

Autonomous Encoding of Irrelevant Goals and Outcomes by Prefrontal Cortex Neurons

Aldo Genovesio,¹ Satoshi Tsujimoto,³ Giulia Navarra,^{1,2} Rossella Falcone,¹ and Steven P. Wise⁴

Departments of ¹Physiology and Pharmacology and ²Basic and Applied Sciences, Engineering, Sapienza University of Rome, 00185 Rome, Italy, ³Nielsen NeuroFocus, Tokyo, 108-0071 Japan, and ⁴Olschefske Institute for the Neurobiology of Knowledge, Potomac, Maryland 20854

Two rhesus monkeys performed a distance discrimination task in which they reported whether a red square or a blue circle had appeared farther from a fixed reference point. Because a new pair of distances was chosen randomly on each trial, and because the monkeys had no opportunity to correct errors, no information from the previous trial was relevant to a current one. Nevertheless, many prefrontal cortex neurons encoded the outcome of the previous trial on current trials. A smaller, intermingled population of cells encoded the spatial goal on the previous trial or the features of the chosen stimuli, such as color or shape. The coding of previous outcomes and goals began at various times during a current trial, and it was selective in that prefrontal cells did not encode other information from the previous trial. The monitoring of previous goals and outcomes often contributes to problem solving, and it can support exploratory behavior. The present results show that such monitoring occurs autonomously and selectively, even when irrelevant to the task at hand.

Key words: memory; monitoring; monkey; prefrontal; spatial

Introduction

Autonomous coding refers to information represented, processed, and stored spontaneously, based on the inherent characteristics of a neural network, as opposed to information recognized to be of behavioral relevance through training, experience, and feedback. For example, primates need no training or conditioning for their visual areas of cortex to encode the features of objects. Visual areas encode visual features autonomously, usually in the form of conjunctions among several features. Understanding more enigmatic areas of cortex, such as the prefrontal cortex (PF), would be advanced by appreciating the information that its networks encode autonomously.

Our dataset on PF activity during a distance discrimination task (Genovesio et al., 2009, 2011, 2012) allowed us to examine autonomous coding in both the dorsolateral and caudal parts of the PF. In this task, each trial required the monkey to make a perceptual decision about which of two stimuli appeared farther from a fixed reference point, with no informational carryover from the previous trial. Specifically, we explored: (1) whether PF neurons monitored irrelevant information about past events, even when this information had never been relevant to the monkeys; and (2) whether information about past events was encoded comprehensively or selectively.

It is well known that neurons in the PF encode behavioral goals and outcomes (Genovesio et al., 2005, 2006a, 2008, 2012;

Saito et al., 2005; Passingham and Wise, 2012; Yamagata et al., 2012). Although many studies have concentrated on the representation of future goals (Rainer et al., 1999) and predicted outcomes (Roesch and Olson, 2003; Wallis and Miller, 2003; Tsujimoto and Sawaguchi, 2005; Kennerley and Wallis, 2009), ample evidence shows that PF also encodes previous goals and outcomes (Barracough et al., 2004; Tsujimoto and Sawaguchi, 2004; Genovesio et al., 2006a). Neurophysiological studies have also demonstrated that PF neurons encode task-irrelevant features of the environment (Mann et al., 1988; Lauwereyns et al., 2001; Genovesio et al., 2006b; Kim et al., 2008). We found, for example, that PF cells encoded the previous goal on a spatial delayed-response task (Tsujimoto et al., 2012), although nothing about the previous trial had any relevance to a current trial. However, in this example and the others cited above, the monkeys had previously performed tasks in which the encoded information had been relevant. Only rarely has training experience involving the irrelevant information been ruled out in studies of this kind (Chen et al., 2001), and that report involved only the encoding of extraneous features of a behavior-guiding cue, with no analysis of previous goals or outcomes. Accordingly, the present analysis is the first to examine the encoding of irrelevant information about previous outcomes and goals in monkeys lacking experience on any task that required the use of such information.

Materials and Methods

Subjects and apparatus. Two adult male rhesus monkeys (*Macaca mulatta*), 8.5 and 8.0 kg, served as subjects in this study. All procedures followed the *Guide for the Care and Use of Laboratory Animals* (1996, SBN 0-309-05377-3) and were approved by the National Institute of Mental Health Animal Care and Use Committee.

The two monkeys sat 29 cm from a video screen (Fig. 1a) with their heads fixed. Within reach there were three 3 × 2 cm switches that detected hand contact with infrared sensors (Fig. 1a7). These switches were

Received July 30, 2013; revised Dec. 12, 2013; accepted Dec. 19, 2013.

Author contributions: A.G. designed research; A.G. and S.T. performed research; A.G., G.N., and R.F. analyzed data; A.G. and S.P.W. wrote the paper.

This work was supported by Division of Intramural Research of the National Institute of Mental Health Grant Z01MH-01092. We thank Dr Andrew Mitz, James Fellows, and Ping-Yu Chen for technical support.

Correspondence should be addressed to Aldo Genovesio, Department of Physiology and Pharmacology, Sapienza University of Rome, Piazzale Aldo Moro 5, 00185 Rome, Italy. E-mail: aldo.genovesio@uniroma1.it.

DOI:10.1523/JNEUROSCI.3228-13.2014

Copyright © 2014 the authors 0270-6474/14/341970-09\$15.00/0

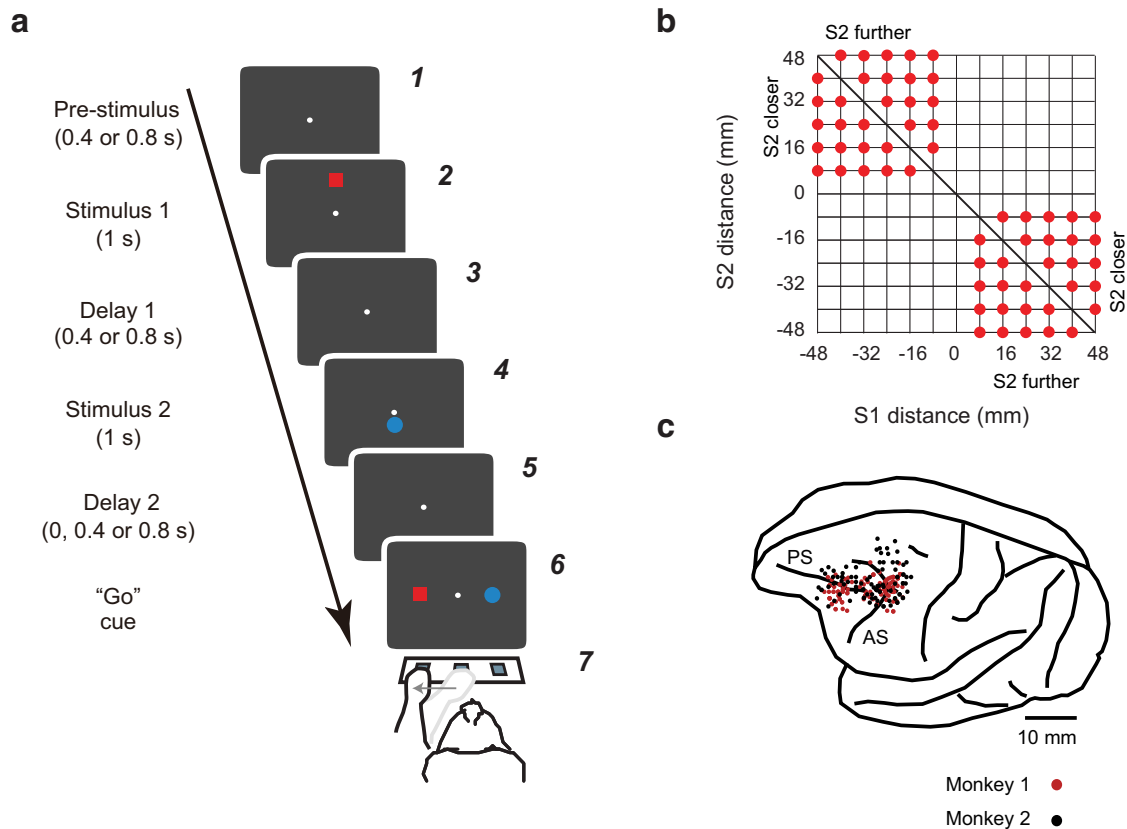


Figure 1. Task, distances used, and recording locations. **a**, Sequence of task events. Each gray rectangle represents the video screen. The array of switches is depicted in **a7**. In the illustrated example, the red square appeared farther from the central reference point (white circle) than did the blue circle. Therefore, when the two stimuli reappeared with the red stimulus to the left and the blue stimulus to the right (**a6**), the correct choice was to press the left switch. **b**, Distances of S1 and S2 from the central reference point. **c**, Penetration sites. Composite from both monkeys, relative to sulcal landmarks. AS, Arcuate sulcus; PS, principal sulcus.

located just beneath the video screen, from left to right, with a separation (center-to-center) of 7 cm. Both monkeys contacted the switches with their left hands.

The stimulus material consisted of a fixed reference point at the center of the video screen, which was a 0.6° (diameter) solid white circle, along with a 3° (diameter) solid blue circle and a $3^\circ \times 3^\circ$ solid red square.

Task events. Figure 1*a* illustrates the task. On each trial, two stimuli, a red square and a blue circle, appeared sequentially at different distances from a reference point. S1 refers to the first stimulus, regardless of its color, shape, or distance from the reference; S2 refers to the second stimulus.

A trial began when the monkey touched the central key, at which time the white circle appeared at the center of the screen (Fig. 1*a1*). After a prestimulus period of either 400 or 800 ms, the two visual stimuli, the red square and the blue circle, appeared in succession for 1.0 s each (Fig. 1*a2,a4*). One stimulus appeared directly below the reference point, and the other appeared directly above it. Both the order of the stimuli (red vs blue) and their location (above vs below screen center) were randomly selected. The relative distance of each stimulus from the reference point ranged from 8 to 48 mm from the reference point, in steps of 8 mm, corresponding to 1.6° , 3.2° , 4.7° , 6.3° , 7.9° , and 9.4° of visual angle (Fig. 1*b*).

A delay period followed both S1 and S2 (Fig. 1*a3,a5*). The first delay period (D1) followed S1, and it usually lasted either 400 or 800 ms (randomly selected). In one set of sessions, however, we added a D1 period of 1200 ms, and in another we used a fixed D1 period of 1200 ms. The second delay period (D2), which followed S2, lasted either 0, 400, or 800 ms (randomly selected). At the end of the D2 period, both the red and blue stimuli reappeared, one 7.8° directly to the left of the fixation point and the other 7.8° to the right, randomly determined (Fig. 1*a6*). This event served as the “go” cue for the monkeys’ reaching movement. We

allowed the monkeys 6 s to touch either the left or right switch, although in practice the monkeys made their choices much faster (Genovesio et al., 2011). The correct choice, the one that would produce a fluid reward, was the switch beneath the stimulus that had appeared farthest from the reference point earlier on that trial (Fig. 1*a7*). A variable intertrial interval (ITI), usually 700–1000 ms, elapsed before the beginning of the next trial.

Surgery. Recording chambers were implanted over the exposed dura mater of the left frontal lobe, using aseptic techniques and isoflurane anesthesia (1–3%, to effect). Monkey 1 had two 18-mm-diameter chambers, one placed over the caudal PF cortex, the other over the dorsolateral PF cortex; Monkey 2 had a single 27×36 mm chamber that encompassed both areas (Fig. 1*c*).

Data collection. We used an infrared oculometer to record eye position (Arrington Recording). Single-cell activity was recorded using quartz-insulated platinum-iridium electrodes (0.5 – 1.5 M Ω at 1 kHz) positioned by a 16-electrode drive assembly (Thomas Recording). The electrodes were arranged in a concentric recording head with 518 mm spacing. We discriminated single-unit potentials online with the Multichannel Acquisition Processor (Plexon) and confirmed each isolated waveform with Off Line Sorter (Plexon). The offline verification of unit isolation used principal component analysis, minimal interspike intervals, and clearly differentiated waveforms inspected individually for every isolated neuron. Inadequately isolated potentials were eliminated from the dataset before any additional analysis.

Eye position was not placed under experimental control because of the tendency of both monkeys to saccade to each stimulus when it appeared. However, the central reference point was the experimentally controlled fixation point in two companion tasks, a duration-discrimination task and a matching-to-sample task (Genovesio et al., 2009, 2012), which

explains why both monkeys tended to fixate the reference point during the prestimulus period and at the time that S1 appeared.

Task periods. We measured activity in several fixed temporal windows. The prestimulus period extended from 80 to 400 ms after the onset of the reference point (Fig. 1a1); the early S1 period was 80–400 ms after the onset of S1 and the late S1 period lasted from 400 to 1000 ms after the onset of the 1000-ms-long S1 (Fig. 1a2). The same time windows applied to the early S2 period and the late S2 period, relative to S2 onset (Fig. 1a4). For activity rates during the D1 period (Fig. 1a3,a5), we took the interval from 80 to 400 ms after onset. The reaction and movement period (RMT) lasted from the “go” cue until goal acquisition.

Neurophysiological analysis. We performed one-way ANOVAs on activity during the prestimulus period to determine whether a given neuron might have encoded any of several factors: the outcome on the previous trial (reward or nonreward), the features of the previous goal (the red square or blue circle), the location of the previous goal (left or right), the ordinal position of the previous goal (first or second), the features of S2 (the red square or blue circle), the location of S2 (above or below screen center), and the distance of S2 from the reference point.

Then, to study the time course of the previous outcome and goal selectivity, we performed a three-way ANOVA with previous outcome, goal location, and features as factors.

We also computed the area under the receiver operating characteristic (ROC) curve (AUC) to measure the selectivity for encoding the previous outcome or goals, with an AUC value of 0.5 indicating no selectivity and 1.0 corresponding to maximal selectivity. A bootstrap analysis estimated the chance-level AUC by assigning the relevant activity to one of the two samples randomly, with repetition, 1000 times, averaged over the population.

For cells encoding some aspect of the goal, we calculated two strategy indices (SI) in each recording session and correlated them with activity in the prestimulus period related to two factors: goal location (left or right) and goal features (e.g., red or blue). We contrasted the number of trials on which the monkeys applied a win–stay/lose–shift strategy (WSt), with the number of trials on which they used a win–shift/lose–stay strategy (WSh). The first index was for goal location: $SI_{\text{location}} = (WSt - WSh)/(WSt + WSh)$, where WSt and WSh refer to goal location. The second index was for goal features: $SI_{\text{features}} = (WSt - WSh)/(WSt + WSh)$, in which WSt and WSh refer to the color and shape of the goal. Then we calculated modulation indices for the two kinds of goal coding. First, we calculated an absolute modulation index for goal location: $|(A_{\text{right}} - A_{\text{left}})/(A_{\text{right}} + A_{\text{left}})|$, where A_{right} was the mean activity rate during trials when the goal had been on the right on the previous trial and A_{left} represents activity for goals to the left. Second, we calculated an absolute goal features index: $|(A_{\text{Blue}} - A_{\text{Red}})/(A_{\text{Blue}} + A_{\text{Red}})|$, where A_{Blue} was the mean activity on trials with a previous blue goal and A_{Red} represents activity for a previous red goal.

Histological analysis. Near the end of recording, we made electrolytic marking lesions (15 μ A for 10 s). Ten days later, the monkeys were deeply anesthetized and perfused through the heart with 10% formol saline. After sectioning the brain and staining the section for Nissl substance, we plotted the recording sites by reference to the marking lesions, pins inserted at the time of the perfusion, and structural magnetic resonance images (MRI) taken periodically before and between recording sessions. Although the entry points for more posterior recordings (Fig. 1c) make it appear that many cells were located in the postarcuate cortex, reconstructions of the recording sites that take into account the angle and depth of penetrations, as well as the structural MRI data, show that nearly all recordings in the caudal PF cortex came from the prearcuate cortex (area 8). Neurons in and near the principal sulcus were mostly located in area 46.

Results

Behavior

Both monkeys performed the task accurately by reaching to a touch pad in the same direction from center as the stimulus that had appeared farther from the reference, with mean scores of 78% and 80% correct for Monkeys 1 and 2, respectively. Easier

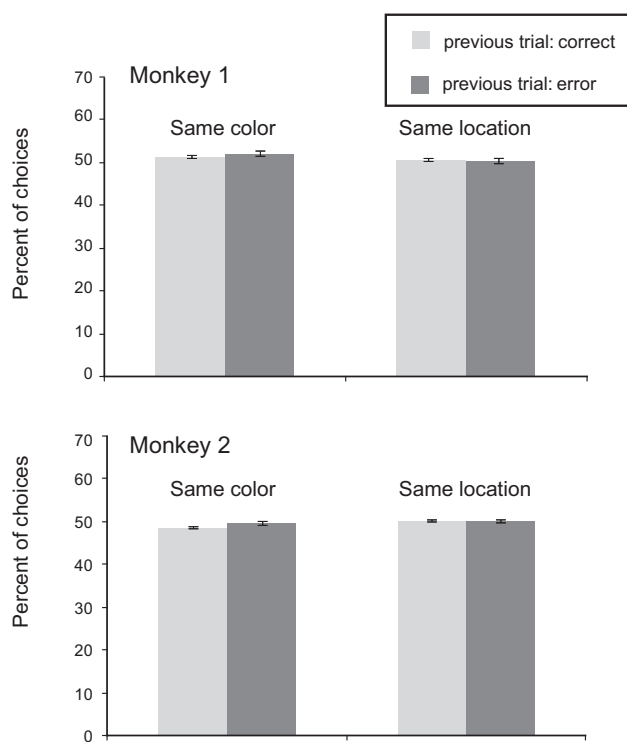


Figure 2. Percentage of trials, by subject, in which the current goal choice is the same as in the previous trial in terms of either goal features (left) or location (right) after correct and error trials. Error bars indicate SEM by session.

discriminations were associated with both higher accuracy and faster response latencies (Genovesio et al., 2011).

An important feature of the experimental design was the fact that there was no correction trial after an incorrect choice. Nevertheless, we could test whether the monkeys attempted to correct previous errors by applying a lose–shift strategy, such as switching from the red square to the blue circle or from the left response to the right. Neither monkey adopted these strategies (Fig. 2).

We also calculated the reaction times (RTs) for different classes of trials to evaluate whether there was a relationship between response speed and previous trial history. We first calculated the RTs after correct and error trials, separately (Table 1). We found significantly faster RTs after correct trials than after error trials in both monkeys (t test; $p < 0.001$), a 31 ms difference in Monkey 1 and a 16 ms difference in Monkey 2. Then, we divided both after-correct and after-error trials into those with the same (stay) or a different (shift) goal as on the previous trial, in terms of either goal location (right or left) or goal features (blue or red). We compared the RTs between stay and shift trials for both types of goals, after both correct and error trials. We found a small but significant difference only in Monkey 1 (Table 1), with a 2 ms faster RT after correct trials when the monkey shifted between red and blue goals (paired t test, $t = 3.03$, $p < 0.05$) and a 14 ms faster RT after error trials when the monkey shifted between left and right goals (paired t test, $t = 3.3$, $p < 0.05$).

Database

The database comprised 1687 neurons that met rigorous isolation criteria: 1301 from caudal PF (PFC) and 386 from dorsolateral PF (PFDl). Of this population, 501 cells came from the monkey 1 and 1186 from Monkey 2. All analyses presented in this report were restricted to neurons that met the following inclusion criteria: activity was recorded for at least 50 trials, including at

Table 1. Mean RTs divided by monkey, and by whether the previous trial had been executed correct or incorrect

	Mean RT, ms (SD)									
			After correct				After error			
	After correct	After error	Goal location		Goal features		Goal location		Goal features	
			Stay	Shift	Stay	Shift	Stay	Shift	Stay	Shift
Mk 1	344 (0.03)	375 (0.05)	343 (0.03)	345 (0.03)	343 ^a (0.03)	345 ^a (0.03)	383 ^a (0.07)	369 ^a (0.06)	372 (0.05)	377 (0.06)
Mk 2	394 (0.04)	410 (0.06)	393 (0.04)	395 (0.04)	395 (0.04)	395 (0.04)	408 (0.07)	412 (0.07)	408 (0.07)	412 (0.07)

^aStatistically significant stay-shift differences in RT.

For both, reaction times are divided into trials in which the monkeys (MK) stayed with their previous goal, correct or incorrect, or shifted to the alternative goal, for both spatial and nonspatial goals: location and features, respectively.

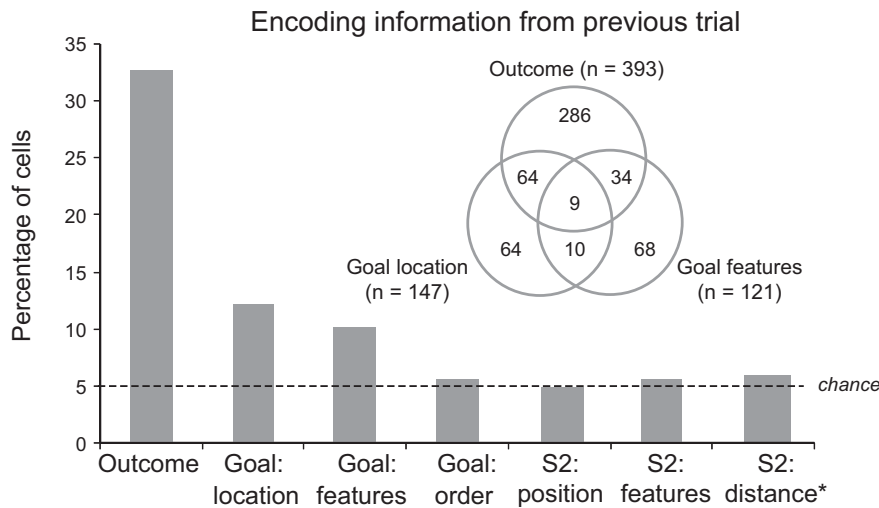


Figure 3. Percentage of cells showing a significant effect of the previous outcome, previous-goal location, previous-goal features, previous-goal order, and the position, features, and distance from the reference point of S2 on the previous trial. The Venn diagram gives the number of cells with significant effects and conjunctions of effects.

Table 2. For the prestimulus period, results of one-way ANOVA for previous outcome (reward), previous goal location (left or right), previous goal features (red square or blue circle), presentation order of the previous goal (first or second), position of the second stimulus (above or below center), features of the second stimulus (red square or blue circle), or distance of the second stimulus from center^a

Area	Outcome	Goal location	Goal features	Goal order	S2 position	S2 features	S2 distance
PFC	34.1	12.0	10.0	5.2	5.0	5.5	5.8
PFdl	27.6	13.1	10.1	6.3	4.5	5.6	6.3

^aPercentage of cells with significant effects.

PFC, Caudal PF; PFdl, dorsolateral PF; S2, second stimulus.

least four trials for each combination of previous goal location, previous goal features and previous outcome; and the cells needed to have an average activity rate of at least 1.0 spikes/s. Of the overall population, 1204 neurons met these criteria: 936 from PFC and 268 from PFdl. Of these neurons, 347 came from Monkey 1 and 857 came from Monkey 2.

Encoding from the previous trial

Figure 3 shows the principal result. Of the several factors tested, only the previous outcome and the previous goal were significantly encoded during the prestimulus period. Approximately one-third of the sampled PF neurons showed a significant effect of the previous outcome (reward or nonreward; 393 of 1204 cells, 33%). Smaller but still significant percentages showed an effect of the location (left or right) of the previous goal (147 cells, 12%) or its color-and-shape features (121 cells, 10%). Both of these percentages exceeded chance levels ($p < 0.001$, binomial test). As

shown by the Venn diagram in Figure 3, many neurons encoded conjunctions of goal and outcome on the previous trial (117 cells, 10% of the tested population; 18% of the neurons with significant effects of previous goal or outcome).

Figure 3 also shows that our sample of PF neurons encoded the previous outcome and goal selectively. These cells did not encode other information from the previous trial, such as the color-and-shape features of S2 (66 cells, 5%), the location of S2 (up or down from center; 59 cells, 5%), the distance of S2 from the reference point (71 cells, 6%), or whether the stimulus ultimately chosen by the monkey had appeared first or second during the previous trial (71 cells, 6%). None of these proportions significantly exceeded those expected by chance ($p > 0.05$, binomial test).

Table 2 gives the percentage of cells showing significant effects of each factor tested, for both PFC and PFdl. Results for the two monkeys were similar.

Figure 4a shows a neuron from the PF cortex encoding the previous outcome. During most of the intertrial interval and the entire prestimulus period, this neuron was significantly more active after error trials than after correct trials. Figure 4b shows another cell that encoded the outcome on the previous trial. This neuron had higher activity after correct trials than after errors, and it showed this effect exclusively during the S1 period. Of the 393 neurons encoding the previous outcome, a higher proportion (χ^2 test, $\chi^2 = 18.9$; $p < 0.001$) had greater activity after an error trial (227 cells, 58%) than after a correct trial (166, 42%), although behavior, in terms of error rate, did not vary according to the correctness of the previous trial.

Figure 4c,d shows the activity of a third example neuron, one that encoded the previous spatial goal. During most of the intertrial interval, the entire prestimulus period and into the first part of the S1 period, this cell had greater activity when the previous goal had been to the left (Fig. 4c). This cell showed no effect of the previous stimulus (S2) location, above or below center (Fig. 4d). Of the 147 neurons encoding the previous goal location, a greater proportion (binomial test, $p = 0.013$) preferred a previous choice to the left (89 cells, 61%) than a previous choice to the right (58, 39%), which corresponded to a ipsilateral bias.

Figure 5 shows an example neuron encoding a conjunction of features from the previous trial. This cell encoded a combination of the nonspatial and spatial features of the previous goal, with a preference for the red square to the right (Fig. 5, top left). This

preference developed during the last part of the intertrial interval and persisted into the first part of the S1 period. Note that this cell also encoded other aspects of the task, as indicated by the activity modulation before goal acquisition (bottom row).

Although there was no evidence of the monkeys applying a consistent strategy of staying or shifting in response to “wins” and “losses” on the previous trial (Fig. 2), it was still possible that their strategy could change over time and correlate with neuronal activity during any given recording session. Accordingly, we tested for significant correlations between goal coding and two strategy indices (SI) (see Materials and Methods) calculated on a session-by-session basis. For these indices, a value of ± 1.0 indicates the consistent use of either a stay or shift strategy, and a value of 0.0 indicates that the monkeys did not use either strategy. The low strategy indices of only 0.08 ± 0.07 for SI_{features} and 0.08 ± 0.06 for SI_{location} show that the monkeys did not use stay or shift strategies during individual recording sessions. Furthermore, there was no significant correlation between the strategy index and the encoding of either goal location ($r = -0.032$; $p = 0.70$) or goal features ($r = 0.052$; $p = 0.57$) for the cells with significant effects.

Time course of previous goal and outcome coding

A three-way ANOVA evaluated the time course of encoding the previous outcome, goal location and goal features. Figure 6 presents results for these three main effects, and Table 3 gives the results divided by area, along with the interactive effects.

Figure 6a shows a progressive decline of each signal, as seen in the population average, over the duration of the current trial. Both location and feature encoding for the previous goal reached a chance level during the S1 period. In contrast, previous-outcome encoding, although experiencing a significant decay, persisted until the end of the S2 period. Figure 6a also shows that the signal for the same three factors was near the chance level when calculated for the trial two trials in the past. An ROC analysis (Fig. 6b), which measured the level of encoding for all the three factors, confirmed the results from ANOVA.

The population plots in Figure 6 do not address how the encoding of previous goals and outcomes either decayed or emerged at the single-cell level during the course of a trial. Figure 7 does so, and to systematically search for carryover effects from the previous trial, we began this analysis with the intertrial interval.

We found that many cells encoded their particular signals during the prestimulus period, as expected from the population plots (Fig. 6). However, approximately half of the neurons that encoded the outcome in the S1 period started to do so only after the prestimulus period (121 of 260 outcome-coding cells, 47%; Fig. 7a) as exemplified by the cell illustrated in Figure 4b. Likewise, a large proportion of neurons encoding the previous goal location in the S1 period began doing so only after the prestimulus period (71 of 94 cells, 76%; Fig. 7b), and a similar result was obtained for neurons encoding the nonspatial features of the previous goal (66 of 84 cells, 79%, Fig. 7c).

We performed a similar analysis for the persistence of goal-encoding signals from the RMT period of the previous trial

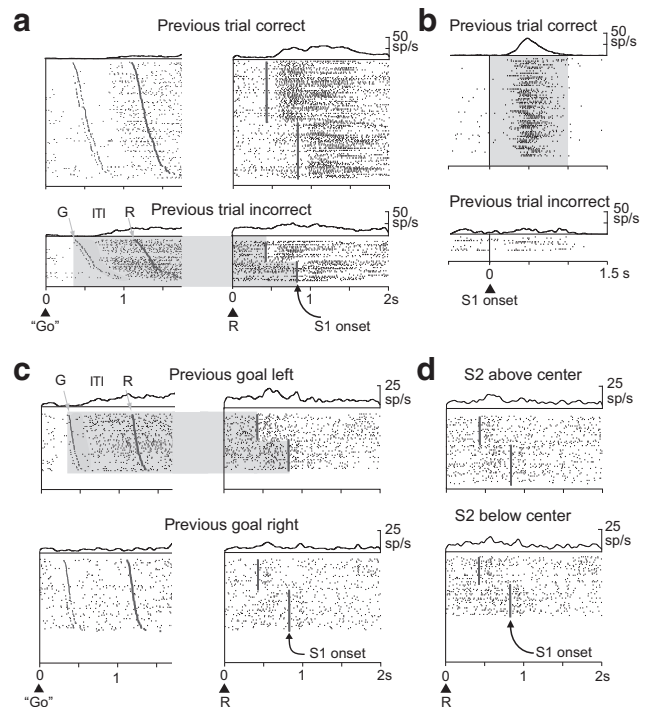


Figure 4. Neurons encoding the previous outcomes or goal locations. Raster displays show spike (sp) times beneath spike-density averages. **a**, Activity of a PF cortex neuron encoding the previous outcome during the intertrial interval (ITI) and the prestimulus period, with a preference for error trials (light shading). Rasters are aligned on the go stimulus during the previous trial (left) and on the onset of the reference point (R) during the current trial (right). The ITI varied from 760 to 820 ms. Note the time gap between goal acquisition on the previous trial (G) and the onset of encoding the previous outcome. **b**, A different PF cortex cell, one that encodes the previous outcome during the S1 period with higher activity after correct trials (gray shading). Rasters are aligned on S1 onset. **c**, A third PF cortex cell, one with activity encoding the previous goal location, with higher activity after previous goals to the left (light gray shading). Shading and raster alignment as in **a**. **d**, The same cell as in **c**, showing that it did not encode the previous S2 position relative to the reference point. Rasters are aligned on the reference point onset (R).

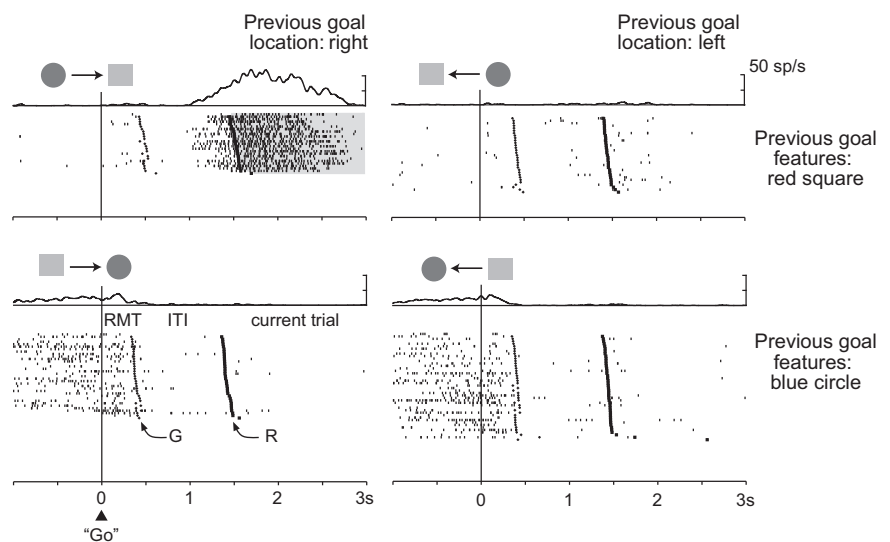


Figure 5. Cell encoding both the location and features of the previous goal, with higher activity when the goal was the red square and when it was to the right. The display above each raster plot indicates the configuration of the two potential goal stimuli on the previous trial (shape only), and the goal selected (arrow). Rasters are aligned on the go signal during the previous trial (triangle); the cross to the right of the alignment line corresponds to goal acquisition on the previous trial (G) and the square mark indicates the onset of the reference point during the current trial (R). Abbreviations as in Figure 4.

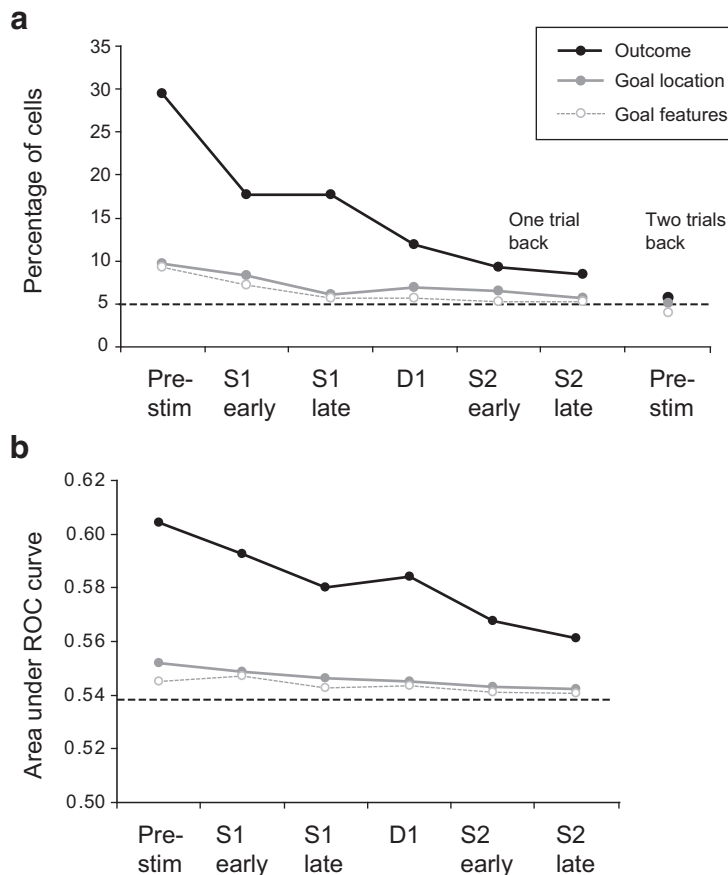


Figure 6. *a*, Percentage of cells with a significant effect of previous outcome, previous goal location and previous goal features (three-way ANOVA, $p < 0.05$) in several task periods. Pre-stim, Prestimulus period. To the right, results are presented for the prestimulus period two trials back. *b*, Results of ROC analysis for the same factors as in *a*. The dashed line represents the value for the AUC expected by chance, as calculated by the bootstrap method.

Table 3. Results of three-way ANOVA for previous outcome and the locations and features of the previous goal^a

Task period	Area	Interactive terms						
		Outcome	Goal location	Goal features	Outcome & location	Outcome & features	Location & features	3-way
Pre-stim	PFC	30.3	10.0	9.2	7.8	8.0	8.0	7.3
	PFdl	26.5	8.6	9.7	6.3	7.1	7.1	6.3
S1 early	PFC	18.1	9.1	7.2	6.8	6.8	7.5	7.2
	PFdl	16.8	5.6	7.5	4.5	5.6	6.0	6.0
S1 late	PFC	17.5	6.7	5.6	6.0	5.1	6.5	5.7
	PFdl	18.3	3.7	6.3	4.5	5.6	7.1	6.7
D1	PFC	11.8	6.4	5.8	5.7	4.8	6.2	5.3
	PFdl	12.3	8.6	5.2	7.5	7.5	6.0	7.5
S2 early	PFC	9.1	6.7	5.7	5.1	5.3	5.2	4.6
	PFdl	9.7	5.6	4.1	4.9	5.6	6.3	5.6
S2 late	PFC	8.7	6.3	5.3	4.9	5.3	6.6	6.4
	PFdl	7.5	3.4	5.2	3.7	3.0	6.0	5.2
RMT	PFC	7.3	6.0	4.8	4.8	5.6	5.2	6.6
	PFdl	7.8	8.2	7.1	5.6	8.6	7.5	8.2

^aPercentages of cells.

PFC, Caudal PF; PFdl, dorsolateral PF; Pre-stim, prestimulus; RMT, reaction and movement period.

to the prestimulus period of a current trial. Only a small minority of cells encoding the previous goal location in the prestimulus period did so during the previous RMT period (13 of 147 cells, 9%, four of which changed their preferred goal location). A somewhat higher proportion of neurons selective for goal features in the prestimulus period also had this property during the

previous RMT period (46 of 121 cells, 38%), but only about half of these cells (24 of 46, 52%) had the same preferred goal feature. The cell shown in Figure 5 exhibited such activity.

Our findings reveal considerable independence in the coding properties of individual neurons across task periods. Figure 7 shows that for all of the variables encoded from the previous trial, many cells first developed their particular signal during the prestimulus period or during subsequent task periods, and not during the ITI. These observations rule out the possibility that the signals result from a simple continuation of activity from the previous trial. Even cells with previous outcome or goal encoding during the ITI do not always have this property in the earliest part of the ITI (Figs. 4*a,c*, 5), a property that also rules out carryover from the previous trial.

Discussion

We found that PF neurons encode and monitor previous outcomes and goals autonomously, i.e., regardless of task relevance and previous training. The goal coding includes both its location and non-spatial features, and the autonomous coding is selective in that it does not include other irrelevant information from the previous trial, such as order information, the position or nonspatial features of the previous cue (S2), or its distance from the central reference point. We also found

that the coding of previous outcomes and goals did not reflect a simple continuation or carryover from the previous trial. Instead, a high proportion of neurons encoded outcome and goal signals at various times during a trial (Fig. 7), as described previously for various within-trial effects (Takeda and Funahashi, 2007; Sigala et al., 2008; Genovesio et al., 2012).

As reviewed in the Introduction, many previous studies have documented the encoding of previous goals and task-irrelevant stimuli by PF neurons, including for the two areas studied here: PFdl and PFC. The present report, however, is the first to demonstrate encoding that is both autonomous, in that the monkeys had not learned to use that information in an operantly conditioned behavior, and selective for previous outcomes and goals.

Although irrelevant to the present task, the monitoring of previous goals and outcomes often contributes to general problem solving. Our results thus support the idea, derived mainly from neuroimaging and lesion studies (Duncan, 2010), that PF contributes to this function.

Irrelevant information

Many neurophysiological studies show that irrelevant information can modulate PF activity (Mann et al., 1988; Chen et al., 2001; Lauwereyns et al., 2001; Genovesio et al., 2006*b*), although some did not (Rainer et al., 1998; Everling et al., 2006).

Often, the irrelevant information has been relevant in a subject's experience. For example, Lauwereyns et al. (2001) described the activity of neurons in a go–no-go task. The monkey's

choice depended on either the direction of visual motion or color. The irrelevant dimension affected cell activity, although much less than did the relevant dimension. Likewise, Kim et al. (2008) reported PF neurons that encoded the number of stimuli, which indicated the duration of a delay period. This property persisted during blocks of trials in which the stimuli ceased to indicate any task-relevant information. Cells in the frontal eye field show an enhanced representation of distractors that had been the target during previous visual search trials (Bichot and Schall, 1999). These examples, and others (Mann et al., 1988; Lauwereyns et al., 2001; Genovesio et al., 2006b; Kim et al., 2008; Tsujimoto et al., 2012), show that the encoding of irrelevancies can occur when information that is temporarily irrelevant has been relevant in the past (and might become relevant again).

Specific task demands could also have an effect. For example, biased competition (Desimone and Duncan, 1995) might suppress the irrelevant signal more within a given processing system, such as the ventral visual stream (Rainer et al., 1998), than for information from different processing streams (Lauwereyns et al., 2001).

In another analysis from the present dataset (Genovesio et al., 2011), only 9% of the cells encoded the distance of S1 from the reference point during the D1 period. This information is not only relevant, but it is necessary for successful task performance because it must be used in a subsequent comparison. Yet this highly relevant signal appears to be less prominent than the irrelevant information about the goal chosen on the previous trial, as reported here.

In the present task, the irrelevant information from the previous trial had little effect on either behavior or cell activity. We found, for example, that trial history had only a minor effect on RT (Table 1), and there was no relationship between trial history and neural modulation.

Irrelevant information about the previous goal could be encoded as part of a high-dimensional representation of experienced events (Rigotti et al., 2013). In this context, the irrelevant information allows a network to adapt flexibly to future tasks simply by modifying the weights of a readout neuron that operates as a classifier.

Goal and output monitoring

Several studies have shown that PF neurons encode future (Rainer et al., 1999; Genovesio et al., 2006a, 2012; Tanji and Hoshi, 2008; Tsujimoto et al., 2008, 2012) and previous (Barraclough et al., 2004; Tsujimoto and Sawaguchi, 2004; Genovesio et al., 2006a) goals. The representation of previous goal choices can be viewed in the context of a network that includes, in addition to PF, the anterior cingulate and posterior parietal cortex, among other areas (Genovesio et al., 2006a; Seo and Lee, 2007; Seo et al., 2007, 2009; Tsujimoto et al., 2012).

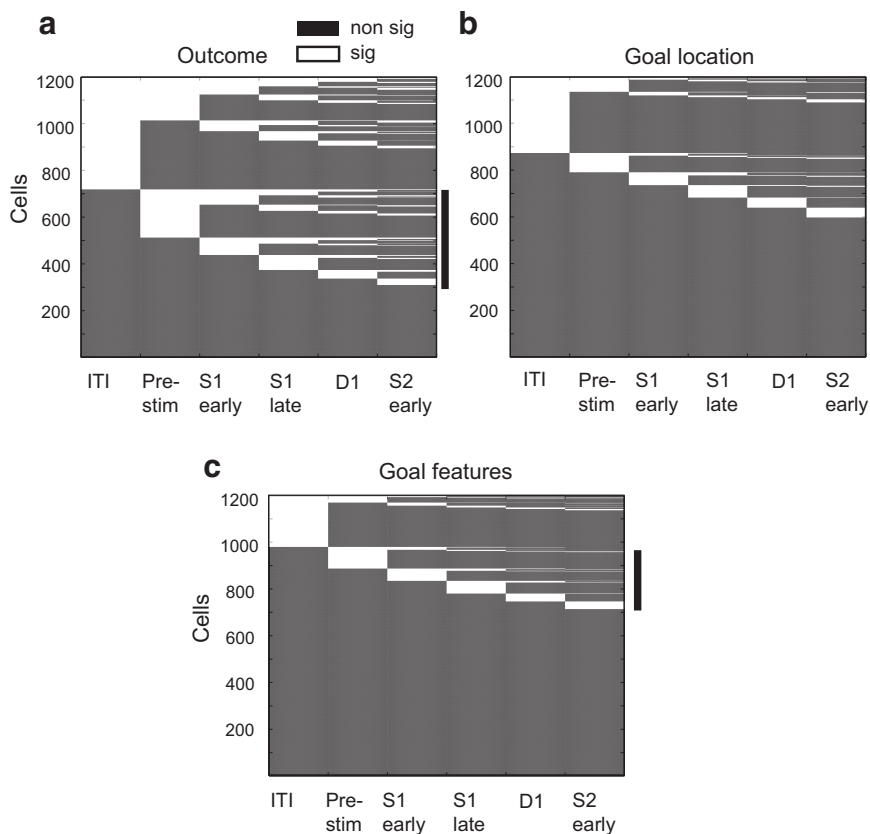


Figure 7. Coding over the course of a trial. The coding properties of each neuron are shown in a line, with statistically significant (sig) coding (one-way ANOVA, $p < 0.05$) of outcome (**a**), goal location (**b**), or goal features (**c**) indicated by white line segments, and task periods without significant coding (non sig) indicated by black line segments. The vertical black bar to the right of each plot marks the cells with no significant previous-trial signal during the intertrial (ITI) but significant coding of either the previous outcome (**a**) or goal (**b, c**) during the prestimulus (Pre-stim) and subsequent task periods.

Tsujimoto and Postle (2012) reported the results of disrupting parts of PFdl on a variant of the spatial delayed-response task. Previous experiments had concluded that small lesions of PFdl caused a mnemonic “scotoma”, i.e., the loss of spatial working memory for a portion of the visual field (Funahashi et al., 1993). Tsujimoto and Postle (2012) found, however, that incorrect responses tended to be associated with the goal location on the previous trial, and after that error the monkeys usually made a corrective movement to the cued location. Thus, the impairment after PFdl lesions seems to be in goal or output monitoring, and is not accurately characterized as a mnemonic scotoma.

PF has been implicated in goal or output monitoring by other tasks, as well. For example, the self-ordered task requires subjects to monitor previously chosen objects in a set of objects. PFdl lesions cause an impairment in this task in monkeys, and cortical activations are observed during the performance of this task in humans (Owen et al., 1996; Petrides, 2005). Output monitoring is also required in tasks involving the switching among abstract rules or relevant stimulus dimension. The effects of lesions, neuronal activity, and imaging activations for these tasks (Wallis et al., 2001; Mansouri et al., 2006) support a role of PFdl in output monitoring, specifically in monitoring chosen goals. Passingham and Wise (2012) advanced a detailed argument to support a role for PFdl in monitoring potential goals in terms of order information, as opposed to the maintenance of spatial working memories.

Surprisingly, we found an ipsilateral preference for previous goals, as reported for postsaccadic activity in PFdl (Rao et al.,

1999). Likewise, Tsujimoto and Postle (2012) reported that when disinhibition of PFdl induces erroneous responses toward the previous target location, it does so especially for the visual field ipsilateral to the injection site. These properties contrast with the contralateral deficit on the delayed-response task observed after PFdl lesions (Funahashi et al., 1993) and a contralateral bias for neuronal activity for future goals (Rao et al., 1999). An ipsilateral preference for one kind of goal combined with a contralateral preference another kind could contribute to representing them separately, which is essential for output monitoring (Genovesio et al., 2006a). Future work might explore this possibility.

Outcome monitoring

Previous-outcome signals have been reported in many studies in the period immediately after reward delivery, in both cortical and subcortical areas (Tsujimoto and Sawaguchi, 2004; Wirth et al., 2009). The representation of previous outcomes occurs in inter-trial and delay periods in PFdl (Tsujimoto and Sawaguchi, 2004; Mansouri et al., 2006), the anterior cingulate cortex (Ito et al., 2003; Seo and Lee, 2007; Kennerley et al., 2011), the lateral intraparietal area (Seo et al., 2009), and the orbitofrontal cortex of rats (Schoenbaum and Eichenbaum, 1995) and monkeys (Simmons and Richmond, 2008; Kennerley and Walton, 2011; Kennerley et al., 2011), as well as in subcortical structures such as the striatum (Samejima et al., 2005; Cai et al., 2011).

Here we found that PF cells encoded the outcome that had occurred on the previous trial more strongly than other signals encoded at the same time, such as the goal chosen on the previous trial or working memory signals for durations and distances (Genovesio et al., 2009, 2011, 2012). Therefore, this signal probably reflects an important aspect of PF function.

It has been proposed that value functions should be continuously revised based on new experiences, comparing predicted and obtained reward (Schultz, 2007). Previous studies have shown that the outcome of a previous trial modulated PFdl activity in a variety of tasks, such as for the optimal strategy in a mixed strategy game (Barraclough et al., 2004; Seo et al., 2007) and for conditional motor learning (Histed et al., 2009). In the context of reinforcement learning models, reward expectation should be based on the previous history of rewards, with a recency weighting that depends on environment stability. In the mixed strategy game (Seo et al., 2009), neurons in both frontal and parietal areas encoded the previous outcome on various timescales. Some timescales were available for filtering fast-changing outcome events in volatile conditions, with others serving a similar function for long timescales in stable environments (Bernacchia et al., 2011). The same neuron could adapt its time constant, probably in relationship to the search for a better solution to the problem posed by the task.

References

- Barraclough DJ, Conroy ML, Lee D (2004) Prefrontal cortex and decision making in a mixed-strategy game. *Nat Neurosci* 7:404–410. [CrossRef Medline](#)
- Bernacchia A, Seo H, Lee D, Wang XJ (2011) A reservoir of time constants for memory traces in cortical neurons. *Nat Neurosci* 14:366–372. [CrossRef Medline](#)
- Bichot NP, Schall JD (1999) Effects of similarity and history on neural mechanisms of visual selection. *Nat Neurosci* 2:549–554. [CrossRef Medline](#)
- Cai X, Kim S, Lee D (2011) Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice. *Neuron* 69:170–182. [CrossRef Medline](#)
- Chen NH, White IM, Wise SP (2001) Neuronal activity in dorsomedial frontal cortex and prefrontal cortex reflecting irrelevant stimulus dimensions. *Exp Brain Res* 139:116–119. [CrossRef Medline](#)
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222. [CrossRef Medline](#)
- Duncan J (2010) The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends Cogn Sci* 14:172–179. [CrossRef Medline](#)
- Everling S, Tinsley CJ, Gaffan D, Duncan J (2006) Selective representation of task-relevant objects and locations in the monkey prefrontal cortex. *Eur J Neurosci* 23:2197–2214. [CrossRef Medline](#)
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1993) Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic “scotomas.” *J Neurosci* 13:1479–1497. [Medline](#)
- Genovesio A, Brasted PJ, Mitz AR, Wise SP (2005) Prefrontal cortex activity related to abstract response strategies. *Neuron* 47:307–320. [CrossRef Medline](#)
- Genovesio A, Brasted PJ, Wise SP (2006a) Representation of future and previous spatial goals by separate neural populations in prefrontal cortex. *J Neurosci* 26:7305–7316. [CrossRef Medline](#)
- Genovesio A, Tsujimoto S, Wise SP (2006b) Neuronal activity related to elapsed time in prefrontal cortex. *J Neurophysiol* 95:3281–3285. [CrossRef Medline](#)
- Genovesio A, Tsujimoto S, Wise SP (2008) Encoding problem-solving strategies in prefrontal cortex: activity during strategic errors. *Eur J Neurosci* 27:984–990. [CrossRef Medline](#)
- Genovesio A, Tsujimoto S, Wise SP (2009) Feature- and order-based timing representations in the frontal cortex. *Neuron* 63:254–266. [CrossRef Medline](#)
- Genovesio A, Tsujimoto S, Wise SP (2011) Prefrontal cortex activity during the discrimination of relative distance. *J Neurosci* 31:3968–3980. [CrossRef Medline](#)
- Genovesio A, Tsujimoto S, Wise SP (2012) Encoding goals but not abstract magnitude in the primate prefrontal cortex. *Neuron* 74:656–662. [CrossRef Medline](#)
- Histed MH, Pasupathy A, Miller EK (2009) Learning substrates in the primate prefrontal cortex and striatum: sustained activity related to successful actions. *Neuron* 63:244–253. [CrossRef Medline](#)
- Ito S, Stuphorn V, Brown JW, Schall JD (2003) Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* 302:120–122. [CrossRef Medline](#)
- Kennerley SW, Wallis JD (2009) Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci* 29:2061–2073. [CrossRef Medline](#)
- Kennerley SW, Walton ME (2011) Decision making and reward in frontal cortex: complementary evidence from neurophysiological and neuropsychological studies. *Behav Neurosci* 125:297–317. [CrossRef Medline](#)
- Kennerley SW, Behrens TE, Wallis JD (2011) Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci* 14:1581–1589. [CrossRef Medline](#)
- Kim S, Hwang J, Lee D (2008) Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron* 59:161–172. [CrossRef Medline](#)
- Lauwereyns J, Sakagami M, Tsutsui K, Kobayashi S, Koizumi M, Hikosaka O (2001) Responses to task-irrelevant visual features by primate prefrontal neurons. *J Neurophysiol* 86:2001–2010. [Medline](#)
- Mann SE, Thau R, Schiller PH (1988) Conditional task-related responses in monkey dorsomedial frontal cortex. *Exp Brain Res* 69:460–468. [Medline](#)
- Mansouri FA, Matsumoto K, Tanaka K (2006) Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin Card Sorting Test analog. *J Neurosci* 26:2745–2756. [CrossRef Medline](#)
- Owen AM, Evans AC, Petrides M (1996) Evidence for a two-stage model of spatial working memory processing within the lateral frontal cortex: a positron emission tomography study. *Cereb Cortex* 6:31–38. [CrossRef Medline](#)
- Passingham RE, Wise SP (2012) *The neurobiology of the prefrontal cortex: anatomy, evolution, and the origin of insight*. Oxford: Oxford UP.
- Petrides M (2005) Lateral prefrontal cortex: architectonic and functional organization. *Philos Trans R Soc Lond B Biol Sci* 360:781–795. [CrossRef Medline](#)
- Rainer G, Asaad WF, Miller EK (1998) Selective representation of relevant

- information by neurons in the primate prefrontal cortex. *Nature* 393:577–579. [CrossRef Medline](#)
- Rainer G, Rao SC, Miller EK (1999) Prospective coding for objects in primate prefrontal cortex. *J Neurosci* 19:5493–5505. [Medline](#)
- Rao SG, Williams GV, Goldman-Rakic PS (1999) Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: evidence for microcolumnar organization in PFC. *J Neurophysiol* 81:1903–1916. [Medline](#)
- Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. *Nature* 497:585–590. [CrossRef Medline](#)
- Roesch MR, Olson CR (2003) Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *J Neurophysiol* 90:1766–1789. [CrossRef Medline](#)
- Saito N, Mushiaki H, Sakamoto K, Itoyama Y, Tanji J (2005) Representation of immediate and final behavioral goals in the monkey prefrontal cortex during an instructed delay period. *Cereb Cortex* 15:1535–1546. [CrossRef Medline](#)
- Samejima K, Ueda Y, Doya K, Kimura M (2005) Representation of action-specific reward values in the striatum. *Science* 310:1337–1340. [CrossRef Medline](#)
- Schoenbaum G, Eichenbaum H (1995) Information coding in the rodent prefrontal cortex. 1. Single-neuron activity in orbitofrontal cortex compared with that in pyriform cortex. *J Neurophysiol* 74:733–750. [Medline](#)
- Schultz W (2007) Multiple dopamine functions at different time courses. *Annu Rev Neurosci* 30:259–288. [CrossRef Medline](#)
- Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J Neurosci* 27:8366–8377. [CrossRef Medline](#)
- Seo H, Barraclough DJ, Lee D (2007) Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cereb Cortex* 17:1110–1117. [CrossRef Medline](#)
- Seo H, Barraclough DJ, Lee D (2009) Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. *J Neurosci* 29:7278–7289. [CrossRef Medline](#)
- Sigala N, Kusunoki M, Nimmo-Smith I, Gaffan D, Duncan J (2008) Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proc Natl Acad Sci U S A* 105:11969–11974. [CrossRef Medline](#)
- Simmons JM, Richmond BJ (2008) Dynamic changes in representations of preceding and upcoming reward in monkey orbitofrontal cortex. *Cereb Cortex* 18:93–103. [CrossRef Medline](#)
- Takeda K, Funahashi S (2007) Relationship between prefrontal task-related activity and information flow during spatial working memory performance. *Cortex* 43:38–52. [CrossRef Medline](#)
- Tanji J, Hoshi E (2008) Role of the lateral prefrontal cortex in executive behavioral control. *Physiol Rev* 88:37–57. [CrossRef Medline](#)
- Tsujimoto S, Postle BR (2012) The prefrontal cortex and oculomotor delayed response: a reconsideration of the “mnemonic scotoma.” *J Cogn Neurosci* 24:627–635. [CrossRef Medline](#)
- Tsujimoto S, Sawaguchi T (2004) Neuronal representation of response-outcome in the primate prefrontal cortex. *Cereb Cortex* 14:47–55. [CrossRef Medline](#)
- Tsujimoto S, Sawaguchi T (2005) Context-dependent representation of response-outcome in monkey prefrontal neurons. *Cereb Cortex* 15:888–898. [CrossRef Medline](#)
- Tsujimoto S, Genovesio A, Wise SP (2008) Transient neuronal correlations underlying goal selection and maintenance in prefrontal cortex. *Cereb Cortex* 18:2748–2761. [CrossRef Medline](#)
- Tsujimoto S, Genovesio A, Wise SP (2012) Neuronal activity during a cued strategy task: comparison of dorsolateral, orbital, and polar prefrontal cortex. *J Neurosci* 32:11017–11031. [CrossRef Medline](#)
- Wallis JD, Miller EK (2003) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci* 18:2069–2081. [CrossRef Medline](#)
- Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411:953–956. [CrossRef Medline](#)
- Wirth S, Avsar E, Chiu CC, Sharma V, Smith AC, Brown E, Suzuki WA (2009) Trial outcome and associative learning signals in the monkey hippocampus. *Neuron* 61:930–940. [CrossRef Medline](#)
- Yamagata T, Nakayama Y, Tanji J, Hoshi E (2012) Distinct information representation and processing for goal-directed behavior in the dorsolateral and ventrolateral prefrontal cortex and the dorsal premotor cortex. *J Neurosci* 32:12934–12949. [CrossRef Medline](#)