nature neuroscience

Activity in prefrontal cortex during dynamic selection of action sequences

Bruno B Averbeck, Jeong-Woo Sohn & Daeyeol Lee

Completing everyday tasks often requires the execution of action sequences matched to a particular problem. To study the neural processes underlying these behaviors, we trained monkeys to produce a series of eye movements according to a sequence that changed unpredictably from one block of trials to the next. We then applied a decoding algorithm to estimate which sequence was being represented by the ensemble activity in prefrontal cortex. We found that the sequence predicted by this analysis changed gradually from the sequence that had been correct in the previous block to the sequence that was correct in the current block, closely following the fraction of executed movements that were consistent with the corresponding sequence. Thus, the neural activity dynamically tracked the monkeys' uncertainty about the correct sequence of actions. These results are consistent with prefrontal involvement in representing subjective knowledge of the correct action sequence.

Many important goal-directed behaviors, such as making coffee or commuting to work, are composed of sequences of actions. Often the appropriate sequence of actions to complete a particular task has to be figured out by trial and error. When the brain is trying to find the appropriate sequence of actions, it has to monitor and integrate information about recent actions and their outcomes. Where in the brain is this information represented? Evidence from lesion and single-cell recording studies suggests that multiple regions in the primate frontal cortex contribute to the proper coordination of sequential actions[1,2]. For example, patients with frontal lobe damage have deficits on the Tower of London task, in which a configuration of beads stacked on three different rods must be reproduced with an efficient sequence of movements[3]. Similarly, many neurons in the frontal cortex of monkeys change their activity according to the temporal order of multiple movements in a sequence[4–12].

When monkeys have to learn arbitrary stimulus-response associations by integrating the outcomes of previous actions, individual neurons in the frontal cortex and the basal ganglia show changes in their activity that mirror the behavioral manifestations of these associations[13–18]. Similarly, when monkeys learn an association between two arbitrary visual stimuli, neural activity in the inferior temporal cortex incorporates this new information with a time course similar to that of the corresponding behavioral change[19]. These studies have provided important insights into how newly acquired information about stimulus-response and stimulus-stimulus associations are integrated and stored in the brain.

In the present study, we investigated how neural activity in prefrontal cortex evolves as the monkey gradually acquires knowledge of the correct movement sequence. Unlike an association between two arbitrary stimuli or between a stimulus and a response, information about the correct movement sequence must be extracted across time and multiple movements. To examine how this dynamic process is reflected in the population activity of prefrontal cortex, we trained monkeys on a sequential eye-movement task. During this task, the correct sequence changed unpredictably, forcing the monkey to discover the new sequence by trial and error. The corresponding changes in neural activity were analyzed using Bayesian decoding analyses. We found that the neural activity predicted the correct sequence with a higher probability—that is, the posterior probability of the correct sequence increased—as the monkey discovered which sequence was correct. This implies that, as the monkey discovered the correct sequence, the activity of the prefrontal cortical ensembles became more similar to the responses produced when the monkey knew which sequence was correct. Furthermore, the time course of this change in neural activity closely tracked the fraction of the monkey's decisions that were correct for the corresponding sequence, indicating that neural activity in prefrontal cortex represented the monkey's subjective or probabilistic knowledge of the correct sequence. In addition, there was a greater increase in posterior probability after the monkey produced a correct movement than after it produced an error, suggesting that the monkey learned more from its successes than from its mistakes.

RESULTS

Two rhesus monkeys were trained on a sequential decision-making task (**Fig. 1**; see Methods). In each trial, they had to execute one of eight possible movement sequences; thus a single trial consisted of a sequence of several movements. (The distinction between individual movements and trials is important; in the results that follow, we always indicate explicitly whether we are referring to a single movement or to an entire

Department of Brain and Cognitive Sciences, Center for Visual Science, University of Rochester, Rochester, New York 14627, USA. Correspondence should be addressed to B.B.A. (baverbeck@cvs.rochester.edu).

Received 25 August 2005; accepted 20 December 2005; published online 22 January 2006; doi:10.1038/nn1634

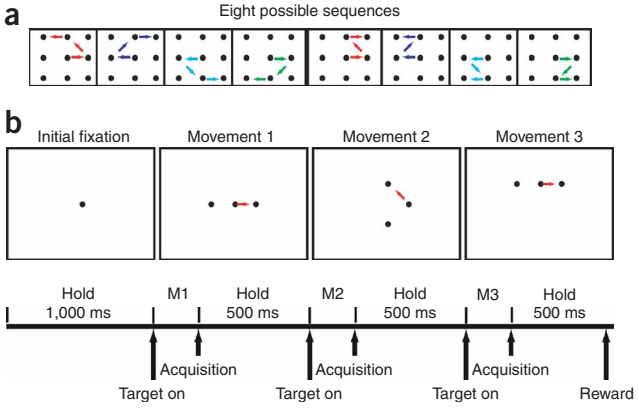

Figure 1 Behavioral task. (**a**) The eight possible sequences the monkeys were trained to execute. Each panel indicates one sequence. (**b**) Temporal sequence of choices in a single trial. The dot in the center of the 'Initial fixation' frame indicates the central fixation point. The two dots left and right of the central dot in the 'Movement 1' frame indicate the first two choice targets. Similarly, the two dots above and below fixation in the 'Movement 2' frame and left and right of fixation in the 'Movement 3' frame indicate the possible targets for the saccade at the corresponding points in the sequence.

trial.) Each movement in a trial corresponded to a saccade to one of two choice targets. If the monkeys selected the correct target, they advanced in the sequence. If they selected the incorrect target, they had to return to their previous target and were again presented with the same two choice targets. This was repeated until they selected the correct target. When the monkeys had chosen all three correct targets and completed the sequence, they received a juice reward. However, a trial was only counted as correct if the monkeys completed the sequence without making any incorrect movements. The correct decision at each stage of the sequence was not cued explicitly; however, it remained fixed until ten trials had been executed correctly. After ten correct trials (one block), a new sequence was selected pseudo-randomly. No sensory cues were provided to indicate when the sequence changed at the start of a new block. Thus, the monkeys had to discover the new sequence of correct decisions by trial and error.

Sequence coding in individual prefrontal neurons

We recorded from 485 neurons around the caudal principal sulcus of macaque prefrontal cortex (308 from monkey 1, 177 from monkey 2). These neurons were recorded in small ensembles of 1–15 neurons, with an average ensemble size of 8. The activity of neurons examined in this study often reflected not only the kinematics of the movements, but also their sequential aspects (**Fig. 2**). For example, one neuron responded strongly to the second movement in all eight sequences, but more strongly to these movements when they were part of sequences 2 and 6 than when they were part of the other sequences (**Fig. 2a**). Another neuron responded selectively to the first and last movements of sequences 2, 3, 5 and 8 (**Fig. 2b**). These patterns of activity are different from what would be expected of neurons with simple directional tuning. To quantify this nonmotoric, sequence-related activity across the population, we carried out a two-way analysis of variance (ANOVA) with movement and sequence as the two factors. In this analysis, the movement was coded as one of the ten possible eye movements with different directions and positions (see **Fig. 1**). We found that the main effect of sequence was significant for 288 out of 485 (59%) neurons ($P < 0.01$, type III sum of squares), and the movement effect was significant ($P < 0.01$) for 330 neurons (68%). Of the 373 neurons in which either sequence or movement was significant ($P < 0.01$), 77% showed a significant effect of sequence. This shows that the response of a majority of our neurons was affected not only by the movement being executed but also by the sequence. These single-cell results on the representation of sequential movements mostly replicate findings from previous studies[4,9].

Ensemble decoding analyses

Next we addressed the question of how the sequence represented in the neural activity of the ensemble changes while the monkey discovers the correct sequence for the current block. To do this, however, we first had to show that the neural activity represents the sequence information when the monkey has discovered the sequence for a particular block and is executing trials correctly. Although the ANOVA results showed that information about sequences and movements was encoded in the activity of individual neurons, they did not reveal how reliably this information was represented in the population of neurons. To examine this quantitatively, we carried out a decoding analysis using a linear decoding algorithm; we treated each movement of every sequence as a unique element and used the activity during the movement, in the simultaneously recorded neural ensembles, to predict which of the 24 possible movements (8 sequences × 3 movements per sequence) was being executed. Only movements from correct trials were included in this analysis; thus, these movements come from trials in which the monkeys presumably knew the correct sequence. Also, neurons were included in this analysis only if their activity was significantly different across these 24 movements (one-way ANOVA, $P < 0.05$). This prescreening resulted in 51 ensembles (**Fig. 3a**), where each ensemble corresponded to a set of cells that we simultaneously recorded from in a single session.

The predictive performance of the ensembles improved in an approximately linear fashion as the size of the ensemble increased (**Fig. 3b**); further, it was above chance ($1/24 = 4.2\%$) in essentially all cases, as expected given that individual neurons were already prescreened. The linear increase in decoding performance with ensemble size has been seen in other coding studies[20,21]. Of course, the increase would eventually become sublinear with more neurons, as performance cannot exceed 100%. This initial decoding analysis, combined with the ANOVA results, demonstrated that when the monkey knew the correct sequence, the neural activity could be used to predict both the correct sequence and which movement in that sequence was being executed.

To examine the possibility that nearby neurons might code for similar movements and sequences, we computed the correlation coefficient between the signal correlation[22] and the distance between neurons, across all pairs of simultaneously recorded neurons. The signal correlation is the correlation coefficient between the mean responses of the neurons to each movement of the different sequences; it measures the similarity in the responses of the neurons to the sequential movements. The correlation between the signal correlation and the distance between neurons was small and not significant ($\rho = -0.058$, $P > 0.05$). Thus, there was no detectable anatomical organization of neurons with respect to the factors relevant to this task.

Sequence discovery and representation of uncertainty

The above single-cell and ensemble decoding analyses showed that individual neurons often responded differently to the same movement in different sequences. Thus, knowledge of the correct sequence affected neural activity. This raises the question of how the neurons responded when the monkey was trying to figure out which sequence was correct in a new block. During this exploration period, the monkeys executed

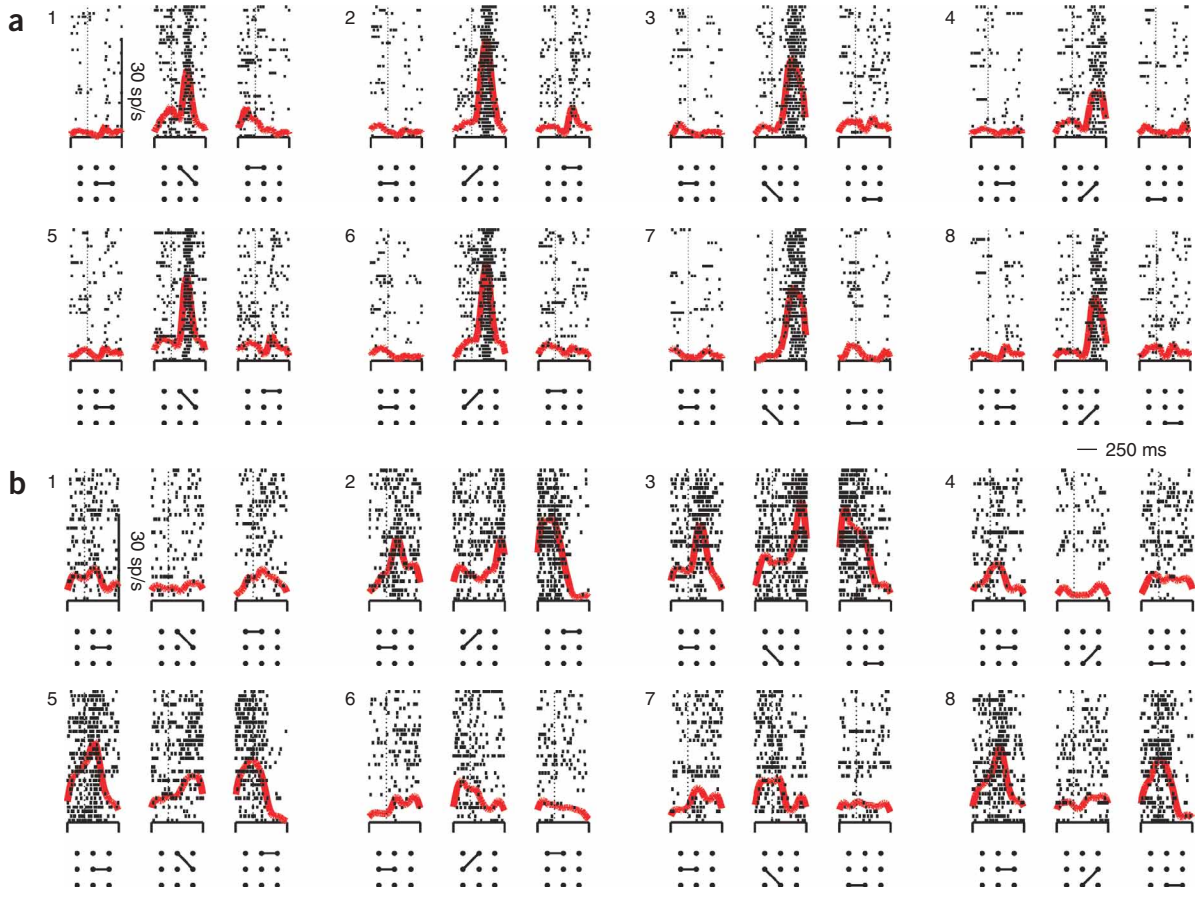

Figure 2 Raster plots and spike density functions for two example neurons. (**a**,**b**) In both neurons, a two-way ANOVA showed a significant effect of movement and a significant effect of sequence ($P < 0.01$, type III sum of squares). Each group of three rasters corresponds to a single sequence; each individual raster shows one of the three movements in the sequence. The direction of movement is indicated below the raster by a bar in the 3 × 3 grid; each dot in this grid corresponds to a potential target. Thus, the first movement of the sequence numbered '1' is from the central target to the right target, the second movement is to the top center target, and so on. The dotted vertical line in the rasters indicates target onset.

movements without always knowing whether they were correct or, most importantly, to which sequence they belonged.

We found that while the monkeys discovered the correct sequence for a particular block of trials, the patterns of neural activity in the recorded ensembles gradually evolved to become more similar to the patterns observed when the monkeys were executing trials correctly. To illustrate a typical series of neural responses during sequence discovery, we show rasters for the data from two simultaneously recorded neurons during two consecutive blocks of trials (**Fig. 4a**). It should be noted that error trials always included correct individual movements, as the monkey was given multiple opportunities to execute the correct movement. We were particularly interested in comparing the responses to the same movements before and after the animal learned the correct sequence. Therefore, for error trials, we show only the responses for the correct movements. In this example (**Fig. 4a**), the monkey made mistakes in the first two trials after the sequence switched from 2 to 4 (labeled Trial 1 and Trial 2) and then executed the next trial (labeled Trial 3) without errors. As the monkey discovered the correct sequence, the response of the illustrated pair of neurons for each movement became more similar to the average response expected for the same movement from the correct trials (**Fig. 4b**). This can be seen by comparing the average activity for the first movement of the sequence from correct trials (m_1) and the activity during the first movement in the first three trials (1, 4 and 7; **Fig. 4b**). Similarly, activity during the second movement (2, 5 and 8) gradually became more similar to the mean for the second movement in the sequence (m_2, **Fig. 4b**). These results can be summarized by plotting the euclidean distance between the neural response for each movement and the corresponding mean response, from correct trials, for that movement (**Fig. 4c**). The graph

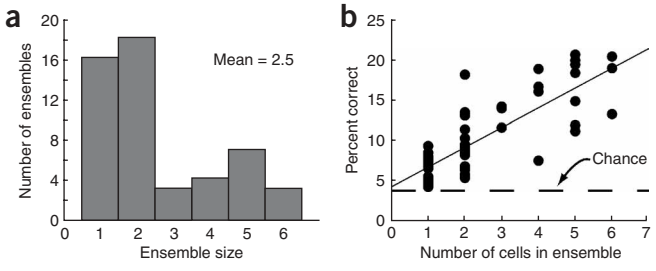

Figure 3 Ensemble size and classification performance. (**a**) Distribution of ensemble size. (**b**) The classification performance is plotted as a function of the number of neurons in an ensemble for which the effect of movement was significant in a one-way ANOVA ($P < 0.01$). The regression line (solid) fit to the data had an intercept of 4.3 and a slope of (mean ± s.e.m.) 2.4 ± 0.0654% per neuron.

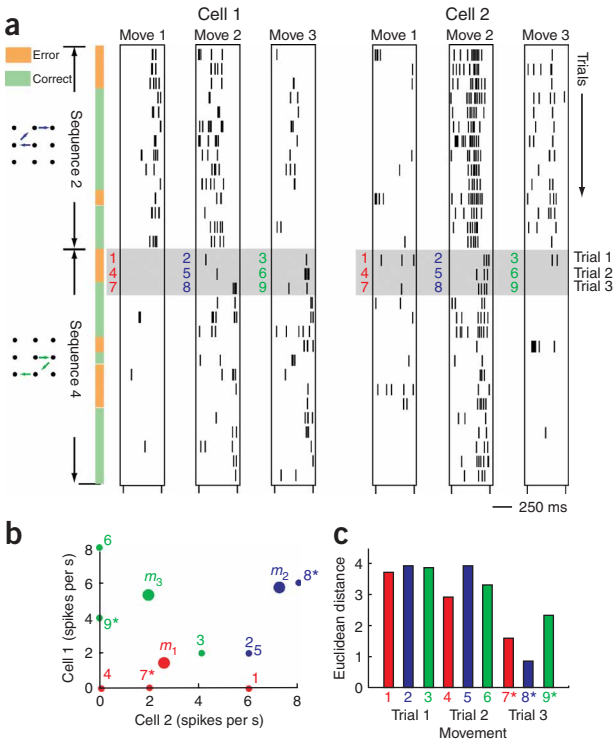

Figure 4 Temporal evolution of activity in a pair of neurons during sequence discovery. (**a**–**c**) Rasters in **a** show the neural activity for two consecutive blocks of trials. In the first block, 'Sequence 2' was correct; in the second block, 'Sequence 4' was correct. The activity shown for the error trials are for the correct movements. Accordingly, the column of responses labeled 'Move 2' for Sequence 2 is always a movement from left to the top center; however, this might not be the second movement of the trial if the monkey made mistakes before this movement. Numbers highlighted in gray correspond to the movements plotted in **b** and **c**. In **b**, the small dots indicate firing rates of the two neurons in the first correct trial and the two previous error trials as highlighted in **a**. The larger colored dots labeled '*m*' are the mean activity averaged across correct trials in all blocks for the corresponding movements in Sequence 4. The number labels on the small dots correspond to the numbers in the raster in **a** and the color of the number indicates the ordinal position in the sequence. The asterisks indicate movements from the first correct trial. The euclidean distance between the activity in single trials and the mean activity for the corresponding movement is shown in **c**. The movement number corresponds to those in the rasters and the rate plot; asterisks again indicate movements from the first correct trial.

shows that this distance decreased across trials for each movement of the sequence.

Although the results from the analysis based on the euclidean distance can be easily visualized, this analysis has two important drawbacks. First, it does not take into account the variability in neuronal activity. Neurons that are more variable should have their distances weighted less. Second, the euclidean distance is not normalized, so ensembles with more neurons will have larger euclidean distances, making averages across ensembles uninterpretable. Both of these drawbacks can be addressed by using a linear decoding model and Bayes' theorem to calculate a posterior probability for each individual movement. The posterior probability is the probability that the monkey executed a particular movement given the observed neural response in an ensemble. We can also think of this as the ability to predict which movement the monkey was executing by observing the neural activity. If the posterior probability is higher, we have more confidence in our ability to predict the movement that was being executed. The posterior probability tended to be inversely related to the euclidean distance (**Fig. 4c**): as the neural response got closer to the mean response and as the euclidean distance decreased, the probability that the response was generated by a distribution with the corresponding mean increased. Because the posterior probability is a proper probability, it is constrained to lie between 0 and 1, which allowed us to average our results across ensembles. When assessing the classification performance of the ensembles (**Fig. 3**), we assigned the neural response in each trial to one of the movements in the sequence—the one that had the largest posterior probability given the neural activity. For the analyses on dynamic changes in neural activity associated with sequence discovery (**Fig. 5**), we calculated the posterior probability given the neural activity and computed its average across ensembles and trials. The posterior probability provides a continuous measure of the neural responses, whereas classification is categorical. Therefore, the posterior probability was better suited to measuring continuous, gradual changes in ensemble activity.

Because our goal was to compare the time course of the changes in neural activity and behavior, we first describe the time course of behavioral changes. After the correct sequence changed, the number of incorrect trials before the first correct trial, averaged across the two monkeys, was 2.66 (**Fig. 5a**). Once a correct trial was executed, indicating that the monkey had discovered the correct sequence for the block, there were about 0.5 error trials per correct trial. This shows that the monkeys were relatively proficient at the task. Monkeys also rarely perseverated. When they executed an incorrect movement, they almost always made a saccade to the other choice target on the next movement. Overall, the rate of perseveration was 6.9% across all trials. When we restricted the analysis to the first trial of a new block, the frequencies of perseveration on the first, second and third movements of the block were 3.9%, 8.7% and 6.3%, respectively.

The neural responses evolved in parallel with the monkeys' behavior. The average posterior probability of the monkeys' movements, predicted by the neural activity in the ensembles, increased after the monkeys had executed the first correct trial in a new block (**Fig. 5a**), and the difference between the posterior probability before and after a correct trial was significant across the population (individual *t*-tests, all $P < 0.01$). To understand this more clearly, we can return to the example discussed above (**Fig. 4a**), and remember that the neural activity related to a particular movement reflects the sequence to which the movement belongs. When the block switched from sequence 2 to sequence 4 (**Fig. 4a**) and the monkey was required to make the second movement downward, it was initially doing so without knowing which sequence it was supposed to execute because it had not yet worked out the correct sequence. It would make this downward movement, however, because if it tried the upward movement it would be forced back to the previous target and given the choices again until it selected the lower target. Because the neural responses reflect not just the movement being made, but also the sequence, the neural activity should reflect the monkey's lack of knowledge of the correct sequence, which it does (**Fig. 5a**). When the monkeys had not yet resolved which sequence was correct for the current block, as indicated by a large number of incorrect movements, the posterior probability was also relatively low, suggesting that the neural activity reflected the monkey's lack of knowledge of the correct sequence.

The monkeys also tried movements that would have been correct for the sequence in the previous block; this allowed us to investigate, at a finer grain, the correspondence between the posterior probability of the movement predicted by the neural response and the percentage of correct decisions made by the monkey. To do so, we 'unpacked' the

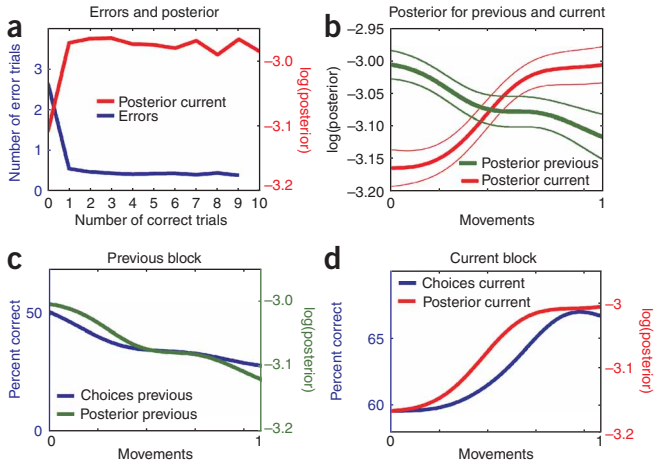

Figure 5 Trial-by-trial and movement-by-movement evolution of the posterior probability and the monkeys' choices. (**a**) Horizontal axis, number of correct trials (*n*). Left vertical axis, number of errors made by the monkeys between the *n*th and (*n*+1)th correct trials. Right vertical axis, population average of the log(posterior probability), predicted by the decoding algorithm from the neural activity during correct movements for each correct trial. The value for 0 correct trial is for the correct movements before the first correct trail. (**b**) Log (posterior probability) evaluated from the neural activity according to the correct sequence in the previous (green) and current (red) block for the movements executed before the first correct trial. The average of these values evaluated for the current block corresponds to the first value plotted in **a**. All the movements from the trials before a correct trial in each individual block were distributed evenly across the *x*-axis, such that the first movement of the block had a value of 0 and the last movement had a value of 1. A moving average was done to estimate the average of the posterior probability along the *x*-axis. Thick lines indicate means; thin lines indicate s.e.m. estimated by bootstrap resampling of the original dataset. (**c**) Average log(posterior probability) for the sequence in the previous block, plotted along with the percentage of the monkeys' decisions that would have been correct in the previous block. (**d**) Same as **c**, except that the movements were evaluated relative to the correct sequence in the current block.

movements that were averaged for the posterior probability estimate of error trials before the first correct trial (trial 0, **Fig. 5a**) and examined the evolution of the population-averaged posterior probability on a movement-by-movement basis (**Fig. 5b**). On average, the monkeys executed 8.7 correct movements before the first correct trial. The progression of the log(posterior probability) predicted by decoding the neural activity of correct movements in the current block ('Posterior current', **Fig. 5b**) shows, at the population level, the effect described above for a single pair of neurons (**Fig. 4c**). Specifically, as the monkeys figured out the correct sequence, the neural activity in the ensembles became more similar to the activity during the correct trials, as reflected in the increasing posterior probability. This curve flattens out at just under –3, which is slightly less than the average log(posterior probability) for the subsequent correct trials of the block (just over –3; **Fig. 5a**). Thus the posterior probability increased slightly in the correct trials.

The monkeys also tried movements that would have been correct for the sequence in the previous block; this allowed us to assess how the posterior probability of these movements evolved as the monkeys discovered the correct sequence and realized that the sequence from the previous block was no longer correct ('Posterior previous', **Fig. 5b**). For the first movement of the new block, the larger value of the posterior probability for the previous block reflects the fact that the monkey was not cued when the sequence switched and therefore it continued to execute the sequence that had been correct in the previous block. The posterior probability then gradually decreased as the monkey discovered the correct sequence for the current block. Although these curves are not linear, they are monotonic and so we tested the significance of their increase or decrease by fitting linear regression models with log(posterior probability) as the dependent variable and the normalized number of movements as the independent variable. The slope of the regression was significantly different from zero for both previous and current blocks ($P < 0.01$).

We also compared the posterior probability of the movements that would have been correct for the previous sequence with the average number of choices the monkey made that would have been correct in the previous block. The average time evolution of the posterior probability closely followed the time evolution of the choices (**Fig. 5c**). Thus, as the monkey made fewer and fewer movements that would have been correct in the previous block, the posterior probability of those movements decreased as well. The results comparing the posterior probability and the decisions for the current trial were in close correspondence (**Fig. 5d**), although in this case the change in the posterior probability preceded that in the behavioral data by at least one movement. In summary, as the monkey made more movements consistent with the current sequence, the neural response to each of those movements became more similar to the neural response from trials when the monkey knew the sequence and was executing the trial without mistakes.

It is important to consider the potential effect that changes in the kinematics of the saccades might have on these results. If a saccade parameter, such as direction, became more variable for a particular movement and the neural activity reflected this variability, the decoding analysis would potentially do worse as the variability increased and this could be manifested as a decrease in the posterior probability. This is unlikely, however, as the posterior probability changed in opposite directions for movements that were correct in the previous and current blocks (**Fig. 5b**). The eye movements would have to become more variable for only those movements that were correct in the previous trial and less variable for correct movements in the current trial. To examine this issue directly, we carried out regression analyses that estimated the effect of the peak velocity, the distance and the direction of the saccades on the neural responses for individual movements. We found that one or more of these variables was significant ($P < 0.01$) in only 6% of the neurons. We carried the analysis one step further by removing the influence of these variables on the activity of neurons for which they had a significant effect and recomputing the posterior probabilities. The results were similar, except that the posterior probabilities were shifted up slightly, as removing the variability in neural activity that was due to the variable kinematics improved the decoding performance.

Given our analytical framework and the correspondence between the posterior probability and the monkey's choice behavior, one can ask the following question: does the monkey learn more about which sequence is correct from making a correct decision or an incorrect decision? To answer this question, we fit the following regression equation to the log(posterior probabilities) evaluated with respect to the sequence in the current block: $p(n) = a_0 + a_1 p(n-2) + a_2 c(n-1) + a_3 m$, where $p(n)$ is the log of the posterior probability for movement n, $c(n-1)$ is a variable that takes a value of 0 for an incorrect decision and 1 for a correct decision for movement $n-1$, a_0, a_1, a_2 and a_3 are regression coefficients and m is shorthand notation for two dummy variables that code the ordinal position of the movement in the sequence ($m = [1\ 0]$, $[0\ 1]$ and $[0\ 0]$ for the first, second and third movements, respectively). The dummy variables for ordinal position were included to eliminate

the possible confound of sequence position on the learning effect. In essence, we asked how much the posterior probability increased between the previous and subsequent movement, when the intervening movement was correct or incorrect. We found that a_2 was significantly different from zero ($P < 0.01$) and that it had a value (mean ± s.e.m.) of 0.206 ± 0.0024. The values of the other variables were as follows: $a_0 = -2.954 \pm 0.0059$, $a_1 = 0.060 \pm 0.0004$ and $a_3 = [0.070 \pm 0.0029, -0.087 \pm 0.0034]$. Thus, the contribution of a correct movement to the log(posterior probability) was 0.206 more than that of an error movement. Furthermore, this effect was considerably larger than the effect of the ordinal position of the movement, given by a_3.

DISCUSSION

Neural activity in the prefrontal cortex related to a sequence of movements reflects not only the movement being made but also the sequence in which this movement is embedded. Such sequence-specific neural activity has been documented in many brain areas and in many protocols[4–9,23–26]. In our protocol, the correct sequence of decisions had to be discovered by trial and error and this provided a unique opportunity to examine how the neural coding of sequential movements changes during this process. Although a similar trial-and-error learning protocol has been used in a previous study[7], the focus there was on differential representation of new and well-learned sequences and the study did not examine how learning was related to changes in ensemble neural activity. In our protocol, all sequences were well learned and there was no new sequence learning. Thus, in each block, when the sequence was switched, the monkeys had to figure out which sequence of the set of eight was correct for the current block. We took advantage of the fact that neurons responded differently to the same movement in different sequences to see how the neural responses tracked the monkeys' knowledge of the correct sequence. We found that, as the monkeys learned the sequence and made more decisions that were correct for the current block, there was an increase in the posterior probability of the movements estimated from the ensemble neural activity. Furthermore, as fewer movements were executed that were correct for the previous sequence, there was a decrease in the posterior probability of these movements evaluated with respect to the previous sequence. Thus, the posterior probability of the movement predicted by the neural activity closely tracked the monkeys' knowledge of which movements were correct.

To compute the action that has the largest expected utility, the theory of optimal decision making requires modeling of both the value of an outcome and the belief that the outcome will be obtained[27]. Thus, the representation of belief or subjective knowledge shown here, along with previous work describing a neural correlate of action outcome in the same brain region[28], implies that the prefrontal cortex can represent the information necessary for optimal decision making. Previous work in parietal cortex has also shown that neural activity can reflect a monkey's confidence in its decision. In a task in which monkeys had to estimate the average direction of dots in a dynamic random-dot display, neurons in the lateral intraparietal area reflected the amount of motion coherence[29,30]; moreover, the monkey's behavior also scaled with the level of motion coherence, such that direction estimation was more accurate when the dots were more coherent.

Learning-related changes in neural activity have been shown in several brain areas. For example, as monkeys learn to associate pairs of stimuli, neural responses to the two stimuli in inferotemporal cortex become more similar[19]. Experiments in prefrontal cortex[13] and the basal ganglia[17] have also shown that neurons reflect behavioral adaptation to arbitrary task contingencies. In these studies, learning an arbitrary stimulus-response mapping caused the representation of movement direction to appear earlier in a delay period before the execution of an eye movement. A gradual change in neural activity related to the learning of stimulus-response contingencies has also been shown in other frontal cortical areas[14–16,18]. Previous studies, however, have not directly investigated whether these changes actually lead to an increase in the information content of the neural responses. The present study focused on changes in neural activity associated with the selection of appropriate sequences of actions. In our task, the response was not contingent on an arbitrary stimulus-response mapping, but rather on which sequence was correct in the current block. Furthermore, we showed directly that the information coded in the ensemble neural activity about the correct sequence increased as the monkey discovered the correct sequence.

We showed that the fraction of the monkey's choices that would be correct for the current sequence increased slowly when a new sequence was introduced. This fraction of choices is an overt measure of the monkey's internal prediction of which movement would be correct and consequently of which sequence it should execute. Our analyses also demonstrated a close link between the patterns of activity in ensembles of simultaneously recorded prefrontal neurons and the fraction of the monkey's choices that were correct for the current or previous sequence. Thus, the dynamics of the posterior probability, which closely follows the monkey's pattern of decisions, reflects the monkey's subjective knowledge of which sequence is correct.

METHODS

General. Two male rhesus macaques were used in this study. All surgical and experimental procedures conformed to the guidelines of the US National Institutes of Health and were approved by the University of Rochester Committee on Animal Research. The recording chamber (18-mm diameter) was placed over the prefrontal cortex in a sterile surgery using stereotaxic coordinates derived from structural magnetic resonance imaging (MRI). Neural responses were recorded using a 16-channel multielectrode recording system (Thomas Recording) and single-unit spikes were sorted online using the Plexon data acquisition system (Plexon). The task was presented to the monkeys on a CRT monitor. A custom Microsoft 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). Location of the frontal eye field (FEF) was verified in both monkeys using microstimulation. Electrode penetrations were considered to be within the FEF when a 50-μA 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 monkey.

Behavioral task. The monkeys 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 × 3 grid of targets spaced by 5.3° of visual angle (**Fig. 1a**). The monkeys began a trial by acquiring a central fixation point (**Fig. 1b**). After a 1-s hold period, two targets were presented to the left and right of fixation. One of the targets was the correct choice in a given block. The monkey was allowed to make a saccade as soon as the target appeared. If the monkey made a saccade to the correct target and maintained fixation for 500 ms, two additional choice targets were presented above and below the central fixation point and the monkey was required to select one of these targets. The third pair of choice targets were presented to the left and right of the previous fixation target. If the monkey made an incorrect decision, the choice targets were extinguished and it had to return to the previous fixation target (not necessarily the beginning of the sequence). It was then presented with the same choice targets again. After making three correct decisions, regardless of the number of intervening incorrect decisions, the monkey was given a juice reward. The total trial length was constrained to be less than 7 s, but this limit was rarely reached. The sequences were presented in a randomized

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

Data analysis. The neural activity for each movement was quantified as the spike rate in a 500-ms window starting 250 ms before the saccade. The ANOVA results for sequence and movement were carried out using type III sum of squares, as the design was not balanced with respect to the two factors.

We used a Gaussian decoding analysis described in detail previously[31]. This analysis assumes that the distribution of neural responses is Gaussian for a particular movement in a sequence. Thus, we have a Gaussian likelihood function given by

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

where r is a vector of spike rates for a movement, μ_i is the vector of mean spike counts for movement i, the superscript T indicates transpose, Q is the noise covariance matrix pooled across conditions and $||$ indicates the determinant of the matrix. For the decoding analyses, i takes on values between 1 and 24.

The posterior probability that a particular target led to the response under consideration is given by Bayes' theorem:

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

where t is one of the targets. In this study we assumed a flat prior probability, and thus $p(t)$ is a constant. The normalization is given by

$$p(r) = \sum_t p(r|t)p(t) \quad (3)$$

In the analyses where we consider the log of the posterior probability, the posterior probability is given directly by the left side of equation. (2). To carry out classification explicitly, movements were predicted by selecting, from the conditional distribution of targets, the movement with the maximum probability given the neural activity:

$$\hat{t} = \arg\max_t p(t|r) \quad (4)$$

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

The decoding analyses on ensemble data for the movements in correct trials (**Fig. 3b**) were done using twofold cross-validation on the data from all correct trials. In other words, the means and covariances in the likelihood (equation (1)) were estimated using half the data and then classification was carried out on the other half. The estimation and classification data were then switched and the analysis was repeated. For the decoding analyses that characterized the difference in the posterior probability between trials before and after a correct trial had been executed (**Fig. 5a**), the posterior probabilities were calculated using cross validation. For the remaining analyses on the movement-by-movement evolution of the posterior probability (**Fig. 5b–d**), the posterior probabilities were determined using means and covariances calculated from all correct trials from all blocks. For the posterior probability and the choices evaluated with respect to the correct movement in the current block (posterior current in **Fig. 5b**), we analyzed neural activity from movements that were correct for the sequence in the current block. For the analyses that calculated the posterior probability for the previous block, we analyzed movements that would have been correct for the sequence in the previous block. The posterior probability for all movements was calculated by entering the corresponding response of simultaneously recorded neurons, r, into equation (1) and using the value of the mean response, μ_i, that corresponded to the movement from the appropriate sequence.

ACKNOWLEDGMENTS

We are grateful to D. Barraclough for help with surgeries, L. Carr for technical assistance and J. Swan-Stone for programming. This work was supported by grants from the US National Institutes of Health (R01-MH59216, T32-MH19942 and P30-EY01319).

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Published online at http://www.nature.com/natureneuroscience/
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions/

1. Petrides, M. & Milner, B. Deficits on subject-ordered tasks after frontal- and temporal-lobe lesions in man. *Neuropsychologia* **20**, 249–262 (1982).
2. Penfield, W. & Evans, J. The frontal lobe in man: a clinical study of maximum removals. *Brain* **58**, 115–133 (1935).
3. Shallice, T. Specific impairments of planning. *Phil. Trans. R. Soc. Lond. B* **298**, 199–209 (1982).
4. Tanji, J. Sequential organization of multiple movements: involvement of cortical motor areas. *Annu. Rev. Neurosci.* **24**, 631–651 (2001).
5. Shima, K. & Tanji, J. Neuronal activity in the supplementary and presupplementary motor areas for temporal organization of multiple movements. *J. Neurophysiol.* **84**, 2148–2160 (2000).
6. Mushiake, H., Inase, M. & Tanji, J. Selective coding of motor sequence in the supplementary motor area of the monkey cerebral cortex. *Exp. Brain Res.* **82**, 208–210 (1990).
7. Nakamura, K., Sakai, K. & Hikosaka, O. Neuronal activity in medial frontal cortex during learning of sequential procedures. *J. Neurophysiol.* **80**, 2671–2687 (1998).
8. Averbeck, B.B., Chafee, M.V., Crowe, D.A. & Georgopoulos, A.P. Neural activity in prefrontal cortex during copying geometrical shapes I. Single cells encode shape, sequence, and metric parameters. *Exp. Brain Res.* **150**, 127–141 (2003).
9. Rhodes, B.J., Bullock, D., Verwey, W.B., Averbeck, B.B. & Page, M.P. Learning and production of movement sequences: behavioral, neurophysiological, and modeling perspectives. *Hum. Mov. Sci.* **23**, 699–746 (2004).
10. Barone, P. & Joseph, J.P. Prefrontal cortex and spatial sequencing in macaque monkey. *Exp. Brain Res.* **78**, 447–464 (1989).
11. Isoda, M. & Tanji, J. Contrasting neuronal activity in the supplementary and frontal eye fields during temporal organization of multiple saccades. *J. Neurophysiol.* **90**, 3054–3065 (2003).
12. Clower, W.T. & Alexander, G.E. Movement sequence-related activity reflecting numerical order of components in supplementary and presupplementary motor areas. *J. Neurophysiol.* **80**, 1562–1566 (1998).
13. Asaad, W.F., Rainer, G. & Miller, E.K. Neural activity in the primate prefrontal cortex during associative learning. *Neuron* **21**, 1399–1407 (1998).
14. Chen, L.L. & Wise, S.P. Conditional oculomotor learning: population vectors in the supplementary eye field. *J. Neurophysiol.* **78**, 1166–1169 (1997).
15. Chen, L.L. & Wise, S.P. Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations. *J. Neurophysiol.* **73**, 1101–1121 (1995).
16. Chen, L.L. & Wise, S.P. Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations. *J. Neurosci.* **16**, 3067–3081 (1996).
17. Pasupathy, A. & Miller, E.K. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* **433**, 873–876 (2005).
18. Mitz, A.R., Godschalk, M. & Wise, S.P. Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. *J. Neurosci.* **11**, 1855–1872 (1991).
19. Messinger, A., Squire, L.R., Zola, S.M. & Albright, T.D. Neuronal representations of stimulus associations develop in the temporal lobe during learning. *Proc. Natl. Acad. Sci. USA* **98**, 12239–12244 (2001).
20. Reich, D.S., Mechler, F. & Victor, J.D. Independent and redundant information in nearby cortical neurons. *Science* **294**, 2566–2568 (2001).
21. Averbeck, B.B., Crowe, D.A., Chafee, M.V. & Georgopoulos, A.P. Neural activity in prefrontal cortex during copying geometrical shapes II. Decoding shape segments from neural ensembles. *Exp. Brain Res.* **150**, 142–153 (2003).
22. Lee, D., Port, N.L., Kruse, W. & Georgopoulos, A.P. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J. Neurosci.* **18**, 1161–1170 (1998).
23. Lu, X. & Ashe, J. Anticipatory activity in primary motor cortex codes memorized movement sequences. *Neuron* **45**, 967–973 (2005).
24. Kermadi, I., Jurquet, Y., Arzi, M. & Joseph, J.P. Neural activity in the caudate nucleus of monkeys during spatial sequencing. *Exp. Brain Res.* **94**, 352–356 (1993).
25. Kermadi, I. & Joseph, J.P. Activity in the caudate nucleus of monkey during spatial sequencing. *J. Neurophysiol.* **74**, 911–933 (1995).
26. Procyk, E., Tanaka, Y.L. & Joseph, J.P. Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat. Neurosci.* **3**, 502–508 (2000).
27. Jaynes, E.T. *Probability Theory: the Logic of Science* (Cambridge Univ. Press, Cambridge, UK, 2003).
28. Barraclough, D.J., Conroy, M.L. & Lee, D. Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* **7**, 404–410 (2004).
29. Shadlen, M.N. & Newsome, W.T. Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci. USA* **93**, 628–633 (1996).
30. Roitman, J.D. & Shadlen, M.N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
31. Averbeck, B.B. & Lee, D. Neural noise and movement-related codes in the macaque supplementary motor area. *J. Neurosci.* **23**, 7630–7641 (2003).