

Prefrontal Neural Correlates of Memory for Sequences

Bruno B. Averbeck^{1,2} and Daeyeol Lee^{1,3}

¹Department of Brain and Cognitive Sciences, Center for Visual Science, University of Rochester, Rochester, New York 14627, ²Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom, and ³Department of Neurobiology, Yale University School of Medicine, New Haven, Connecticut 06510

The sequence of actions appropriate to solve a problem often needs to be discovered by trial and error and recalled in the future when faced with the same problem. Here, we show that when monkeys had to discover and then remember a sequence of decisions across trials, ensembles of prefrontal cortex neurons reflected the sequence of decisions the animal would make throughout the interval between trials. This signal could reflect either an explicit memory process or a sequence-planning process that begins far in advance of the actual sequence execution. This finding extended to error trials such that, when the neural activity during the intertrial interval specified the wrong sequence, the animal also attempted to execute an incorrect sequence. More specifically, we used a decoding analysis to predict the sequence the monkey was planning to execute at the end of the fore-period, just before sequence execution. When this analysis was applied to error trials, we were able to predict where in the sequence the error would occur, up to three movements into the future. This suggests that prefrontal neural activity can retain information about sequences between trials, and that regardless of whether information is remembered correctly or incorrectly, the prefrontal activity veridically reflects the animal's action plan.

Key words: sequence; ensemble; monkey; executive control; neurophysiology; prefrontal cortex

Introduction

Many behaviors in everyday life require the execution of sequential actions, and the correct sequence of actions normally has to be learned or selected through experience. The process of sequence selection, or sequence learning over a finite set, has been examined by a few researchers. We have shown recently that when animals have to discover which sequence is correct within a block of trials, the activity in dorsal lateral prefrontal cortex (dlPFC) predicts closely the fraction of movements made by the animal that are consistent with the new sequence, as well as the fraction of movements that are consistent with the sequence that had been correct in the previous block (Averbeck et al., 2006b). Studies by other laboratories have shown that activity in the medial frontal cortex is different for well learned versus novel sequences (Nakamura et al., 1998), and that activity in the cingulate cortex is different depending on whether animals are selecting which sequence is correct in the current block, or repeating the appropriate sequence after it has been selected (Procyk et al., 2000). Thus, these studies of sequence learning have demonstrated that, when animals learn sequences, frontal cortical areas might be involved in updating information about the correct movements. These previous studies, however, leave open two important questions. First, how are sequences remembered across trials, and second, is dlPFC involved in the executive con-

trol of sequence behavior, or is it mostly playing a role in action monitoring?

Although short-term or working memory is normally studied with paradigms that require the remembrance of items within a trial (Funahashi et al., 1991), remembering items across trials is also important. Remembering the sequence of movements that is correct within a block of trials is similar to, although clearly different from, remembering task set, and as such our task shares elements with experiments requiring animals to remember arbitrary stimulus-response mappings (Chen and Wise, 1995a,b, 1996; Asaad et al., 1998; Pasupathy and Miller, 2005), or the categorization rule in the Wisconsin card sorting task (Mansouri et al., 2006). Where in the brain then is the signal that tracks which sequence is correct in the current block? In the present study, we show that the activity of dlPFC neurons during the intertrial interval carries information about the correct sequence in a particular context. We also show that prefrontal cortex is closely involved in the executive control of sequential actions. More specifically, we found that the sequence represented in the activity of ensembles of dlPFC neurons, before the actual execution of the sequence, predicted the sequence the animal executed. However, our findings do not distinguish between memory and planning processes, in that a plan held on-line throughout the intertrial interval would resemble a memory trace, and either mechanism would fulfill the role of maintaining the relevant information between trials. Additionally, we found that the sequence represented in dlPFC predicted the sequence the monkey would execute not only in correct trials but also in error trials, such that when the animal was planning the wrong sequence, it also tried to execute the wrong sequence. This shows that there is a strong correspondence between prefrontal activity and the animal's action plans.

Received Oct. 13, 2006; revised Dec. 20, 2006; accepted Jan. 16, 2007.

This work was supported by National Institutes of Health Grants R01-MH59216, T32-MH19942, and P30-EY01319. We are grateful to Jeong-woo Sohn for his help with the experiment, Dominic Barraclough for his help with surgeries, Lindsay Carr for her technical assistance, and John Swan-Stone for programming.

Correspondence should be addressed to Dr. Bruno B. Averbeck, University College London Institute of Neurology, Sobell Department, Box 28, Queen Square, London WC1N 3BG, UK. E-mail: b.averbeck@ion.ucl.ac.uk.

DOI:10.1523/JNEUROSCI.4483-06.2007

Copyright © 2007 Society for Neuroscience 0270-6474/07/272204-08\$15.00/0

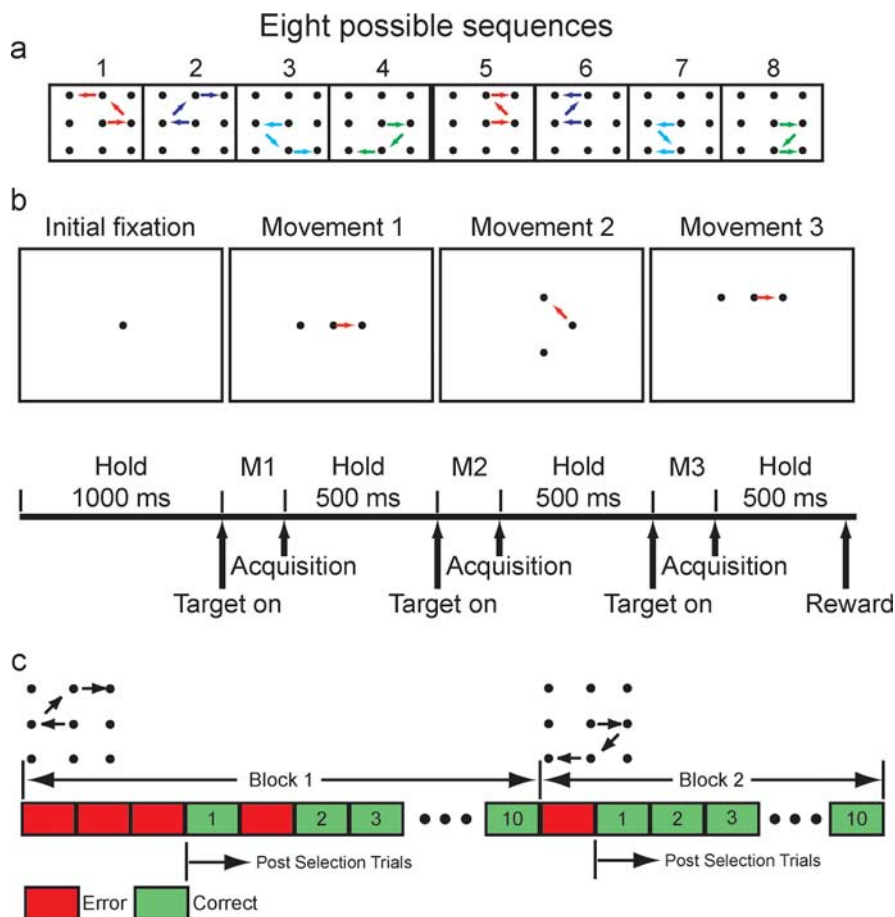


Figure 1. Task. *a*, The eight possible sequences the monkeys were trained to execute. Each panel indicates one sequence. *b*, Temporal sequence of choices in a single trial. The dot in the center of the initial fixation frame indicates the central fixation point. The two dots to the left and right of the central dot in the Movement 1 frame indicate the first two choice targets from which the animal can select. Similarly, the two dots above and below and to the left and right of the fixation in the Movement 2 and Movement 3 frames indicate the possible targets for the saccade at the corresponding points in the sequence. *c*, Example sequence of trials from two blocks. The data analyzed in this manuscript are the first correct trial and all subsequent trials in each block, indicated as post-selection trials in the figure.

Materials and Methods

General. Two male rhesus macaques were used in this study. All surgical and experimental procedures conformed to the National Institutes of Health guidelines and were approved by the University of Rochester Committee on Animal Research. The recording chamber (18 mm diameter) was placed over the dorsolateral prefrontal cortex in a sterile surgery using stereotaxic coordinates derived from structural magnetic resonance imaging (MRI). Neural activity was recorded using a 16-channel multielectrode recording system (Thomas Recording, Marburg Germany), and single-unit spikes were sorted on-line using the Plexon (Dallas, TX) data acquisition system. The task was presented to the animals on a CRT monitor, and a custom windows-based program was written to control the task and coordinate data acquisition with the Plexon system. Eye movements were monitored using a video eye-tracking system (ET-49; Thomas Recording). The location of the frontal eye field (FEF) was verified in both animals using microstimulation. Electrode penetrations were considered within the FEF when a 50 μ A peak-to-peak bipolar current elicited an eye movement at least 50% of the time. The location of the FEF found using microstimulation corresponded to its predicted location in the chamber derived from MRI coordinates. All penetrations were anterior to the FEF except one penetration in each animal.

Behavioral task. The animals were trained on a sequential decision-making task (Fig. 1). In this task, there were eight possible correct sequences of eye movements presented on a 3 \times 3 grid of targets spaced by 5.3° of visual angle (Fig. 1*a*), and a single trial was composed of a se-

quence of at least three movements. We distinguish between trials and movements through out.

The animals began a trial by acquiring a central fixation point (Fig. 1*b*). After a 1 s fore-period, two targets were presented to the left and right of fixation, and the animal was allowed to make a saccade as soon as the targets appeared. Within a block of trials, one of the targets was correct at each stage of the sequence. If the animal made a saccade to the correct target and maintained fixation for 500 ms, the next two choice targets in the sequence were presented. This was repeated until the animal selected three correct targets, reaching the end of the sequence, at which point the trial ended and the animal was given a juice reward. If they chose the wrong target at any point in the sequence, they were forced back to their previous fixation point, and they were shown the two choice targets again. This was repeated until they selected the correct target. If they completed the sequence, even if they had selected a wrong target at some point, they were given a juice reward. Thus, the monkeys always had incentive to finish the sequence. However, the trial was only counted as correct for the analyses if they completed the sequence without selecting any of the wrong targets. The correct sequence remained fixed for a block of 10 correct trials. After 10 trials were completed successfully, not necessarily consecutively, a new sequence was introduced. The change in the sequence was not cued, and thus the animals only discovered the change in the sequence after selecting a target that had been correct in the previous block, and being forced back to the previous fixation target and given the choice targets again. At this point, the animals had to work out the new correct sequence by trial and error.

The total duration of a trial was constrained to be <7 s, but this limit was rarely reached.

The sequences were presented in a randomized block design, such that one block of 10 trials had to be executed for each of eight sequences before the same sequence was presented again. Data from a recording session were included in the analysis only if at least two blocks were completed for each sequence. Animals were trained for 3–5 months on the task before recordings began.

Data analysis. The present study focused on the neural correlates of errors that were caused by the animal’s failure to maintain information about the correct sequence. Therefore, we excluded the data from the trials before the animals selected the correct sequence and analyzed only the postselection trials, which were the first correct trial and subsequent trials from a given block (Fig. 1*c*). Furthermore, because we performed extensive analyses of error trials, we only included the data from sessions with at least 15 postselection error trials across all sequences and blocks. We also excluded all error trials in which the animal failed to acquire one of the choice targets because of an inaccurate saccade. Therefore, all errors analyzed are those in which the animal made saccades to incorrect targets.

The ANOVA results given in this paper were all assessed using type-III sums of squares, within a general linear model framework, because we had an unequal number of data points (i.e., an unbalanced design) for all comparisons. All analyses were performed on binned spike counts. No prior smoothing was done.

We used a Gaussian decoding analysis as described in detail previously (Averbeck et al., 2003b). This analysis assumes that the distribution of

neural responses in a time bin is Gaussian for a particular sequence. Thus, we have a Gaussian likelihood function given by the following:

$$p(r|s = i) = |2\pi Q|^{-1/2} \exp\left[-\frac{1}{2}(r - \mu_i)^T Q^{-1}(r - \mu_i)\right], \quad (1)$$

where r is a vector of spike rates for a sequence, μ_i is the vector of mean spike counts for sequence i , the superscript T indicates transpose, Q is the noise covariance matrix pooled across conditions, and $||$ indicates the determinant of the matrix. For the present analysis, i takes on values between 1 and 8, because there are eight sequences.

The posterior probability that a particular sequence led to the neural activity under consideration is given by Bayes' theorem, as follows:

$$p(s|r) = \frac{p(r|s)p(s)}{p(r)}, \quad (2)$$

where r is one of the sequences. In this study, we assumed a flat prior, and thus $p(s)$ is a constant. The normalization is given by the following:

$$p(r) = \sum_s p(r|s)p(s). \quad (3)$$

The decoding analysis is performed by first estimating the average response vector, μ_p , for each condition to be decoded (for example, for each of the eight sequences in the sequence-decoding analysis). Each element of this vector is the average response of one of the simultaneously recorded cells to one of the sequences. Then, for each trial, this average response is subtracted from the actual response to get the residual or noise for that trial for each neuron. These noise values are then used to calculate the pooled noise covariance matrix, Q across all conditions to be decoded. We then generate a separate Equation 1 for each condition. They all have the same Q , but each has a different μ_p , one for each condition to be decoded. To perform the analysis, the response on an individual trial, r , is plugged into each Equation 1. This gives the likelihood that the response under consideration was generated by the condition that corresponds to the particular equation. Each of these likelihood values is then plugged into Equation 2, giving the posterior probability that the condition to be decoded gave rise to the response we are considering, where the normalization factor in Equation 3 is the sum of these likelihood values. In the analyses where we consider the log of the posterior, the posterior is given directly by the left side of Equation 2. To perform classification explicitly, sequences were predicted by selecting the sequence with the maximum probability from the conditional distribution given the neural activity as follows:

$$\hat{s} = \arg \max_s p(s|r). \quad (4)$$

Because we used a flat prior, decoding with either maximum likelihood estimation (i.e., picking the stimulus that maximizes Eq. 1) or maximum a posteriori estimation (Eq. 4) gave the same results.

The decoding analyses presented in Figure 4*b* were performed using twofold cross-validation for both the correct and error trials. Importantly, activity in none of the time bins for the correct trials were classified against a model that was estimated using the same trial, although we were analyzing different bins. Including the same trial when estimating the decoding model inflates the estimate of percentage correct classification performance for the correct trial condition, presumably because of slow drift in neural activity. For the analyses in Figure 4*b*, we used 300 ms bins with an interbin interval of 150 ms. For the analysis used to generate Table 1, we used spike counts in 500 ms bins to improve signal-to-noise ratio, because we were not interested in the time course of the planning.

Results

Behavioral performance and neural database

The reaction time and error performance of the animals as a function of where they were in the block of trials is shown in Figure 2. One of the animals was slightly faster and more accurate. The fast saccadic reaction times, which varied little across the block, were attributable to the fact that the timing of the onset

Table 1. Predicted movement on which error occurs

Measured/null	Predicted movement 1	Predicted movement 2	Predicted movement 3
Actual movement 1	1034/1003	433/470	263/257
Actual movement 2	306/321	181/151	67/82
Actual movement 3	112/128	67/60	42/33

Numbers represent measured number of errors/the number predicted by the null hypothesis ($\chi^2 = 19.14$; df, 4; $p < 0.001$).

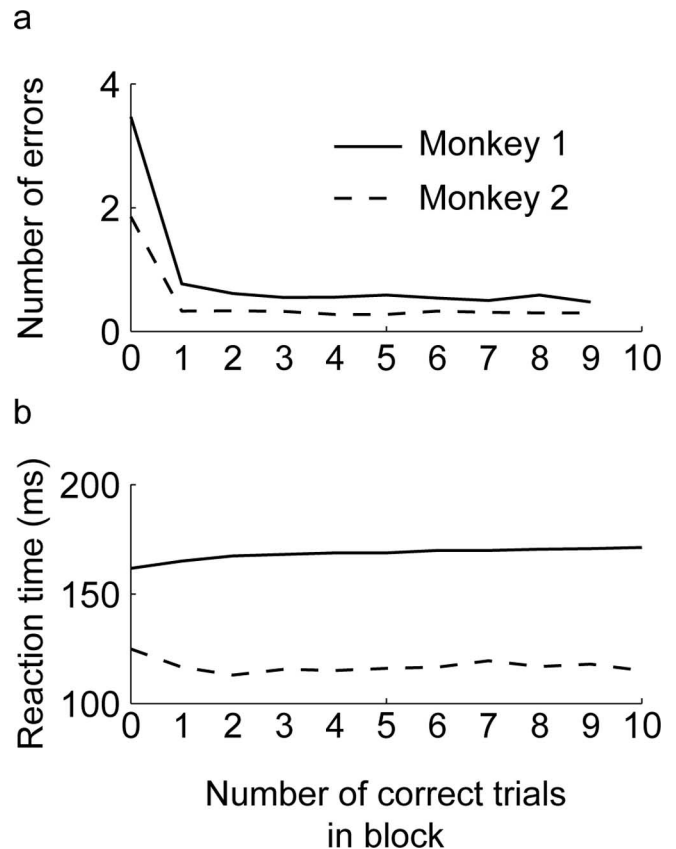


Figure 2. Behavioral performance. *a*, Error rate as a function of the number of correct trials in the block. The vertical axis represents the average number of errors that were committed before getting a trial correct, as a function of the number of correct trials. Trial 0 is all trials before one correct and presumably represents the period where the animal is working out the correct sequence. *b*, Mean reaction times plotted as a function of the number of correct trials in the block.

of the targets was completely predictable. We analyzed the activity of 442 neurons in 60 ensembles that were recorded while animals performed the oculomotor sequence task.

Sequence activity during the intertrial interval

Accurate performance in our task required the animals to remember, or maintain a plan of, which sequence of decisions had been correct in the previous trial. Thus, this information must be maintained somewhere in the brain. Our first analysis searched for evidence of this information in the activity of dIPFC neurons. There were single neurons that changed their activity as a function of which sequence was correct in the current block. This difference was present throughout the intertrial interval, as well as during fixation before the animal began producing the sequence of eye movements (Fig. 3). In the following, the time interval that includes both the intertrial interval and fore-period is referred to as the presequence period. A two-way ANOVA

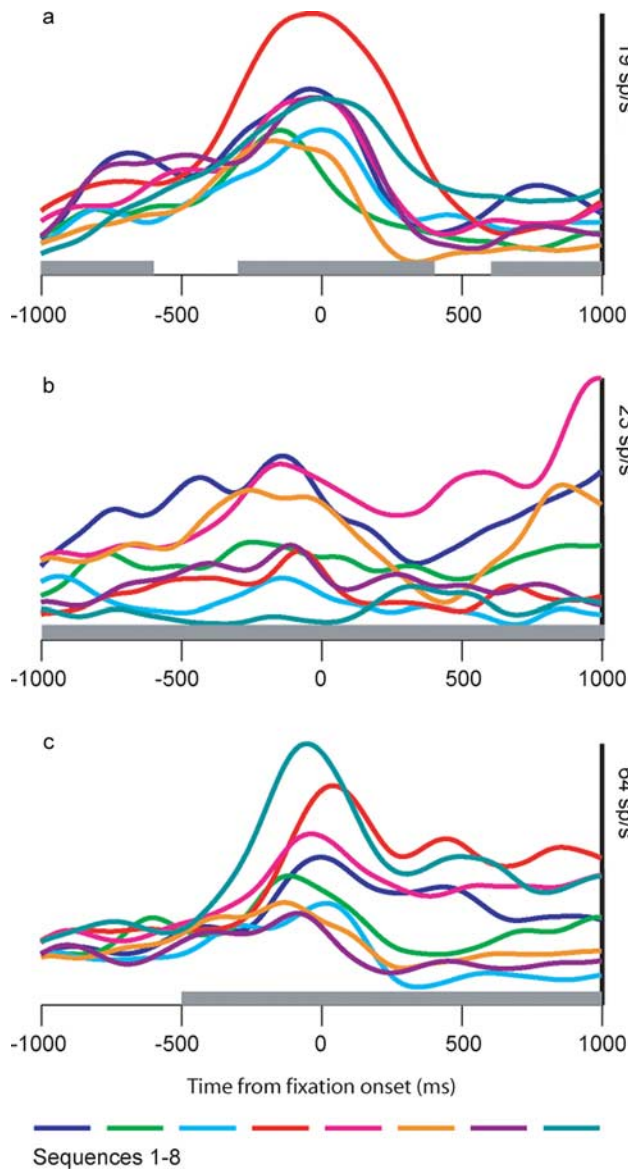


Figure 3. Spike density functions for three example neurons, showing different activity patterns for different sequences during intertrial interval and fore-period. The intertrial interval is 1 s, and thus -1000 ms approximately corresponds to the end of the previous trial. Time 0 is fixation onset. At 1000 ms, the first pair of targets for the sequence is presented. The gray bar at the bottom of each plot indicates whether there was a significant main effect ($p < 0.05$) of sequence in a two-way ANOVA (300 ms bins) with sequence and first movement direction as factors.

(type III sums of squares to control for unbalanced design) was performed on spike counts in a 300 ms sliding window with steps of 100 ms, with sequence and first movement direction as factors. The first movement factor controls for simple movement planning effects. The data were aligned to fixation onset (time 0), and the reaction time to acquire the initial fixation target after it appeared averaged 365 ms (SEM, 2.69; $n = 20,307$). All correct trials were included in the analysis, regardless of whether they were preceded by correct trials. The results showed that, of the neurons with a significant main effect of sequence at the end of the fore-period (87 of 442; 20%), $\sim 40\%$ also had a significant main effect during the intertrial interval (Fig. 4a). Almost all neurons that had a significant effect during the intertrial interval also showed an effect at the end of the fore-period. For the entire population of recorded neurons, the number of significant neurons never




Figure 4. Representation of sequence during intertrial interval and fore-period. **a**, Percentage of individual neurons that showed a significant effect of sequence during the intertrial interval and the fore-period, expressed as a fraction of those that had a significant effect at the end of the fore-period. The time values indicate the start of the 300 ms bin. Time 0 is the bin beginning at fixation onset, and -1000 is approximately the end of the previous trial. **b**, Population average of percentage correct classification performance for correct and error trials. The vertical axis is percentage correct classification. **c**, $\ln(p)$ from likelihood ratio test of difference between correct and error curve shown in **b**. The dashed line indicates $\ln(0.05)$.

dropped below the number expected by chance ($p < 0.05$, binomial test; minimum of 8% at -900 ms). It can also be seen that this signal about sequence started to increase rapidly just before fixation onset (bin starting at -300 ms) and continued to increase in strength during the hold period, before the beginning of sequence execution. Thus, a small but significant portion of the frontal network maintained information about the correct sequence in the current block between trials, and before execution of the sequence a larger portion of the network began to represent this information.

Although the ANOVA analysis showed that information is present in the activity of single neurons about the correct sequence in the current block, it does not provide detailed information about the fidelity of this signal. Furthermore, because the ANOVA is run separately on each time bin, it does not tell us whether the sequence information tends to be coded in the same way across the population during this time. Specifically, does the

pattern of activity across an ensemble of simultaneously recorded neurons that is related to a particular sequence change or remain the same throughout the fore-period? Although it can be seen from Fig. 3 that there are cases where the information is not coded consistently (Fig. 3*a*), in other cases, it is more consistent (Fig. 3*c*) (0–1000 ms). To assess the strength of the representation of the sequence and the consistency of the pattern in small ensembles of neurons, we performed a decoding analysis. We first defined a decoding model for each sequence (see Materials and Methods) using the neural activity in the final 300 ms bin of the fore-period from all of the correct trials in a given recording session. We included 141 neurons (average ensemble size, 2.4) in this analysis. Neurons were included if they had a significant main effect of sequence at the end of the fore-period, not controlling for first movement, as we did in the analysis above, which is why the number of neurons is larger in this analysis than the number used in the above ANOVA ($n = 87$). We then classified the activity in each bin of the presequence period with respect to this decoding model. If the pattern of activity in an ensemble that represents a sequence changes during the presequence period, it will not be classified correctly by the decoding analysis, because this analysis assumes the pattern in an ensemble related to a specific sequence and remains unchanged until the end of the fixation period. This analysis showed that the average strength of the sequence representation in small ensembles during the intertrial interval was just above chance, and grew stronger, peaking just before sequence initiation (Fig. 4*b*). The relatively low classification performance is a result of the small size of the ensembles (2.4 on average), and the large number of categories (eight) we were trying to decode. The same analysis was also performed on error trials in which the monkeys selected an incorrect target at some point in the sequence. The results showed that, although the decoding performance at the beginning of the intertrial interval was essentially the same for both error and correct trials, the representation of the correct sequence did not become stronger in error trials and remained only slightly above chance throughout the presequence period. Thus, ~ 200 ms before fixation onset (bin starting at -500 ms), signals in the correct and error trials began to diverge, and the difference reached significance ($p < 0.05$; likelihood ratio test) around the time of fixation onset (Fig. 4*c*). This suggests that when the animal was going to make a mistake in the subsequent trial, it did not properly maintain the information about the correct sequence, and this was reflected in the dlPFC activity just before fixation onset.

We carried this analysis one step further by asking whether or not we could find evidence in the trial preceding the error trial that the animal would make a mistake in the next trial. This might happen if the animal's attention to the task was already beginning to drift in the trial preceding the error trial. To answer this question, we performed a one-way ANOVA, with respect to whether the next trial would be an error, on the log-posterior transformed neural activity (see Materials and Methods) during movement execution from correct trials. The log-posterior transformation measures how probable it is that neural activity came from the distribution of interest, and as such it measures how similar the neural activity is to the average neural activity from a particular task condition. Thus, a large negative log-posterior indicates that the neural activity is very different from the average activity for a particular condition. Because the log-posterior easily handles activity in small ensembles of neurons and takes into account differences in the mean as well as variability in neural activity, we can use it to pool data across movement directions as well as recording sessions. The ANOVA showed that movement-related neural




Figure 5. Spike density function examples of sequence misplanning. This is the same neuron as that shown in Figure 3*c*. The colored lines indicate the average spike density functions for the indicated sequences (same as in Fig. 3*c*). The black line indicates the spike density function for the individual trial in which the sequence was misplanned. At the left is a case in which the monkey was supposed to execute sequence 8 (blue line), but it was planning sequence 7 (purple line). At the right is a case in which the monkey was supposed to execute sequence 4 (red line), but instead it was planning sequence 2 (green line). In both cases, the single-trial response was more similar to the sequence the monkey was planning than to the sequence the monkey was supposed to execute.

activity in correct trials was not significantly different ($p > 0.05$), depending on whether the trial was followed by a correct or error trial, suggesting that the difference in neural activity between correct and error trials did not arise until the intertrial interval. This is consistent with the decoding analysis, which showed that the neural activity did not differ between correct and error trials until well into the intertrial interval.

Errors in sequence execution

Errors in sequence execution could occur for two reasons. First, the animals could plan the wrong sequence, having forgotten the correct sequence, or second, they could plan the correct sequence but execute it incorrectly by, for example, forgetting one of the movements within the sequence. For an example of incorrect sequence planning, in which the incorrect sequence the monkey was planning was known, we can examine the response of a single neuron in the first trial of a new block after the sequence had switched but before the monkey knew it had switched. In this case, it can be seen that the neural response was more similar to the neural response for the sequence from the previous block, than it was to the neural response for the sequence from the current block, that the monkey was supposed to execute (Fig. 5), because the animal did not yet know the sequence had switched.

To determine, at a population level, whether errors generally occurred because of incorrect sequence planning or incorrect execution of correct sequence plans, we performed a decoding analysis in which we predicted the animal's sequence plan at the end of the fore-period, using the ensemble neural activity (see Materials and Methods). From this analysis, error trials were divided into those in which the animals planned the correct sequence (440 trials) and those in which the animals planned the incorrect sequence (2505 trials), where we took the sequence predicted by the decoding analysis as the sequence being planned by the animal. Thus, a majority of errors were caused by planning the sequence incorrectly. Next, to test further the hypothesis that neural activity predicting the wrong sequence at the end of the hold period reflected an incorrect action plan, we predicted the specific movement in the sequence on which the error would occur, based on the first movement in the sequence that differed between the decoded sequence (i.e., the alleged planned sequence) and the sequence the animal was supposed to execute (i.e., correct sequence). For example, if the animal was supposed to execute sequence 1 (Fig. 1*a*), and it actually planned and executed sequence 2, there would be a mistake on the first movement. If, however, the animal planned and executed sequence 5, a mistake would occur on the last movement. After predicting

where the error would occur, we compiled a contingency table (Table 1), where each location in the table was defined by the movement on which we predicted the animal would make a mistake (columns) and the movement on which the animal actually made a mistake (rows). By compiling the data across all of the error trials in which the activity did not predict the correct sequence, we found that we could predict significantly where the mistake would occur in the sequence ($\chi^2 = 19.14$; $df = 4$; $p < 0.001$). In this table, the numbers on the left of the slash in the diagonal boxes are the cases we predicted correctly. The numbers on the right are the number of times we would expect to predict correctly based only on knowing how many times the errors occurred on each movement, that is to say, by guessing optimally under the assumption that there is no relationship between the neural data and the behavior. Thus, when the number on the left of the slash (/) is higher than the number on the right for the diagonal entries, we are doing better than chance, where chance is indicated by the number on the right. As can be seen, errors were predicted above chance level even for the third movement. Thus, in a statistically significant number of error trials, prefrontal activity represented the incorrect sequence, and this error in sequence planning led to the predicted mistake in sequence execution.

We further characterized the incorrect action plan by asking whether in error trials the animals were more likely to plan the sequence that had been correct in the previous block. We found that the probability that the neural activity in incorrect trials would predict the sequence from the previous block was 0.143, and that this was significantly greater than chance, which would be 0.125 (binomial test, $p < 0.01$; $n = 2505$). Thus, there is a weak but significant tendency for mistakes in action plans to be reversions to the action plan that had been correct in the previous block, possibly because of an incomplete suppression of the action plan of the previous block (Averbeck et al., 2006b).

To understand the predictions about neural activity during sequence execution in error trials, we note that neural activity in dlPFC during sequence execution is different for the same movement depending on the sequence in which the movement is embedded (Averbeck et al., 2006b). Thus, the same movements executed under different action plans should result in different neural activity. This implies that when the animal proceeded according to the incorrect sequence in error trials, the neural activity during the correct movements preceding the first mistake should be different from the neural activity during the same movements in error trials when the animal proceeded according to the correct sequence plan. To test this, we compared the log-posterior of the perisaccadic neural activity among the following: (1) movements from correct trials, (2) correct movements from error trials with a correctly planned sequence, and (3) correct movements from error trials with an incorrectly planned sequence, where the sequence plan in error trials was determined using the decoding analysis. Because the log-posterior decreases as the neural activity in a given trial deviates from the average neural activity in correct trials, it should be smallest in error trials with an incorrect action plan. Importantly, in all cases, we compared neural activity across the same movements, executed according to different sequence plans. As predicted, we found that the log posterior from error trials with an incorrect action plan was smallest, whereas the log posterior from error trials with a correct action plan was more similar to neural activity from correct trials (Fig. 6). A two-way ANOVA, with main effects of movement number and trial type (correct, error with correct plan, error with incorrect plan) showed a significant main effect

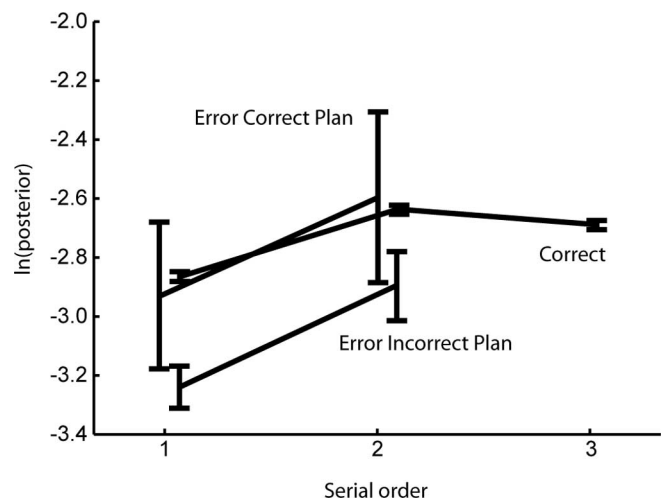


Figure 6. Average and SEM of log posterior of neural activity, for the three movements of the sequence, and the three trial types. Correct, Correct trials; Error Correct Plan, error trials with a correct action plan; Error Incorrect Plan, error trials with an incorrect action plan.

of trial type. *Post hoc* analyses showed that only correct trials and error trials with an incorrect plan were significantly different ($p < 0.05$; Tukey's honestly significant difference test). Log-posteriors in error trials with a correctly planned sequence were not significantly different from those in either of the other trial types, although their means were more similar to the correct trials. These results show that prefrontal neurons carry a signal that reflects the animal's action plans, before and during sequence execution. Furthermore, they suggest that errors are of two types. Either the animal can generate the incorrect sequence plan, or it can generate the correct sequence plan but fail to implement it correctly. The most common type of error, however, was an error in planning the sequence.

Discussion

We have shown that dlPFC carries a signal, during the intertrial interval and the pretrial hold period, related to the correct sequence in the current block of trials. This signal is critical for correct task performance, because the sequence to be executed is not cued explicitly and it has to be remembered across trials in a block. This signal could, however, be related to sequence planning instead of sequence memory. Our task cannot disentangle these two possibilities. We also found that, when the animal made mistakes during the execution of the sequence, the point in the sequence where the mistake occurred could be predicted by the prefrontal neural activity observed before the animal began executing the sequence. Finally, we found that most mistakes resulted from errors in action plan or memory rather than from erroneous execution of correct action plans.

Sequence tasks used in previous studies have required animals to learn a sequence either through a series of visually cued trials (Shima and Tanji, 2000) or by trial-and-error (Procyk et al., 2000; Lu et al., 2002) and then execute the sequence from memory. However, these previous studies have not demonstrated a signal present throughout the interval between trials that indicated which sequence was correct in the current block. It is possible that the sequence memory signal described in this study is confined to dlPFC as opposed to the medial cortical areas often examined in previous studies, including the supplementary eye fields (Lu et al., 2002), the supplementary motor area (Shima and Tanji, 2000), and the anterior cingulate cortex (Procyk et al., 2000). This

would be consistent with the known role of prefrontal cortex in working memory (Funahashi and Kubota, 1994), but it needs to be tested directly in future studies. Previous studies have also shown that, when animals have to make decisions based on a behavioral rule relevant in a block of trials, dlPFC neurons can represent the rule during the fixation interval before the beginning of the trial (Asaad et al., 2000) or during the intertrial interval (Mansouri et al., 2006). Our data show that dlPFC activity is not only relevant for remembering which behavioral rule is active within a block of trials but also for remembering which sequence of movements is correct within a block of trials.

We have also shown that dlPFC neural activity can predict mistakes up to three movements into the future. This finding has several implications. First, it shows that dlPFC activity is closely related to the behavior actually executed by the animal, suggesting a role in executive control. The close link between dlPFC activity and behavior is consistent with recent findings in studies of stimulus-driven decision making, which show that dlPFC activity correlates with the animal's decision (Kim and Shadlen, 1999; de Lafuente and Romo, 2005), whereas primary sensory areas correlate with the sensory stimulus and are little affected by the decision (Britten et al., 1992; de Lafuente and Romo, 2005). Other studies have shown that the parietal cortex, which is anatomically linked with dlPFC (Cavada and Goldman-Rakic, 1989), carries a signal more closely linked to the animal's behavior than to the stimuli that drive that behavior (Chafee et al., 2005), whereas it has been shown that the inferotemporal cortex can compute the correct answer in a task, even when the animal responds incorrectly (Messinger et al., 2005). These data suggest a hierarchical structure for computations in decision-making tasks, with sensory areas having accurate representations of stimuli and parietal and frontal areas having accurate representations of decisions. An interesting question from this perspective is whether another area, for example the hippocampus, might carry a more accurate memory of the sequence, or whether dlPFC is responsible for working memory as well as implementing the action plan. These previous studies also suggest that mistakes in task performance may not always be attributable to noise in sensory representations, but rather they may come about through computations performed by structures downstream from sensory representations, which may be suboptimal in some respect (Averbeck et al., 2006a). Accordingly, these findings are at variance with the hypothesis that noise in visual motor processing is 93% sensory (Osborne et al., 2005). There are also interesting computational links between attractors networks for short-term memory (Compte et al., 2000), and networks that can perform optimal Bayesian inference for extracting information from upstream structures (Pouget et al., 1998; Wang, 2002), which suggests that prefrontal cortex might have the computational and anatomical architecture for both processes. Wherever these computations take place, dlPFC ultimately carries a signal related to the output of those computations, because it correlates more strongly with behavior than with the sensory stimuli driving behavior. Understanding why some cortical areas contain veridical estimates of the outside world, whereas other areas contain veridical predictors of behavior, and why the computations linking these processes sometimes break down, will be an important topic for future study.

The second implication of this result is that dlPFC activity is important for sequence planning, because activity before sequence implementation is related to the sequence the animal ultimately executes. This is consistent with neuropsychological studies, which have shown that damage to prefrontal cortex that

would decrease the ability of the network to correctly plan sequences can cause problems with implementing sequences of actions (Shallice, 1982). Previous single-unit studies in dlPFC have also shown sequence-planning activity before sequence execution (Averbeck et al., 2002; Mushiakhe et al., 2006). However, these studies did not address errors in sequence planning, and sequences did not have to be recalled from memory, because they were cued by sensory stimuli. Our data suggest that, in most cases, these errors in sequence planning arise through planning to execute the wrong sequence of movements, as opposed to incorrectly executing a correctly planned sequence of actions. In a less constrained behavioral setting, this mis-remembering could result in the execution of behaviors that appear to have little to do with the current task.

One potential confound in our findings was that we did not monitor eye movements during the intertrial interval, and therefore, we could not test directly the possibility that the sequence-dependent activity was caused by differential eye movements during this period. However, the sequence-selective neural activity smoothly increased from just before fixation onset (−400 ms) into the fixation interval (Fig. 4a), while at the same time the unconstrained eye movements of the intertrial interval changed dramatically to fixation. Thus, we believe that the sequence-selective activity during the intertrial interval is a reflection of a mental strategy, likely memory, to maintain the sequence information across trials.

It is also useful to make a brief point about the log-posterior values plotted in Figure 6, which were also reported in our previous study (Averbeck et al., 2006b). The average values are quite low because of the way the analysis was performed and not because the effect we are analyzing is small. Specifically, these are logs of the posterior probabilities in a decoding analysis with 24 possible outcomes. Although posterior probabilities can be quite high (as well as classification accuracy) when small neural ensembles are being used to predict only two or, at most, a few directions (Averbeck et al., 2003b; Averbeck and Lee, 2006), classification accuracy drops quickly when a large number of possibilities are being predicted (Averbeck et al., 2003b). The responses of the neurons in our task tend to be rather sparse, such that they only respond to a few different movements across the 24 possible movements (eight sequences, three movements each). Thus, a randomly selected ensemble will only distinguish among a subset of the movements effectively, and when the other movements are being decoded, the ensemble will be at almost chance performance. When we average across all movements, the effects look like they are just above chance. However, as can be seen easily from the ANOVA and decoding analyses presented by Averbeck et al. (2006b), the effects we are reporting are highly robust.

Our results demonstrate that prefrontal activity carries signals relevant for remembering sequences of actions within a block of trials, and that when these signals were not properly represented in prefrontal cortex, the animal executed the wrong sequence. Therefore, it seems that dlPFC is an important node in the network of areas responsible for the selection, planning, and correct execution of sequences of actions.

References

- Asaad WF, Rainer G, Miller EK (1998) Neural activity in the primate prefrontal cortex during associative learning. *Neuron* 21:1399–1407.
- Asaad WF, Rainer G, Miller EK (2000) Task-specific neural activity in the primate prefrontal cortex. *J Neurophysiol* 84:451–459.
- Averbeck BB, Lee D (2006) Effects of noise correlations on information encoding and decoding. *J Neurophysiol* 95:3633–3644.
- Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP (2002) Parallel

- processing of serial movements in prefrontal cortex. *Proc Natl Acad Sci USA* 99:13172–13177.
- Averbeck BB, Chafee MV, Crowe DA, Georgopoulos AP (2003a) Neural activity in prefrontal cortex during copying geometrical shapes I. Single cells encode shape, sequence, and metric parameters. *Exp Brain Res* 150:127–141.
- Averbeck BB, Crowe DA, Chafee MV, Georgopoulos AP (2003b) Neural activity in prefrontal cortex during copying geometrical shapes II. Decoding shape segments from neural ensembles. *Exp Brain Res* 150:142–153.
- Averbeck BB, Latham PE, Pouget A (2006a) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
- Averbeck BB, Sohn JW, Lee D (2006b) Activity in prefrontal cortex during dynamic selection of action sequences. *Nat Neurosci* 9:276–282.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Cavada C, Goldman-Rakic PS (1989) Posterior parietal cortex in rhesus monkey: II. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *J Comp Neurol* 287:422–445.
- Chafee MV, Crowe DA, Averbeck BB, Georgopoulos AP (2005) Neural correlates of spatial judgement during object construction in parietal cortex. *Cereb Cortex* 15:1393–1413.
- Chen LL, Wise SP (1995a) Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurophysiol* 73:1101–1121.
- Chen LL, Wise SP (1995b) Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations. *J Neurophysiol* 73:1122–1134.
- Chen LL, Wise SP (1996) Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations. *J Neurosci* 16:3067–3081.
- 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.
- de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. *Nat Neurosci* 8:1698–1703.
- Funahashi S, Kubota K (1994) Working memory and prefrontal cortex. *Neurosci Res* 21:1–11.
- Funahashi S, Bruce CJ, Goldman-Rakic PS (1991) Neuronal activity related to saccadic eye movements in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 65:1464–1483.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Lu X, Matsuzawa M, Hikosaka O (2002) A neural correlate of oculomotor sequences in supplementary eye field. *Neuron* 34:317–325.
- 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.
- Messinger A, Squire LR, Zola SM, Albright TD (2005) Neural correlates of knowledge: stable representation of stimulus associations across variations in behavioral performance. *Neuron* 48:359–371.
- Mushiake H, Saito N, Sakamoto K, Itoyama Y, Tanji J (2006) Activity in the lateral prefrontal cortex reflects multiple steps of future events in action plans. *Neuron* 50:631–641.
- Nakamura K, Sakai K, Hikosaka O (1998) Neuronal activity in medial frontal cortex during learning of sequential procedures. *J Neurophysiol* 80:2671–2687.
- Osborne LC, Lisberger SG, Bialek W (2005) A sensory source for motor variation. *Nature* 437:412–416.
- Pasupathy A, Miller EK (2005) Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433:873–876.
- Pouget A, Zhang K, Deneve S, Latham PE (1998) Statistically efficient estimation using population coding. *Neural Comput* 10:373–401.
- Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* 3:502–508.
- Shallice T (1982) Specific impairments of planning. *Philos Trans R Soc Lond B Biol Sci* 298:199–209.
- Shima K, Tanji J (2000) Neuronal activity in the supplementary and pre-supplementary motor areas for temporal organization of multiple movements. *J Neurophysiol* 84:2148–2160.
- Wang XJ (2002) Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36:955–968.