

Models for monotonic persistent activity have recently been developed, focusing on eye position behavior (Seung *et al.*, 2000; Koulakov *et al.*, 2002) and on vibrotactile frequency memorization (Miller *et al.*, 2003). In the case of vibrotactile frequency, positive monotonic neurons coexist and are inter-

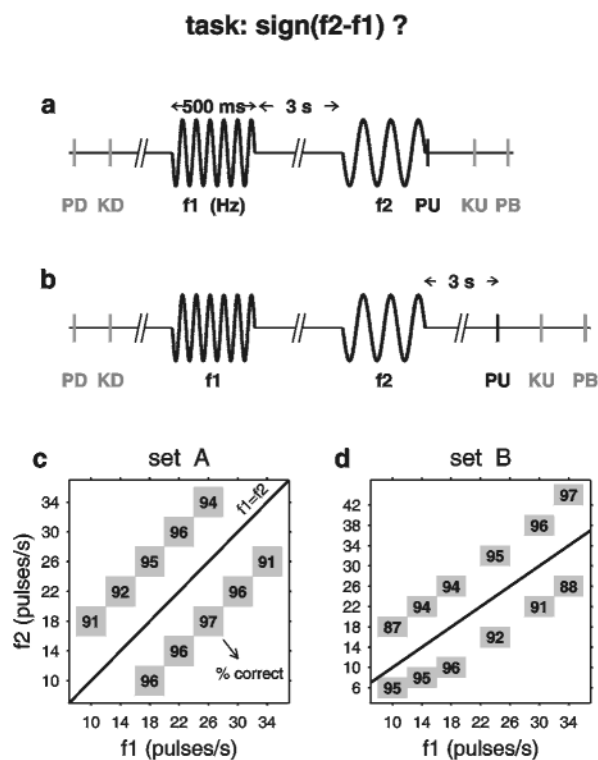


Figure 1. Behavioral task and stimulus sets. (a, b) Sequence of events during the task. PD: mechanical probe moves down onto glabrous (hairless) skin of one finger of a restrained hand. KD: monkey's free hand presses key to indicate readiness for trial. f1: 500 ms long mechanical vibration, followed by a delay period, typically 3 s long. f2: 500 ms long mechanical vibration. PU: mechanical probe is raised, indicating that the monkey may now report its choice regarding the sign of $f_2 - f_1$. KU: monkey's free hand leaves readiness key. PB: Monkey presses one of two pushbuttons to indicate its choice. (a) Standard task, used with four monkeys. (b) Task used with a fifth monkey. In this variant, there is a 3 s delay between the end of f2 and PU. (c, d) Stimulus sets used during recordings. Each grey box indicates an f_1/f_2 pair used; numbers inside grey box indicate overall percentage correct for that f_1/f_2 pair.

persed in the same hemisphere with negative monotonic neurons. Both this aspect and the time-dependence of many responses are considered in the model of Miller *et al.*

Part of the results presented here were described previously in shorter form (Romo *et al.*, 1999). In the present paper, we present additional data (from a further monkey, performing the variant of the task shown in Fig. 1b), as well as additional and more extensive analysis of the data.

Materials and Methods

Five monkeys (*Macaca mulatta*) were trained to perform the vibrotactile discrimination task up to their psychophysical thresholds, which for the range of stimulus frequencies used is on the order of a 2–4 Hz difference between base and comparison (Hernández *et al.*, 1997). All stimuli were kept within the frequency range which, in humans, gives rise to the percept known as TMflutter. Vibration frequencies outside this range give rise to qualitatively different percepts and are transduced, in both humans and monkeys, by cutaneous receptors different to those which transduce flutter (Talbot *et al.*, 1968). Neurophysiological recordings began after training was complete. Two variants of the task were used. In the first (Fig. 1a), four monkeys were trained to respond immediately after the end of the second stimulus (f2). Data collected from these

monkeys were briefly described in Romo *et al.* (1999). A fifth monkey was trained in a second variant of the task (Fig. 1b), in which the monkey had to wait for 3 s after f2 before reporting its decision; a cue (PU, Probe Up in Fig. 1b) then indicated to the monkey that it was free to respond. This variant of the task has two separate short-term memory phases, during the second of which subjects need only to remember the binary outcome of the comparison of f2 to f1. This second delay period thus separates the f_1/f_2 comparison from the motor response. Here we focus on responses during the first delay period. Neuronal responses in the fifth monkey during this first delay period were found to be indistinguishable from responses during the analogous delay period in the first four monkeys; data have therefore been collapsed across all five monkeys.

Several different stimulus sets were used, but the majority of recordings were carried out using the two sets illustrated in Figure 1c,d. To avoid variations in task difficulty, large differences between f2 and f1, compared to the monkeys' psychophysical threshold, were used in these stimulus sets. Stimulus set A was used with monkeys 1–3 (trained on task variant of Fig. 1a); stimulus set B was used with monkeys 4–5 (monkey 4 trained on task variant of Figure 1a, monkey 5 trained on task variant of Fig. 1b). Stimulus set B was specifically designed so that no f1 value allowed subjects to predict the appropriate motor response. Results across all stimulus sets were similar.

Single neuronal activity was recorded through seven independently movable microelectrodes (2–3 M Ω) (Romo *et al.*, 1999; Salinas *et al.*, 2000; Hernández *et al.*, 2002), separated from each other by 500 μ m and inserted in parallel into the inferior convexity of the PFC of the left hemisphere of four monkeys and of the right hemisphere in three monkeys. Results from both right and left PFCs were similar and data have been collapsed across them. Recordings sites changed from session to session. After advancing the electrodes into cortex, window discriminators were used to isolate one neuron per electrode while the monkeys performed the task. Typically, one or two neurons judged to be task-dependent, either because of responses that seemed correlated in time with task events or because of responses that seemed dependent on values of the vibrotactile stimuli, would be thus identified. Signals from the electrodes used to record these neurons would receive the focus of the experimenters' attention during subsequent recording of a full block of trials of a stimulus set (typically 10 repeats per f_1/f_2 pair of stimulus sets shown in Fig. 1c,d). After recording a full set, all seven electrodes would be advanced further into cortex to search for further neurons; 1162 neurons were recorded in this manner. For the analysis of Figure 6, signals from electrodes that had not been the focus of attention during recording were used and treated as if they came from single neurons, provided they passed all three of the following tests: (i) the fraction of interspike intervals < 1.5 ms was $< 2\%$; (ii) overall firing rate remained stable throughout the recording of the full set (trials were grouped into units of 10 consecutive trials each, average firing rate being deemed not different across units if an analysis of variance (ANOVA) test gave $P > 0.1$); and (iii) firing rate was never > 80 Hz and, when averaged over f1 stimuli and the entire delay period, not < 2 Hz. This led to a further 217 recordings being identified as individual units, for a grand total of 1379 recordings treated as originating from individual neurons. Of these, 557 (40%) were found to have a significantly monotonic f_1 -dependent firing rate during the delay period.

We used standard histological procedures to construct surface maps of all penetrations in the PFC. This was made first by marking the edges of the small chamber (7 mm diameter) placed above the PFC. Additionally, in the last recording sessions, we made small lesions at different depths. Electrode penetrations were normalized against the arcuate sulcus and principal sulcus. The penetrations that gave the results reported here were located in the inferior convexity of the PFC.

Data Analysis

Single-neuron spike trains were convolved with Gaussian kernels to obtain time-dependent spike density functions for each trial. Kernels were corrected for edge effects and stimulus periods densities were computed separately from densities of other periods; stimulus periods kernel $\sigma = 50$ ms, other periods $\sigma = 150$ ms: thus, the effective time scale of the delay period analysis was ≈ 300 – 600 ms. Edge correction was carried out by smoothing the firing rate function $f(t)$ for each period with edges at T_0 and T_1 using

$$s(t) = \frac{1}{Z(t)} \int_{(t-\tau=T_0)}^{(t-\tau=T_1)} f(t)h(t-\tau)d\tau$$

$$Z(t) = \int_{(t-\tau=T_0)}^{(t-\tau=T_1)} h(t-\tau)d\tau$$

where $s(t)$ is the smoothed function and $h(t)$ is the smoothing Gaussian. A time-dependent spike rate mean and a time-dependent spike rate standard error of the mean were computed from the set of density functions for each base stimulus condition. At each point in time, we computed the best linear fit of mean response as a function of base stimulus frequency; the significance of the slope of the linear fit being different from zero was then computed from the standard errors (Ross, 1987; Press *et al.*, 1992). We also computed the χ^2 -based goodness-of-fit probability Q of both a linear (2 degrees of freedom) and a sigmoidal (4 degrees of freedom) fit to the data (Press *et al.*, 1992). Times where the linear regression slope was significantly different from zero ($P < 0.01$) and either or both of the linear and sigmoidal fits were acceptable ($Q > 0.05$) were marked as ‘significantly monotonic’. Sigmoidal fits can replicate linear fits and, thus, the optimal sigmoidal fit always has a χ^2 value smaller than or equal to that of the optimal linear fit. But Q is a function of both χ^2 and the number of degrees of freedom and, for a given value of χ^2 , is higher for the fit with the smaller number of degrees of freedom (Press *et al.*, 1992). Significantly monotonic times were thus marked as either ‘linear’ or ‘sigmoidal’, according to which fit had the higher Q . A similar classification procedure was followed for firing rates averaged over the entire delay period (Fig. 3). To evaluate the net effect of our marking methods, we took 318 neurons that had at least one delay period moment marked as significantly monotonic and shuffled the base frequency labels among trials of each neuron. Upon reanalysis, only 13 of the shuffled neurons had one or more moments during the delay period marked as significantly monotonic. Thus, the net probability of marking a neuron as monotonic by chance was $\sim P = 13/318 \approx 0.04$.

Neurons were classified as ‘early’ if they had at least 100 ms of significant monotonic stimulus-dependence during the first second of the delay period and no significant stimulus-dependence during the last second of the delay period; as ‘persistent’ if they had at least 100 ms of significant monotonic stimulus-dependence during each of the 3 s of the delay period; and as ‘late’ if they had no significant stimulus-dependence during the first second of the delay period and at least 100 ms of significant stimulus-dependence during the last second of the delay period.

Results

Figure 2 illustrates the responses of a PFC neuron that has f1-dependent firing rates throughout the entire delay period. Responses with such a persistent stimulus-dependence are natural candidates for being considered the neuronal substrate of the memory of f1 in the vibrotactile discrimination task. Two features of the responses shown in Figure 2 are characteristic and representative of the neuronal population. First, the f1 \rightarrow firing rate mapping is monotonic, fairly smooth and centered at the midrange of f1 frequencies used (sigmoidal fit in Fig. 2c: $\beta = 0.11$ s/pulse, $\theta = 24$ pulses/s). Secondly, even though the stimulus-dependence of the firing rate persists throughout the delay period, the average firing rate of the neuron is not steady, but undergoes marked changes as the delay period unfolds.

Smoothness of Monotonic Encoding

The smooth monotonic encoding described here is different to two other types of delay period encoding that have been found. The first is similar to a labelled-line code: neurons that fire persistently during the delay period do so maximally for only one or a few of the possible cues and different neurons respond preferentially to different cues (Fuster, 1973; Funahashi *et al.*, 1989,

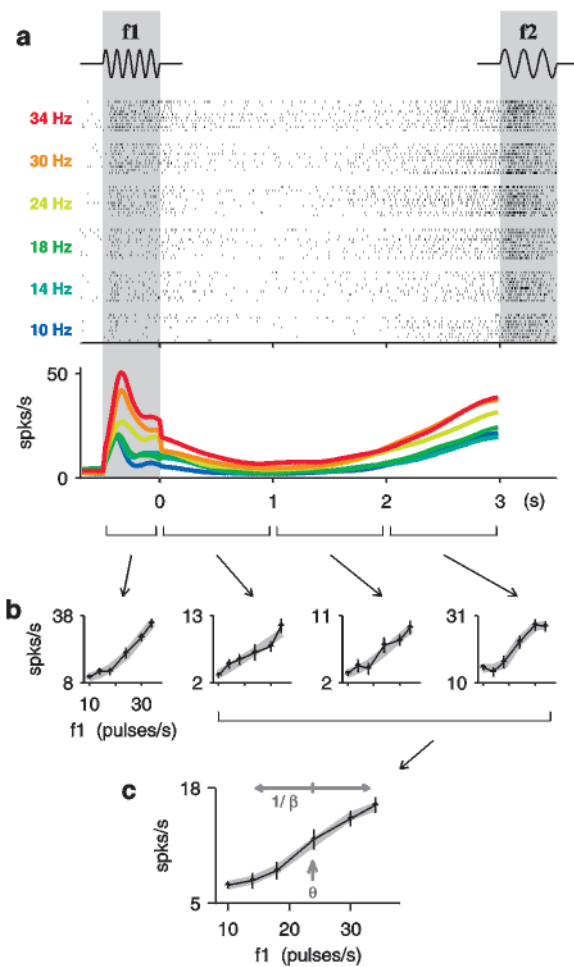


Figure 2. A persistent, positive monotonic neuron. (a) Rasters: each row of ticks represents a trial and each tick an action potential. Trials were presented in random order but have been sorted here into blocks of equal f1 frequency, indicated at left. Color code for f1 frequencies indicates corresponding smoothed peristimulus time histograms (PSTHs), shown below the rasters. (b) Firing rate, as a function of f1 frequency, averaged over different parts of the trial, indicated at the base of the arrows pointing to each panel. Notice the different y-axis scaling for each panel. Grey lines are sigmoidal fits to data in each panel (firing rate = $a * \tanh [\beta(f1 - \theta)] + c$); error bars, in this and in subsequent figures, are standard errors. (c) Firing rate as a function of f1 averaged over the entire delay period. $\beta = 0.11$ s/pulse, $\theta = 24$ pulses/s.

1990; Miller *et al.*, 1993, 1996; Fuster *et al.*, 2000). Each neuron may therefore be labelled as having a particular preferred cue and this label is a key component to decoding the responses of the neuronal population. The neurons may fire appreciably, if less than maximally, for a relatively broad range of cues, leading to such codes being often thought of in terms of Gaussian population coding (Deneve *et al.*, 2001). But because of the key role played by the neurons’ labels, we may think of this code as similar to a labelled-line code. Note that in the well-studied task of remembering one of a set of spatial positions arranged in a ring (Gnadt and Andersen, 1988; Funahashi *et al.*, 1989, 1990), there is no well-defined ordinal relationship between positions in the ring and thus we would not expect monotonic encoding. Another type of encoding with which monotonic encoding should be contrasted is temporal encoding. In this type of code, it is the temporal properties of the neural activity (e.g. the pres-

ence or absence of oscillations within a particular frequency band), rather than firing rate *per se*, that encode the memory of the stimulus (Pesaran *et al.*, 2002). Temporal encoding has been observed coexisting with labelled-line firing rate encoding (Pesaran *et al.*, 2002). Whether one causes the other is not yet clear.

We selected those neurons that were classified as persistent [see Materials and Methods; $n = 126$ out of 557 (23%) delay period stimulus-dependent neurons] and fit their firing rates, averaged over the whole delay period, as either linear (2 degrees of freedom) or sigmoidal (4 degrees of freedom) functions of $f1$. Taking into account both the number of degrees of freedom and the quality of the fit, neurons were classified as either linear or sigmoidal (see Materials and Methods). Fifty-four (43% of 126) persistent neurons were classified as linear. By definition, the $f1 \rightarrow$ firing rate function of linear neurons is smooth. Figure 3*a* shows a histogram of the slopes found for these neurons. The absolute value of the slope ranged from 0.15 to 0.9 spikes/s per pulses/s, with a median value of 0.32. Thus, for a typical linear persistent neuron, a change in $f1$ from 10 to 34 pulses/s (full range used in the experiments) led to a change in the delay period firing rate of 8 spikes/s. Overall median firing rate during

the delay period, compiled over all stimuli and all linear persistent neurons, was 16 spikes/s.

Seventy-two (57% of 126) persistent neurons were classified as having a sigmoidal delay period tuning. The delay period firing rates of these neurons were fit with a function of the form $a \cdot \tanh[\beta(f1 - \theta)] + c$. The hyperbolic tangent function, $\tanh(x)$, is roughly linear for x within $[-1, 1]$ and begins to saturate at higher or lower values. Thus, high values of β indicate a function that approximates a sharp (non-smooth) step function; conversely, very low values of β approximate smooth, linear tuning. The θ parameter indicates the $f1$ centerpoint of the sigmoidal function. Figure 3*b* shows a scatterplot of β versus θ fit values, with one data point per sigmoidal persistent neuron. The centerpoints θ cluster around the midpoint of the $f1$ range used (median value of $\theta = 23$ pulses/s). The values of β were close to zero, with a median absolute value of 0.15. Thus, the typical sigmoidal persistent neuron had fairly smooth tuning, being linear over a range of ~ 13 pulses/s ($\approx 2/0.15$), an appreciable part of the range used in the experiments. The median firing rate change for sigmoidal persistent neurons, for $f1$ changing from 10 to 34 pulses/s, was 10 spikes/s. The overall median firing rate of these neurons was 18 spikes/s.

Some outliers from the distribution shown in Figure 3*b* were found. Values of θ far from 22 generally corresponded to sigmoidal fits similar to that shown in the inset of the upper left of Figure 3*b*, for which the slightly curved, but nevertheless smooth, shape is fit better by the corner of a soft sigmoid than by a straight line. Eleven neurons had a θ value outside the range shown in the graph. All of these 11 neurons had very small β values ($0.02 < |\beta| < 0.11$), indicating a high degree of smoothness. Nine neurons had sigmoidal fits with a value of $|\beta| > 1$. These were fits similar to that shown in the inset at the top right of Figure 3*b*; although the best-fitting sigmoid was a sharp step function, for all nine of these neurons the step function had a data point located at the midpoint of its high slope portion, indicating that large changes in firing rate required a change of more than the 4 pulses/sc that typically separated values of $f1$ used (see Fig. 1*c,d*).

For large values of $|\beta|$, the centerpoint α represents the value of $f1$ at which the neuron's response changes abruptly from a low to a high firing rate. An alternative to the smooth encoding we have focused on here would be to represent $f1$ by the number of neurons that are firing above such an abrupt threshold. However, for small values of $|\beta|$, even though the fit's centerpoint θ remains well-defined, there is no clear threshold value at which the firing rate changes abruptly (see, for example, the insets at top left, lower left and lower right in Fig. 3*b*). Threshold-based encoding is thus not well-defined for small $|\beta|$. The inset at the lower center of Figure 3*b* shows a relatively smooth fit which has a β value of -0.25 . We conservatively took 0.25 as the minimum $|\beta|$ for which a threshold-based encoding could be considered. Including all outliers, only 22 neurons (31% of 77 sigmoidal neurons, only 17% of all 126 persistent neurons) had $|\beta| > 0.25$, indicating that a strong majority of neurons encoded $f1$ in a smooth monotonic fashion. We touch again on this issue in the Choice Probability and Discussion sections below.

We re-performed the smoothness analysis on the firing rates of persistent neurons, but now examining separately each of the three non-overlapping 1 s long segments of time within the delay period (e.g. the three right-hand panels of Fig. 2*b*). Distributions and results of linear and sigmoidal fits of firing rates as a

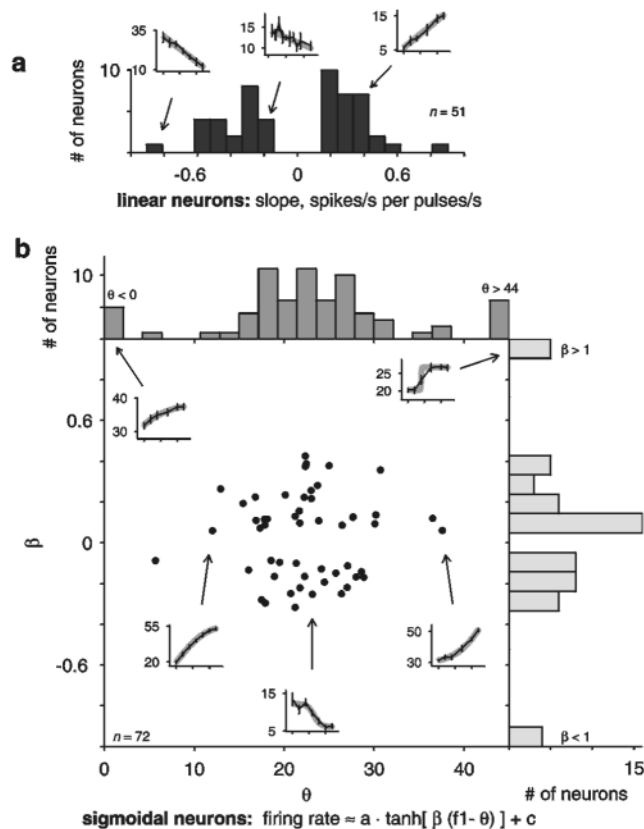


Figure 3. Persistent neurons: linear and sigmoidal fits of firing rate, averaged over the entire delay period, as a function of $f1$. (*a*) Histogram of slopes of neurons classified as linear. Insets show examples of fits for individual neurons. Tickmarks on x -axis of insets indicate 10, 20 and 30 pulses/s of stimulus $f1$; y -axis units are spikes/s and thick grey lines are linear fits. (*b*) Scatterplot of slopes (β) and midpoints (θ) of neurons classified as sigmoidal. Each data point represents one neuron. Insets follow same conventions as in (*a*), but grey lines are sigmoidal fits. Histograms of values of θ (above) and β (right) show, in the extreme bins, number of outliers not shown in the central scatterplot.

function of f_1 , for the individual 1 s long segments, were similar to those shown in Figure 3 (data not shown). Thus, the overall picture that emerged was that the monotonic stimulus-dependence of delay period firing rates of persistent neurons was a smooth function of f_1 , with many of the $f_1 \rightarrow$ firing rate mappings best fit by straight lines and with sigmoidal fits having shallow slopes and centerpoints clustered near 22 pulses/s.

Lack of Non-monotonic Encoding

Was monotonic encoding the only type of f_1 -dependence displayed by PFC neuron firing rates during the delay period? To search for f_1 stimulus-dependent firing rates that were not necessarily monotonic functions of f_1 during the delay period, we grouped trials according to the value of f_1 and tested whether there was any significant difference between responses to different f_1 values by using a one-way ANOVA test ($P < 0.05$). The ANOVA test does not assume monotonicity or any other type of functional form. We took records of neurons studied with 3 s long delay periods and divided the delay period into three non-overlapping seconds (three rightmost panels of Fig. 2*b*). For each of these 3 s, we asked first whether firing rates were a monotonic function of f_1 (see Materials and Methods) and, secondly, whether the ANOVA test indicated an f_1 -dependence. Records were marked as monotonic or non-monotonic according to the result of the monotonic test ($P < 0.05$). Records marked as non-monotonic were further marked as significantly f_1 -dependent, in a non-monotonic manner, if the ANOVA test was positive ($P < 0.05$). Out of all non-monotonic records, only 3.5% were marked as f_1 -dependent by the ANOVA test. At $P < 0.05$, this is comparable to the number that would be expected purely by chance. We thus concluded that, within the resolution afforded us by the finite number of f_1 values used and the finite number of trials recorded per condition (typically 10 trials per f_1/f_2 pair), there was no significantly detectable non-monotonic encoding of f_1 in the firing rates of prefrontal cortex neurons.

Time-varying f_1 -dependence

Neurons studied with a 3 s long delay period were classified into different groups according to the portion of the delay period during which their firing rates were significantly f_1 stimulus-dependent (Fig. 4; see Materials and Methods). This classification is based on stimulus-dependence, not on overall temporal trends in firing rate. To emphasize this point, the neurons chosen for illustration in Figure 4*a–c* were all specifically chosen as neurons that, after averaging across f_1 stimuli, show no significant temporal trends in their delay period firing rate. Monkeys 1–4 performed the task as shown in Figure 1*a*, with no enforced delay between the end of stimulus f_2 and the start of the motor response. From these four monkeys, 345 neurons that were stimulus-dependent during the delay period were studied with a 3 s long delay period. Of these, 29% ($n = 101$; Fig. 4*a*) were classified as late, 26% ($n = 90$; Fig. 4*b*) were classified as persistent, and 35% ($n = 121$; Fig. 4*c*) were classified as early. The remainder could not be classified within any of these three broad groups.

Monkey 5 was studied while it performed the variant of the task shown in Figure 1*b*, in which there is an enforced, 3 s long separation between the end of stimulus f_2 and the initiation of the motor response. For monkey 5, 189 stimulus-dependent neurons were studied with a 3 s long delay period, of which $n = 75$ (40%), $n = 36$ (19%), and $n = 46$ (24%) were classified as late,

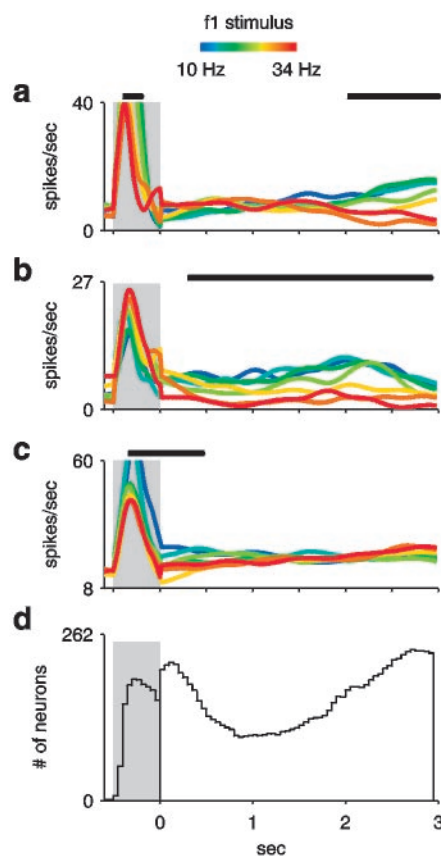


Figure 4. Different neurons are f_1 -dependent at different times of the delay period. Panels (a–c) show PSTHs, color-coded according to the frequency of the f_1 stimulus applied, for the f_1 stimulus period (grey box) and the delay period. Thick horizontal black bars indicate periods of time when the neuron’s firing rate was significantly f_1 -dependent (see Materials and Methods). (a) A late neuron. (b) A persistent neuron. (c) An early neuron. (d) Total number of recorded neurons (during fixed 3 s delay period runs) carrying a significant signal about the base stimulus, as a function of time relative to the beginning of the delay period. Individual neurons may participate in more than one bin.

persistent and early, respectively. Thus, the fraction of neurons classified as late was as high in monkey 5 as in the other four monkeys. This shows that the presence and timing of late neuron responses is not dependent on the close temporal proximity of a motor act or reward. Both early and late neurons had coding properties (proportion of neurons classified as linear, tendency for sigmoidal fits to be smooth) similar to those shown in Figure 3 for persistent neurons.

Choice Probability Analysis

If the monotonic firing rate encoding we have described is the neuronal basis for representing the short-term memory of f_1 during the task, then trial-by-trial variations in firing rates of monotonic neurons should be accompanied by correlated trial-by-trial variations in performance. For example, if on a particular $f_1 = 22$ pulses/s trial a positive monotonic neuron fired more strongly than average, we would expect that on that trial the monkey would remember f_1 as having been slightly higher than 22 pulses/s. This effect on memory would be revealed by a bias, on such trials, for the monkey to respond by pressing the ‘ $f_1 > f_2$ ’ button rather than the ‘ $f_1 < f_2$ ’ button.

We quantified the correlation between firing rate and behavior by taking trials for each particular f1/f2 stimulus pair and dividing them into two sets, those in which the monkey finished by pressing the 'f1 > f2' button and those in which the monkey finished by pressing the 'f1 < f2' button. We then asked, for each neuron, whether the neuronal responses of the two sets of trials were similar or different from each other. With the f1 and f2 stimuli held fixed, it is only the monkey's pushbutton choice that distinguishes the two sets of trials from each other. (One of the sets will be composed of error trials, while the other will be composed of correct trials.) When the difference between the two sets of responses is evaluated by computing a receiver operating characteristic (ROC) curve area, this type of analysis is termed 'choice probability analysis' (Green and Swets, 1966; Britten *et al.*, 1996). Here, a choice probability >0.5 indicates that the sign of the separation between the two sets of responses is the same as predicted by whether, based on correct trials only, the neuron was positive or negative monotonic (as in the f1 = 22 pulses/s example of the paragraph above). A choice probability of 1 (its maximum value) further indicates that the distributions of the two sets of responses are completely non-overlapping; a choice probability of 0.5 indicates that the distributions of the two sets of responses overlap very closely; and a choice probability <0.5 indicates that the sign of the difference between the two distributions is opposite to that predicted by their tuning.

We performed this analysis for each individual f1/f2 stimulus pair for each neuron and then averaged across f1/f2 pairs and across each of the neuronal populations of early, persistent and late neurons. The results are shown in Figure 5. The average choice probability (ACP) of persistent neurons is significantly higher than 0.5, even during the f1 stimulus period, grows as the delay period unfolds and is higher than that of any other type of neuron throughout the entire delay period. The ACP of early neurons is significantly higher than 0.5 only during the beginning of the delay period, while the ACP of late neurons becomes significantly higher than 0.5 only towards the end of the delay

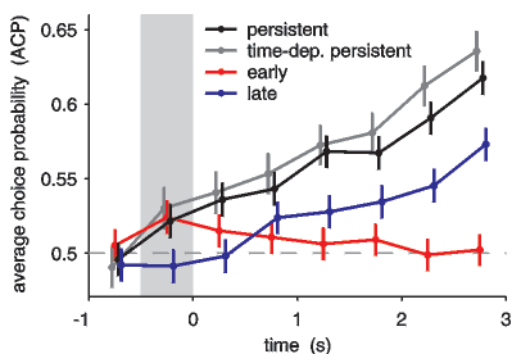


Figure 5. Choice probability grows during the delay period and is strongest for persistent neurons. Each data point represents an average over the indicated neuronal population and is computed from spikes fired in a 500 ms-wide window centered at the data point. Different data points are separated by 500 ms and are thus taken from non-overlapping temporal windows. Vertical bars denote standard errors. Different lines have been staggered slightly in time for visibility. Time-dependent persistent neurons (grey) are a subset of the persistent neurons (black): they are the 67% of persistent neurons that have a significantly time-dependent firing rate during the two central seconds of the delay period (see Fig. 6 and accompanying text).

period. This is consistent with the separately defined temporal properties of each group's f1-dependence.

If the encoding of f1 was not smooth, but preferentially based on sigmoidal neurons with a high value of $|\beta|$ (see 'Smooth Monotonic Encoding' section above), we might expect that sigmoidal neurons with high $|\beta|$ would be the neurons most closely linked to behavior and would therefore tend to have the highest ACP values. However, both linear neurons (which have the smoothest possible encoding) and neurons with $|\beta| < 0.25$ had ACP values statistically indistinguishable from those of neurons with $|\beta| > 0.25$ (data not shown). This does not support a preferential role for high $|\beta|$ sigmoidal neurons in encoding the memory of f1.

One of the differences between the analysis of Figure 5 and previous choice probability analysis of PFC neuron responses during delayed visuomotor tasks (Kim and Shadlen, 1999; Constantinidis *et al.*, 2001) is that here the analysis is performed on responses from a delay period that occurs before the monkey can decide on its motor response to the trial. The analysis of Figure 5 therefore more closely reflects a purely memory-based correlation of firing rates with behavior, as opposed to a potential mix of memory and decision effects (Constantinidis *et al.*, 2001). When stimulus f2 is finally presented and the monkey makes its choice, ACP values of responses in both PFC and in secondary somatosensory cortex rise sharply, reaching values as high as 0.81, much higher than any value shown in Figure 5 (Romo *et al.*, 2002; Romo *et al.*, unpublished results).

Time-dependent Average Firing Rates

To study overall temporal trends during the delay period, we averaged firing rates across f1 stimuli and computed for each neuron the firing rate during each of the three non-overlapping seconds of the delay period (three right-hand panels of Fig. 2b). We then tested for differences in firing rate across these three successive seconds. A very large proportion of all neurons thus tested (560 of 777, 72%) were found to have a significant time-dependence in their firing rates during the delay period (one-way ANOVA, significance limit a highly conservative $P < 0.001$ for each neuron; P -value confirmed in a shuffle test that randomly permuted data between the three different seconds). Of the 777 neurons tested, 217 were units recorded from electrodes that during the experiments had been not been judged to carry signals related to the task, either in the timing or in the stimulus-dependence of their responses [see Materials and Methods; the 217 apparently not task-related neurons should not be confused with a different, overlapping set that coincidentally has the same number of members, namely the 217 neurons (777 - 560) that did not have significantly time-dependent responses]. In contrast to most of the recorded neurons, the apparently non-task-related subset of the population had an explicit selection bias against having f1-dependent firing rates. Yet a high proportion of neurons of this subpopulation also had significantly time-dependent firing rates (112/217, 52%). This suggests that firing rates that vary systematically with time as the delay period unfolds are very common in the PFC. This also shows that many of the 217 neurons that had no f1-dependence nevertheless did have responses related to the task, albeit to timing aspects of the task.

A large variety of time-dependent patterns in firing rate were found. To enable us to visualize this variety, we computed two numbers for each neuron: (i) the difference in average firing rates between the middle and the first second of the delay

period ('do firing rates start the delay period by rising or falling?'); and (ii) the difference in average firing rates between the last and the middle second of the delay period ('do firing rates end the delay period by rising or falling?'). To facilitate comparison across neurons, we then normalized both of the computed numbers by the neuron's firing rate, averaged over the entirety of the delay period. The two normalized numbers are plotted against each other, for the 560 significantly time-dependent neurons, in Figure 6. Insets in this figure show examples of peristimulus time histograms (PSTHs) for individual neurons. We emphasize that the firing rates and PSTHs in this figure are computed averaged over f1 stimuli: here we are focusing purely on time-dependent, not f1-dependent, effects. For clarity, the 217 non-significantly time-dependent neurons are not shown; they would all cluster near the origin of the graph.

A plurality of neurons had data points in the top right quadrant of Figure 6 (236 neurons, 42%). Neurons corresponding to data points in this quadrant have firing rates that are rising during both the start and the end the delay period. The quadrant with the next highest number of data points (171 neurons, 31%) is the lower left quadrant, where neurons have monotonically decaying firing rates. Note that in both the upper right and lower left quadrants there are many data points far from the

cardinal axes. For such neurons, the decay (or rise) in firing rates occurs from the beginning through to the end of the delay period. The upper left quadrant, in which a substantial number of neurons (119, 21%) have their data points, indicates falling followed by rising firing rates. Finally, the odd quadrant out is the lower right quadrant, in which a very small number of neurons (34, 6%) had data points. Neurons with firing rates that at first rise and then fall during the delay period are thus rare. If the data points fell randomly in the different quadrants, 140 ± 10 (mean \pm SD) neurons per quadrant would be expected; since all differences in number of neurons between the four quadrants are much greater than 10, the differences are thus highly significant (confirmed by permutation test, $P < 0.001$).

Of the 126 persistently stimulus-dependent neurons, a remarkably high proportion (113, 90%) had significantly time-varying firing rates ($P < 0.001$ for each neuron). This fraction is higher even than the fraction in the general population (which was 72%, as indicated above). We speculated that the origin of this result with persistent neurons could lie in temporal variations occurring at the very beginning of the delay period, when the memory signal may be in the course of being established, or at the very end of the delay period, when anticipation of the end of the delay period may lead to changes in firing rates. We therefore repeated the analysis of time-dependency of persistent neurons, but after having excluded both the first 500 and the last 500 ms of the delay period from consideration. This analysis, focusing on the central 2 s of the delay period, still resulted in a very high fraction of the persistent neurons (84 of 126, 67%) being marked as having significantly time-dependent firing rates (again, $P < 0.001$ for each neuron). These results indicate that neurons with steady persistent firing rates, such as the example neuron shown in Figure 4*b*, are markedly in the minority. Instead, persistent neurons with time-dependent changes in their firing rates, such as those shown in Figures 2 or 8*i*, are much more common.

Effect of Changing the Delay Period Length

Are time-dependent firing patterns in the PFC dependent on the length of the delay period, or are they a response due to a fixed dynamics, having a fixed time scale that is independent of the delay period length? Following Kojima and Goldman-Rakic (1982), we had the opportunity, for a small number of neurons ($n = 51$), of recording from them during a block of trials with a delay period of 6 s, after having recording from them during a standard block of trials with the 3 s long delay period. The monkeys had previously trained on both 3 and 6 s delay periods (although 3 s delay periods were predominantly used). All sets of trials were done in blocks. Thus, one or a few trials with a particular delay period length were sufficient to allow the monkeys to know which delay period length the next trial would most likely have.

Figure 7*a* shows the spike trains of a late neuron during both the last trials of the 3 s delay period block and the first trials of the 6 s delay period block. During the first few trials of the 6 s block, the neuron responded with a timing similar to its timing during the 3 s block, firing vigorously slightly before the 3 s mark. But as the trials progressed, the response of the neuron adjusted, moving its firing to a much later part of the delay period. Most late neurons in which a similar change in response profile was visible on a trial-to-trial basis adjusted much more rapidly, usually within one or two trials, such as the neuron with




Figure 6. Time-dependent firing rate patterns during the delay period. Each data point represents one neuron. The position along the the x-axis indicates whether firing rates (averaged over all f1-stimuli) were rising (right half of x-axis) or falling (left half) between the first and the middle seconds of the 3 s delay period. Position along the y-axis indicates rising (top half) or falling (bottom half) firing rates between the middle and final seconds of the delay period. Only neurons with significant temporal modulation are shown. Numbers at the corner of each quadrant indicate number of data points in that quadrant. Insets are PSTHs, averaged over all f1 stimuli, for individual example neurons; corresponding data points are indicated by arrows. Grey box in PSTHs indicates 500 ms f1 stimulus period. Inset y-axes have units of spikes/s and the bottom is 0 spikes/s in all insets. Tickmarks along inset x-axes indicate seconds. Red, black and blue dots indicate early, persistent and late neurons, respectively (see full legend inside graph).




Figure 7. Responses of two different neurons, recorded during a block of 3 s delay period trials and during an immediately subsequent block of 6 s delay period trials. Upper panels show the last 40 trials of the 3 s block, in the sequence in which they were presented. Lower panels show the first 40 trials of the 6 s delay period block.

responses illustrated in Figure 7*b*. There was a small but significant tendency for the monkeys to make more mistakes than usual on the first trial of the 6 s blocks (percentage correct on first trials of 6 s blocks, 79%; on second trial, 95%; on all trials subsequent to the first, 93%; $P < 0.05$ in a permutation test for first trial having lower percentage correct than subsequent trials). No such tendency was found on the first trial of 3 s blocks (percentage correct on first trials of 3 s blocks, 92%; on all trials subsequent to the first, 94%; $P > 0.45$ for first trial having lower percentage correct than subsequent trials).

Do late neurons adjust to the 6 s delay period block by simply shifting their f1-dependent phase to later in the delay period? Or do they adjust by scaling time, with their response characteristics evolving half as fast during 6 s delay period trials as they evolved during blocks of 3 s delay period trials? Figure 8*a* shows PSTHs, color-coded according to f1 stimulus value, for a late neuron recorded during both a 3 s delay period block (upper panel of Fig. 8*a*) and a subsequent 6 s delay period block (lower panel of Fig. 8*a*). The two panels have very different time axes, one being scaled by a factor of two with respect to the other. Yet, after scaling and without any shifts, the two graphs are remarkably similar. Six other neurons classified as late neurons (based on data from 3 s delay period blocks) were recorded from during subsequent 6 s delay period blocks. PSTHs for these neurons are shown in Figure 8*b–g*. In each of these panels, the graphs of the two conditions look similar to each other after scaling time by a factor of two. Allowing both shifting and scaling, we asked what combination of these two transformations of the time axis would make the 3 s block delay period PSTHs best match the 6 s block PSTHs, in a squared error sense. That is, for each f1 frequency, we computed

$$E_{f1}(\omega, s) = \frac{1}{T} \int [PSTH_6(t) - PSTH_3(\omega t + s)]^2 dt$$

where ω is a dimensionless scaling factor, s is a temporal shift with units of seconds and T represents the portion of the 6 s delay period that overlaps with the 3 s delay period after scaling and shifting. We then averaged over f1 frequencies, to obtain a total squared error measure $E(\omega, s) = \langle E_{f1}(\omega, s) \rangle_{f1}$. The values of




Figure 8. PSTHs of individual late and persistent neurons recorded during both a block of trials with a 3 s delay period (upper panel of each pair of panels) and a subsequent block of 6 s delay period trials (lower panel in each pair). Value of f1 stimulus indicated by color scale at upper right. Each pair of panels has the optimal scaling and shifting of the time axis of the 3 s PSTH that produces the smallest squared error match with the 6 s PSTH (see text). (*a–g*) Seven different late neurons. (*h–i*) Two different persistent neurons.

ω and s that minimized E were defined as the optimal warp and shift.

The optimal warp and shift for each neuron are indicated in the corresponding panels of Figure 8*a–g*. Time shifts are generally small ($|s| < 0.31$ s), while most values of ω are close to one-half ($|\omega - 0.5| < 0.1$). This confirms that the transformation of these late neuron responses from the 3 s delay period condition to the 6 s delay period condition is characterized mostly by a time stretch (stretch factor ~ 2), with time shifts being relatively

unimportant. For persistent neurons with firing rates that remain relatively constant during the delay period (Fig. 8*h*), there is no meaningful optimal stretch or shift. But an optimal stretch and shift can be defined for persistent neurons with time-varying firing rates; an example of such a persistent neuron is shown in Figure 8*i*. Once again, the shift is small (100 ms) and the transform is dominated by stretching.

Many of the late neurons in Figure 8 have low or constant firing rates during the first third of the delay period. The analysis is therefore mostly dependent on dynamics of responses during the final, rather than the earlier, part of the delay period. In contrast to late neurons, early neurons generally have strong dynamics in their responses during the beginning of the delay period (Fig. 6). Five neurons classified as early, based on data from 3 s delay period blocks, were recorded during subsequent 6 s delay period blocks. These recordings allowed us to approach the question of what effect changing the delay period length might have on the early, rather than the late, parts of the delay period. Unlike the late neurons of Figure 8, some of these five early neurons changed their maximum firing rates markedly between the two delay period conditions. Some of them also changed the overall shape of their PSTHs. For these early neurons, matching PSTHs from the two conditions was thus a less straightforward proposition than for the late neurons. We addressed this problem in two ways. First, we focused on early parts of the response by requiring the match to be based only on the first second of the 3 s delay period responses; PSTH shapes were roughly similar in the two conditions for this part of the response. Secondly, we focused only on the PSTH shapes, not their maxima or minima, by allowing shifting and stretching of the firing rate axis, in addition to the time axis when performing the match. Figure 9 shows overlaid PSTHs, averaged over f1

stimuli, from the 6 s delay period condition and, after stretching and shifting, from the 3 s delay period condition. Good matches in the shapes of the early parts of the responses are obtained. Similarly to the late neurons of Figure 8, shifts in the time axis are small ($|s| < 0.2$ s). In contrast to the late neurons, however, the best shape matches were obtained with time-stretch factors ω that are close to 1, instead of close to a half ($|\omega - 1| < 0.17$).

Discussion

We have used a task in which monkeys must remember a scalar value (f_1 , the frequency of the first vibrotactile stimulus) over a delay period lasting several seconds and must then make an ordinal comparison of f_1 to the frequency of a second stimulus (f_2). We found neurons in the inferior convexity of the prefrontal cortex with firing rates during the delay period that depend monotonically on f_1 . We refer to this as a monotonic encoding of f_1 . No significant number of neurons with a statistically significant non-monotonic f_1 -dependence were found.

Different neurons were found to have f_1 -dependent firing rates during different portions of the delay (Fig. 4). Most neurons with f_1 -dependent delay period firing rates could be described as belonging to one of three types: ‘early’ neurons (31%), which had f_1 -dependent firing rates during the early, but not the late, part of the delay period; ‘persistent’ neurons (24%), which had f_1 -dependent firing rates throughout the delay period; and ‘late’ neurons (33%), which had f_1 -dependent firing rates during the late, but not the early, part of the delay period. This classification is not intended to describe three widely separate groups, but is merely an aid in characterizing the variety of neurons found. Neurons that had f_1 -dependence only during the middle portion of the delay period were rare; neurons that had f_1 -dependence only during the beginning and end, but not the middle, of the delay period, were also rare.

Neurons with f_1 -dependent delay period activity are also found in brain areas other than the PFC. In particular, secondary somatosensory cortex (area S2; Salinas *et al.*, 2000) contains neurons that have early delay period f_1 -dependence. No persistent neurons are found in area S2 (Romo *et al.*, 2002). Early, persistent and late neurons are all found in medial premotor cortex (MPC), although late neurons predominate in the MPC (Hernández *et al.*, 2002). The similarity between MPC and PFC during the delay period of the vibrotactile task may be analogous to the similarity between lateral intraparietal cortex (LIP) and PFC in a delayed saccade visuomotor task (Chafee and Goldman-Rakic, 1998), although in the vibrotactile task persistent neurons are more easily found in PFC than in MPC. In contrast to all of these areas, no delay period f_1 -dependent activity was found in S1, the primary somatosensory cortex (Salinas *et al.*, 2000).

Supporting a role for monotonic encoding as the neural substrate of working memory in the vibrotactile discrimination task, the firing rates of monotonic neurons were found to correlate with the monkeys’ behavior in a manner predicted by their stimulus-dependence (Fig. 5). That is, for positive monotonic neurons, trials with above-average firing rates resulted in a bias in favor of monkeys pressing the ‘ $f_1 > f_2$ ’ button rather than the ‘ $f_1 < f_2$ ’ button, consistent with the monkeys remembering f_1 as higher than it actually was. The converse was true for negative monotonic neurons. The portions of the delay period during which this correlation was detected depended on the portions of the delay period during which the neurons monotonically

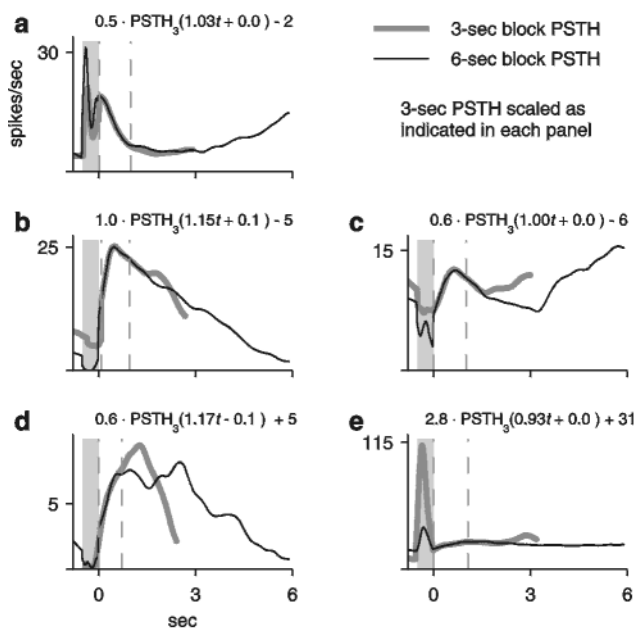


Figure 9. PSTHs of five different early cells recorded during both a block of trials with a 3 s delay period and a subsequent block of trials with a 6 s delay period. The PSTHs have been averaged across f_1 stimuli and the 3 s delay period PSTHs have been scaled and shifted in both time and firing rate, by factors indicated in each panel. After scaling and shifting, the 3 s delay period PSTHs (thick grey lines) are overlaid onto the 6 s delay period PSTHs (thin black lines).

encoded f1: for early neurons, the correlation was present only during the early part of the delay period; for late neurons only at the end of the delay period; and for persistent neurons, it was present throughout. These results are unlikely to be due to modulations in arousal or attentional state during the delay period: such modulations would lead to opposite results for trials in which $f2 > f1$ compared to trials in which $f2 < f1$, and no such difference between these two groups of trials was observed (comparison data not shown).

The gradual increase in choice probability seen in Figure 5 is consistent with a memory representation of f1 that is gradually degrading, with noise that accumulates during the delay period. In such a scenario, the later moments of the delay period would have the representation of f1 most highly correlated with the representation actually present when the memory is required for the task (that is, during presentation of f2, immediately after the delay period). The later moments of the delay period would thus have the neuronal responses most highly correlated with the monkey's choice, while neuronal responses from the early part of the delay period, while still correlated with the monkey's choice, would be less strongly correlated, as seen in Figure 5.

The monotonic encoding described here is smooth, in the sense that changes in f1 of a magnitude detectable by the monkey (2–4 Hz for the range of f1 stimuli used; Hernández *et al.*, 1997) are encoded by gradual changes in the firing rates of monotonic neurons. Of persistent neurons, 43% had firing rates that were a linear function of f1 (Fig. 3*a*); the remaining 57% had firing rates that were well fit by sigmoidal functions of f1 (firing rate $a * \tanh[\beta(f1 - \theta)] + c$; Fig. 3*b*). Sixty-nine per cent of the sigmoidal neurons were fit by very gradual, smooth functions of f1 ($|\beta| < 0.25$). Since there is no single preferred f1 stimulus in a smooth monotonic code, it is fundamentally different from a code based on labelled lines.

Do our data support an f1 encoding scheme different to smooth monotonic encoding? One possibility would be to encode f1 through the identity of sigmoidal neurons firing above their centerpoints θ —the distribution of centerpoints θ for the population of sigmoidal neurons was found to be broad enough to span most of the range of f1 values used in the experiment. However, only 17% of all persistent neurons were fit by sigmoids with $|\beta| > 0.25$. This is a conservative minimum value of $|\beta|$ for which the sigmoid is sharp enough that θ , the f1 centerpoint of the fit, can be reasonably interpreted as a clear 'threshold' f1 value, dividing low from high firing rate responses. It seems unlikely that f1 would be preferentially encoded by such a small fraction of neurons. Furthermore, such an encoding would predict that high $|\beta|$ neurons would have a higher correlation with behavior than other neurons, but no such difference was found. Alternatively, a sharp threshold could be part of the decoding procedure rather than the encoding firing rates. For example, f1 could be encoded by the identity of neurons firing above 16 spikes/s, regardless of whether these neurons are linear, smooth sigmoidal neurons, or otherwise. We cannot conclusively rule out such a possibility. However, this decoding threshold-based scenario would suggest that, for linear neurons, high choice probability values would be found for only one or two f1 values (those bordering on the decoding threshold), while other f1 values would be less tightly linked with behavior and would have lower choice probability values. No such difference was found, once again suggesting smooth monotonic encoding as a more likely option.

Our results on monotonic encoding (Romo *et al.*, 1999; present paper) apparently contrast with a recent study of short-term memory encoding of numerosity (Nieder *et al.*, 2002; Nieder and Miller, 2003), in which Gaussian-shaped tuning curves for numerosity were found in the PFC, suggesting a labelled-line based population code. However, numerosity was explicitly chosen for that study as being highly abstracted from any sensory representation and composed of well-defined discrete steps (the natural numbers), in contrast to the sensory continuum of vibrotactile frequency used here. The two tasks are thus quite different; this and other differences may explain the contrasting results.

If the late f1-dependence of 'late' neurons were associated with the anticipation of an immediately upcoming f2 stimulus, it might have been expected that the onset of their stimulus-dependence would have a fixed timing relative to the end of the delay period. For example, one might expect a late neuron to start its f1-dependent response 1 s before the end of the delay period, regardless of whether the trial is part of a block of trials with 3 s delay periods or a block with 6 s delay periods. However, we found that late neurons did not shift their response together with the end of the delay period, but instead scaled the timing of their response: their response characteristics evolved half as fast during blocks of 6 s delay period trials as they did during blocks of 3 s delay period trials [Fig. 8; see also Komura *et al.* (2001)] for a similar response scaling in thalamus]. If time-varying neuronal responses are used to represent time itself, the result is consistent with 'scalar timing', the equivalent of Weber's law in the time domain (Allan, 1979). The scaling result also suggests that long delay periods cannot be used to temporally isolate persistent neuron f1-dependence from late neuron f1-dependence. Late neurons may thus be as intrinsic a part of a time-dependent memory representation as persistent neurons. It must be emphasized, however, that our results on delay period lengthening are based on an as yet small number of neurons. While strongly suggestive, they can therefore not yet be taken as fully conclusive.

Starting with the first reports of recordings from prefrontal cortex during working memory tasks (Walter, 1964; Fuster, 1973), it has been known that PFC neuron firing rates often change systematically over the course of delay periods. Separately from temporal effects, the dependency of PFC neuron firing rates on stimuli being held in memory has been considered a very attractive candidate for the neural substrate of working memory. As such, it is this stimulus-dependent aspect, rather than temporal dependency, that has received the greater focus of attention. Current models of persistent activity as a basis for short-term memory assume that the memory is held in a pattern of firing rates in which each neuron's firing rate is essentially constant throughout the delay period, except for undesired and minor drifts due to noise (for review, see Wang, 2001; see also Miller *et al.*, 2003).

Nevertheless, our experimental observation is that only a very small number of neurons have firing rates that are both f1-dependent and constant during the delay period. Persistent neurons are thought to be fundamental for carrying the memory signal throughout the delay period. Yet an overwhelming fraction of persistent neurons (90%) have firing rates that vary systematically with time during the delay period. Restricting the analysis to a central portion of the delay period necessarily reduces the number of neurons detected as significantly time-dependent (to see this, note that in the limit in which the central

portion is reduced to a single instant, there can be no time-dependence). But when the analysis is restricted to the central 2/3 of the delay period, in order to exclude any edge effects associated with the beginning or end of the delay period, a full 67% of persistent neurons are still found to have highly significantly ($P < 0.001$) time-dependent firing rates.

For 10% ($n = 12$) of the persistent neurons, the time-dependency analysis indicated a P -value of >0.1 in favor of the null hypothesis of static firing rates (if the analysis is restricted to the central 2 s of the delay period; when the full delay period was considered, only 3%, $n = 4$, neurons had $P > 0.1$). This subset of neurons had firing rates that were roughly constant during the delay period. It may be that such a small fraction of f1-dependent neurons is fully responsible for the maintenance of the memory signal and that other neurons merely reflect the activity of these key neurons. However, analysis of error trials showed that time-dependent persistent neurons are as closely correlated with behavior as time-independent persistent neurons (Fig. 5). This suggests that both time-dependent and -independent persistent neurons may play a similar, and equally important, role in determining memory-based behavior.

Rainer and Miller (2002), analyzing PFC responses during an object delayed-matching-to-sample task, have also found strong time-dependence during the delay period and have also emphasized the need to take this into account in computational models of working memory. Our results using a parametric working memory task (Romo *et al.*, 1999; present paper) support theirs and reinforce concerns regarding the need to incorporate time-dependence into computational modeling studies.

A time-varying memory code might at first seem inconsistent with decoding a memory that should be constant. But the persistence of a memory does not necessarily require static persistence of individual neuron firing rates. It requires a looser condition, namely that some decodable, stimulus-dependent neural property remain invariant during delay periods. Static persistence of firing rates is the simplest and most natural candidate for this, but it is not the only one. A different possibility would have some combination of firing rates that must remain invariant—for example, the sum of firing rates of neurons with opposite types of time-dependence. Other possibilities exist, perhaps involving an explicit representation of time by time-dependent neurons (Leon and Shadlen, 2003). If time-dependent delay period responses are the neuronal substrate of working memory, two issues must be addressed. First, how are the time-dependent firing rates of the population decoded in order to represent a memory that, in comparison to the rapidly and strongly varying firing rates, is constant over time spans on the order of seconds? Secondly, what are the mechanisms that constrain time-varying delay period activities to firing rates that would allow the decoding of a persistent memory?

Individually static, stimulus-dependent firing rates are a particularly simple way to represent working memory in a neuronal population and thus deserve a great deal of attention. The optimal strategy to elucidate mechanisms underlying working memory may be first to fully understand mechanisms for such static representations and in some systems, such as the oculomotor integrator in the goldfish, static firing rates may indeed be the basis for representing a static memory (Aksay *et al.*, 2000, 2001). However, the preponderance of time-dependent firing rates in primate prefrontal cortex, as an almost ubiquitous feature of experimentally observed delay period responses,

suggests that a full understanding of mechanisms of working memory will require understanding the role of time-dependent representations and the mechanisms behind them.

Notes

This article is dedicated to the memory of our friend and colleague, Patricia Goldman-Rakic, a pioneer and leader in the study of persistent activity.

R.R.'s research was partially supported by an International Research Scholar award from the Howard Hughes Medical Institute and grants from the Millennium Science Initiative-CONACYT and DGAPA-UNAM. C.D.B. is supported by startup funds from Cold Spring Harbor Laboratory.

Address correspondence to Ranulfo Romo, Instituto de Fisiología Celular, Universidad Nacional Autónoma de México, Apartado Postal 70-253, C.P. 04510 México, D.F. México. Email: rromo@ifc.unam.mx.

References

- Aksay E, Baker R, Seung HS, Tank DW (2000) Anatomy and discharge properties of pre-motor neurons in the goldfish medulla that have eye-position signals during fixations. *J Neurophysiol* 84:1035–1049.
- Aksay E, Gamkrelidze G, Seung HS, Baker R, Tank DW (2001) *In vivo* intracellular recording and perturbation of persistent activity in a neural integrator. *Nat Neurosci* 4:184–193.
- Allan LG (1979) The perception of time. *Percept Psychophys* 26:340–354.
- Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb Cortex* 7:237–252.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Vis Neurosci* 13:87–100.
- Cannon SC, Robinson DA, Shamma S (1983) A proposed neural network for the integrator of the oculomotor system. *Biol Cybern* 49:127–136.
- Chafee MV, Goldman-Rakic PS (1998) Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol* 79:2919–2940.
- Compte A, Brunel N, Goldman-Rakic PS, Wang XJ (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex* 10:910–923.
- Constantinidis C, Franowicz MN, Goldman-Rakic PS (2001) The sensory nature of mnemonic representation in the primate prefrontal cortex. *Nat Neurosci* 4:311–316.
- Deneve S, Latham PE, Pouget A (2001) Efficient computation and cue integration with noisy population codes. *Nat Neurosci* 4:826–831.
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1989) Mnemonic coding of visual space in the monkeys dorsolateral prefrontal cortex. *J Neurophysiol* 61:331–349.
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1990) Visuospatial coding in primate prefrontal neurons revealed by oculomotor paradigms. *J Neurophysiol* 63:814–831.
- Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol* 36:61–78.
- Fuster JM, Bodner M, Kroger JK (2000) Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature*, 405:347–351.
- Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70:216–220.
- Green DM, Swets JA (1966) Signal detection theory and psychophysics. New York: John Wiley.
- Hernández A, Salinas E, García R, Romo R (1997) Discrimination in the sense of flutter: new psychophysical measurements in monkeys. *J Neurosci* 17:6391–6400.
- Hernández A, Zainos A, Romo R (2002) Temporal evolution of a decision-making process in medial premotor cortex. *Neuron* 33:959–972.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Kojima S, Goldman-Rakic PS (1982) Delay-related activity of prefrontal neurons in rhesus-monkeys performing delayed-response. *Brain Res* 248:43–49.

- Komura Y, Tamura R, Uwano T, Nishijo H, Kaga K, Ono T (2001) Retrospective and prospective coding for predicted reward in the sensory thalamus. *Nature* 412:546–549.
- Koulakov AA, Raghavachari S, Kepecs A, Lisman JE (2002) Model for a robust neural integrator. *Nat Neurosci* 5:775–782.
- Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 38:317–327.
- Lopez-Barneo J, Darlot C, Berthoz A, Baker R (1982) Neuronal activity in prepositus nucleus correlated with eye movement in the alert cat. *J Neurophysiol* 47:329–352.
- Miller EK, Li L, Desimone R (1993) Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13:1460–1478.
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154–5167.
- Miller P, Brody CD, Romo R, Wang X-J (2003) A recurrent network model of somatosensory parametric working memory in the prefrontal cortex. *Cereb Cortex* 13:1208–1218.
- Nieder A, Miller EK (2003) Coding of cognitive magnitude: compressed scaling of numerical information in the primate prefrontal cortex. *Neuron* 37:149–157.
- Nieder A, Freedman DJ, Miller EK (2002) Representation of the quantity of visual items in the primate prefrontal cortex. *Science* 297:1708–1711.
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*, 5:805–811.
- Press W, Teukolsky SA, Vetterling WT, Flannery BP (1992) *Numerical recipes in C*, 2nd edn. Cambridge: Cambridge University Press.
- Prut Y, Fetz EE (1999) Primate spinal interneurons show pre-movement instructed delay activity. *Nature* 401:590–594.
- Rainer G, Miller EK (2002) Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task. *Eur J Neurosci* 15:1244–1254.
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473.
- Romo R, Hernández A, Lemus L, Zainos A, Brody CD (2002) Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat Neurosci* 5:1217–1225.
- Ross SM (1987) *Introduction to probability for scientists and engineers*. New York: John Wiley.
- Salinas E, Hernández A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J Neurosci* 20:5503–5515.
- Seung HS (1996) How the brain keeps the eyes still. *Proc Natl Acad Sci USA* 93:13339–13344.
- Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* 26:259–271.
- Talbot WH, Darian-Smith I, Kornhuber HH, Mountcastle VB (1968) Sense of flutter-vibration—comparison of human capacity with response patterns of mechanoreceptive afferents from monkey hand. *J Neurophysiol* 31:301–334.
- Taube JS (1998) Head direction cells and the neurophysiological basis for a sense of direction. *Prog Neurobiol* 55:225–256.
- Walter WG (1964) Slow potential waves in the human brain associated with expectancy, attention and decision. *Arch Psychiatry Nervenkr* 206:309–322.
- Wang XJ (2001) Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* 24:455–463.
- Zipser D, Kehoe B, Littlewort G, Fuster JM (1993) A spiking network model of short-term active memory. *J Neurosci* 13:3406–3420.