

Dynamics of Cortical Neuronal Ensembles Transit from Decision Making to Storage for Later Report

Adrián Ponce-Alvarez,^{1,2} Verónica Nácher,² Rogelio Luna,² Alexa Riehle,¹ and Ranulfo Romo^{2,3}

¹Institut de Neurosciences de la Timone, CNRS–Aix-Marseille Université, 13005 Marseille, France, ²Instituto de Fisiología Celular-Neurociencias, Universidad Nacional Autónoma de México, 04510 Mexico D.F., Mexico, and ³El Colegio Nacional, 06020 Mexico D.F., Mexico

Decisions based on sensory evaluation during single trials may depend on the collective activity of neurons distributed across brain circuits. Previous studies have deepened our understanding of how the activity of individual neurons relates to the formation of a decision and its storage for later report. However, little is known about how decision-making and decision maintenance processes evolve in single trials. We addressed this problem by studying the activity of simultaneously recorded neurons from different somatosensory and frontal lobe cortices of monkeys performing a vibrotactile discrimination task. We used the hidden Markov model to describe the spatiotemporal pattern of activity in single trials as a sequence of firing rate states. We show that the animal's decision was reliably maintained in frontal lobe activity through a selective state sequence, initiated by an abrupt state transition, during which many neurons changed their activity in a concomitant way, and for which both latency and variability depended on task difficulty. Indeed, transitions were more delayed and more variable for difficult trials compared with easy trials. In contrast, state sequences in somatosensory cortices were weakly decision related, had less variable transitions, and were not affected by the difficulty of the task. In summary, our results suggest that the decision process and its subsequent maintenance are dynamically linked by a cascade of transient events in frontal lobe cortices.

Introduction

Previous studies in monkeys performing a discrimination task have provided new insights into how cortical activity relates to sensory evaluation, working memory, decision making, and the subsequent maintenance of the decision output (Romo et al., 1999, 2002, 2004; Hernández et al., 2000, 2002, 2010; Salinas et al., 2000; Brody et al., 2003; Romo and Salinas, 2003; Luna et al., 2005; Lemus et al., 2007). In this task, monkeys compared the frequency of two mechanical vibrations sequentially delivered to their fingertip and report their perceptual evaluation by pressing one of two buttons. It was shown that, while the values of the base (f_1) and comparison (f_2) frequencies are encoded in the primary somatosensory cortex (S1), neuronal activity in the secondary somatosensory cortex (S2), prefrontal, premotor, and motor cortices evolves to a function of $f_2 - f_1$. These differential responses correlate with the animal's decision and can be maintained until a cue triggers the motor report.

These results were obtained using across-trial averages of single-unit activities and choice probability measures (Britten et

al., 1996; Romo et al., 2002). Despite its efficiency, this approach has the disadvantage of neglecting nonstationarities due to the trial-by-trial variability in the time course of neuronal activity. Dismissing this variability can overlook features of neural processes that are not time-locked to any event or stimulus, but are linked to internal dynamics (Seidemann et al., 1996; Renoult et al., 2006). Here, we investigated the evolution of the collective activity of sets of simultaneously recorded neurons, called neuronal ensembles (NEs), during single trials of perceptual discrimination. Specifically, we asked how decision-making and decision maintenance processes evolve in single trials.

To address these questions, we used the hidden Markov model (HMM) to characterize the NE activity from sensory (S1 and S2) and frontal lobe cortices [dorsal premotor cortex (DPC); medial premotor cortex (MPC); and primary motor cortex (M1)] during a postponed discrimination task. In this model, a latent variable, or “hidden state,” determines the firing rate of each neuron of the NE. In each trial, the model transits through a state sequence with transitions that can occur at variable times from trial to trial. The main hypothesis is that variability is shared by many neurons, so that single-trial NE dynamics can be extracted accurately and with fewer degrees of freedom (i.e., a finite number of states). The HMM was successfully applied to describe cortical activity during movement withholding and preparation (Seidemann et al., 1996; Kemere et al., 2008) and taste processing (Jones et al., 2007).

The present work demonstrates that the state sequences can depend on the result of the comparison between the two stimuli. Comparison selectivity is distributed across frontal lobe cortices, and also, albeit weakly, in S2. We show that the state transitions are produced by rapid and coherent changes in the firing rates of

Received Dec. 13, 2011; revised June 8, 2012; accepted July 5, 2012.

Author contributions: R.R. designed research; V.N., R.L., and R.R. performed research; A.P.-A. and A.R. analyzed data; A.P.-A., A.R., and R.R. wrote the paper.

A.P.-A. was supported by a doctoral fellowship from the Consejo Nacional de Ciencia y Tecnología (CONACYT). A.R. was partially supported by French National Research Agency Grant ANR-05-NEURO45-01. The research of R.R. was partially supported by an International Research Scholars Award from the Howard Hughes Medical Institute, and grants from the Dirección del Personal Académico de la Universidad Nacional Autónoma de México and CONACYT. We thank Nicolas Brunel, Moshe Abeles, Bjorg Kilavik, and Joachim Confais for critical comments and ideas to earlier versions of this work, and Adrián Hernández, Antonio Zainos, and Manuel Álvarez for technical assistance.

Correspondence should be addressed to Ranulfo Romo at the above address. E-mail: rromo@ifc.unam.mx.

A. Ponce-Alvarez's present address: Computational Neuroscience Group, Department of Technology, Universitat Pompeu Fabra, Roc Boronat 138, 08018 Barcelona, Spain.

DOI:10.1523/JNEUROSCI.6176-11.2012

Copyright © 2012 the authors 0270-6474/12/3211956-14\$15.00/0

many neurons. While single-trial dynamics are time-locked to the comparison stimulus in sensory areas, in frontal lobe areas, dynamics are highly variable from trial to trial in a way that depends on the difficulty of the task.

Materials and Methods

Multiple single-neuron activity was recorded simultaneously in several cortical areas of two male monkeys (*Macaca mulatta*) trained to perform a vibrotactile discrimination task (Lemus et al., 2007; Hernández et al., 2010). Simultaneous recordings were obtained by using a new procedure described by Hernández et al. (2008). Monkeys were handled according to institutional standards that meet those of the National Institutes of Health and the Society for Neuroscience.

Postponed discrimination task. The sensory discrimination task used here has already been described previously (Lemus et al., 2007) (see Fig. 1A). Two monkeys were trained to discriminate the difference in frequency between two consecutive mechanical vibrations delivered to one fingertip. Monkeys were asked to report the discrimination by pressing one of two push buttons with the nonstimulated hand, after a fixed delay of 3 s between the end of f2 and a cue (PU) that triggered the beginning of the motor report. Stimuli were delivered to the skin of the distal segments of one digit of the right, restrained hand by means of a computer-controlled stimulator. Vibrotactile stimuli were trains of short mechanical pulses lasting 20 ms during a period of 0.5 s. We collected data using the stimulus set of Figure 1B, usually 7–10 trials per stimulus pair. Stimulus amplitudes were adjusted to equal subjective intensities (Mountcastle et al., 1990; Hernández et al., 1997). The report buttons were located in front of the animal, 25 cm away from the shoulder and at eye level, and 7 and 10.5 cm to the left of the midsagittal plane, respectively, leading to a difference between medial and lateral movements of $\sim 11^\circ$. As pointed out by Hernández et al. (2010), under these conditions some activity related to arm motion may be expected, but should be highly similar for the two arm movements. The animal was rewarded with a drop of water for correct discriminations and randomly at the end of ambiguous trials (f1 = f2).

Control task. The monkeys learned also to perform another version of the task. In this control task, the button that has to be pushed for the correct answer was illuminated at the beginning of the trial. Trials in this test began exactly as described above and in Figure 1A, except that when the probe touched the skin, one of the target switches was illuminated. The monkey had to respond by holding the immovable key. Then, after a variable delay period during which the light was kept on and the two consecutive stimuli were delivered exactly as in the vibrotactile discrimination task (set 1; usually five to seven trials per stimulus pair), the light was turned off and the probe was simultaneously lifted from the skin after a fixed delay period of 3 s between f2 and PU. The monkey was rewarded for pressing the previously illuminated push button. Arm movements in this situation were identical with those in the vibrotactile discrimination task but were cued by visual stimuli. In monkey 1, we recorded the activity of the same cells during both the discrimination task and the control task. The monkey completed the control bloc following the bloc of discrimination trials, while isolation of the same cells was maintained. In monkey 2, different cells were recorded during the control task compared with the ones recorded in the discrimination task.

Recording techniques. Neuronal recordings were obtained with an array of seven independent microelectrodes (1–3 M Ω) (Romo et al., 1999) inserted into each cortical area. Recordings in S1, S2, and DPC were made in the hemisphere contralateral to the stimulated hand, those in M1 were made in the other hemisphere, contralateral to the responding arm, and recordings in MPC were made either in the hemisphere contralateral or ipsilateral to the responding arm. We used well established electrophysiological and anatomical criteria to distinguish between cortical areas (Salinas and Romo, 1998; Romo et al., 1999, 2002, 2004; Salinas et al., 2000; Hernández et al., 2002). In S1, we recorded single neurons with cutaneous receptive fields confined to the distal segments of the glabrous skin of fingertips 2, 3, or 4 and had quickly adapting properties. All neurons recorded in S2 had large cutaneous receptive fields confined to the hand contralateral to the recording site. All neurons in MPC were

recorded in the pre-SMA, located rostral to a line passing from the midline to the posterior edge of the arcuate sulcus (Matsuzaka et al., 1992). All recordings in DPC were made in the arm region of F2 (Rizzolatti and Luppino, 2001), located in front of M1 (F1), lateral to the central dimple, posterior to F7 and genu of arcuate sulcus (Rizzolatti and Luppino, 2001). Recordings in M1 were confined to the arm region of M1, confined to the anterior bank and crown of the central sulcus, medial to the level of the posterior genu of the arcuate sulcus and lateral to the central dimple (Rizzolatti and Luppino, 2001). On all recordings sessions in M1, acceptable penetration sites were first identified. The criterion was that, throughout the penetration track, neurons were found that responded both during the task and to passive movements of the contralateral arm. The passive responses had to be related to shoulder and elbow joints; when they were associated with wrist and finger movements, the penetration was discarded. If these conditions were met, other neurons with different characteristics but recorded in the same penetration were also studied and considered in the analysis. Recording sites changed from session to session.

Hidden Markov model. Here, we review some aspects of the HMM theory that provides the means for the analysis of NE spiking activity. Under the HMM, a system of N recorded neurons is assumed to be in one of a predetermined number of hidden firing rate states. Each state is defined as a vector of N firing rates, one for each neuron. In each state, the neurons are assumed to discharge as stationary Poisson processes (i.e., the firing of neuron j in state i is fully described by its instantaneous firing probability E_{ij}). The transition matrix $A = \{A_{ij}\}$ gives the probability to move from state i to state j . The transition probabilities, A_{ij} , are independent of time (i.e., at any time, the probability of making a transition from state i to state j depends only on the identities of states i and j). Thus, the HMM model uses a discrete hidden state at time t to summarize all the information before t . In this sense, the hidden state time sequence in an HMM is a Markov chain. The HMM is fully described by the firing emission and the transition matrices: $\lambda = \{E, A\}$. To choose the model that better explains the data, we estimated the parameters λ that maximize the probability of observing the data given the model (training step). After finding the optimal model parameters, we computed the state sequence that maximizes the likelihood of the observed spike trains in a given trial (testing step).

Training step. In this section, we describe how to calculate the probability of observing the data given the model and the procedure to maximize it. Assume the system has M states. Let S_t be the state at time bin t . Let T be the length of the observation window. Neglecting the simultaneous firing of two or more of the N neurons, we build in each trial the observation sequence $O = o_1 o_2 \dots o_T$, with $o_t = 0$ if any of the neurons fired at bin t and $o_t = i$ if neuron i fired at bin t . This neglect is justified because the size of the bin was 1 ms and neither N nor the firing rate of the neurons were excessively high, making synchronous events rare: in the whole dataset, the probability of finding a bin containing more than one spike is equal to 0.0065. For bins containing spikes from more than one neuron, the observation was chosen randomly among the firing neurons.

Let π_i be the prior probability that the system starts in state i . We set $\pi = 1$ for state 1. The HMM is fully described by the parameters $\lambda = \{E, A\}$, with $E_{ij}(j) = P(O_t = j | S_t = i)$ and $A_{ij} = P(S_{t+1} = j | S_t = i)$. To estimate the emission matrix and the transition matrix, we used the Baum–Welch algorithm (Baum et al., 1970; Rabiner, 1989). The goal of this algorithm is to find the optimal model parameters λ , given an observation sequence O (training data), that maximize the likelihood $P(O|\lambda)$ (i.e., the probability of the observation sequence given the model). For finite observations, there is no analytical way of estimating the model parameters. However, the Baum–Welch algorithm can locally maximize the likelihood by using an iterative expectation maximization (EM) procedure.

In the expectation step, we need to calculate the probability $\xi_t(i, j)$ of being in state i at time t and state j at time $t + 1$, given the model parameters and the observation sequence. It is defined as follows: $\xi_t(i, j) = P(S_t = i, S_{t+1} = j | O, \lambda)$. To calculate this probability, we need to define two other probabilities: the forward probability and the backward probability. The forward probability $\alpha_t(i)$ is defined as the probability of the partial history of the observation sequence up to time t , noted

$O_1 O_2 \dots O_T$, with state at time t being i , given the model [i.e., $\alpha_t = P(o_1 o_2 \dots o_T, S_t = i | \lambda)$]. The forward probability can be calculated inductively as follows:

- (1) Initialization: $\alpha_1(i) = \pi_i \cdot E_i(o_1)$, $1 \leq i \leq M$
- (2) Induction: $\alpha_{t+1}(i) = E_i(o_t) \sum_{j=1}^M A_{ij} \alpha_t(j)$,
 $2 \leq t \leq T$, $1 \leq i \leq M$.

The probability of the complete history of the observation gives the likelihood $P(O|\lambda)$, which can be calculated by means of the forward probability as follows:

$$P(O|\lambda) = \sum_{i=1}^M \alpha_T(i). \quad (1)$$

The backward probability $\beta_t(i)$ is defined as the probability of the partial observation sequence from time t up to time T , $O_t O_{t+1} \dots O_T$, given that the state at time t is i and given the model [i.e., $\beta_t(i) = P(o_t o_{t+1} \dots o_T | S_t = i, \lambda)$]. The backward probability can also be calculated inductively as follows:

- (1) Initialization: $\beta_{T+1}(i) = 1$, $1 \leq i \leq M$
- (2) Induction: $\beta_t(i) = \sum_{j=1}^M A_{ij} E_j(o_{t+1}) \beta_{t+1}(j)$,
 $1 \leq t \leq T$, $1 \leq i \leq M$.

The probability $\xi_t(i, j)$ can be written in terms of the forward and backward probabilities as follows:

$$\xi_t(i, j) = \frac{\alpha_t(i) A_{ij} E_j(o_{t+1}) \beta_{t+1}(j)}{P(O|\lambda)}. \quad (2)$$

Next, let $\gamma_t(i)$ be the probability of being in state i at time t , given the observation sequence and the model; $\xi_t(i, j)$ and $\gamma_t(i)$ are related by the following:

$$\gamma_t(i) = \sum_{j=1}^M \xi_t(i, j) = \frac{\alpha_t(i) \beta_t(i)}{P(O|\lambda)}. \quad (3)$$

If we sum $\gamma_t(i)$ over time, we get the expected number of times that the system visits state i , or equivalently, the expected number of transitions from state i . If we sum $\xi_t(i, j)$ over time, we get the expected number of transitions from state i to state j . Thus, the model parameters $\lambda = \{E, A\}$ can be reestimated in the maximization step as follows:

$$A_{ij}^{\text{new}} = \frac{\sum_{t=1}^{T-1} \xi_t(i, j)}{\sum_{t=1}^{T-1} \gamma_t(i)} \quad (4)$$

$$E_i(j)^{\text{new}} = \frac{\sum_{t=1}^{T-1} \gamma_t(i)}{\sum_{t=1}^{T-1} \gamma_t(i)} \quad (5)$$

The numerator of the reestimated A_{ij} is the expected number of transitions from state i to state j , while the denominator is the expected number of transitions from i to any state. Then the ratio is the probability of transitioning from state i to state j , which is the desired value of A_{ij} . The numerator of the reestimated $E_i(j)$ is the expected number of times the system is in state i with the observation $o = j$, while the denominator is the expected number of times the system is in state i . The ratio is the probability of observing $o = j$, given that the system is in state i , which is the desired value of $E_i(j)$.

It has been shown (Dempster et al., 1977) that the reestimated model is more likely than the initial model: $P(O|\lambda^{\text{new}}) > P(O|\lambda)$. Thus, iterations of the reestimation leads to a maximum-likelihood estimate of the model. In our study, we stopped the reestimations once the increase in the log of the likelihood is less than a tolerance factor (10^{-6}) or if this

tolerance was not reached within a maximum number of iterations (500). Since the algorithm converges to a local maximum likelihood that is not necessarily the global maximum, we reran the reestimation algorithm five times, each time with new initial parameters, and we chose the model with the highest likelihood. The initial components of the emission matrix were random, while the components of the transition matrix were initialized as in Jones et al. (2007): diagonal elements were initialized in the range $D = 0.99-0.999$, and nondiagonal elements were all equal to $(1 - D)/(M - 1)$.

Testing step. Given the optimal parameters that were estimated in the training step and given an observation sequence O , two useful quantities can be obtained: the probability $\gamma_t(i)$ of being in state i at time t and the most likely sequence of states. The probability $\gamma_t(i)$ is calculated with Equation 3 by means of the forward and backward probabilities, evaluated with the optimal model parameters, and gives the time-varying probability of being in each state at any time bin. We used it to calculate the duration of transition periods for which none of the states had likelihood larger than a threshold equal to 0.8. The most likely sequence of states given the model and the observation sequence was calculated with the following algorithm, called Viterbi algorithm (Rabiner, 1989). Let $s_1 s_2 \dots s_t$ be the state sequence until time t , and $\delta_t(i)$ the sequence with highest probability given the first t observations and with $S_t = i$, in other words, the following:

$$\delta_t(i) = \max_{s_1 s_2 \dots s_{t-1}} \{P(S_1 S_2 \dots S_{t-1}, S_t = i, o_1 o_2 \dots o_t | \lambda)\}.$$

By using the Markov property, we get the following: $\delta_{t+1}(j) = \max_i \{\delta_t(i) A_{ij}\} \cdot E_j(o_{t+1})$. The Viterbi algorithm goes through the following induction:

- (1) Initialization: set $\delta_1(i) = \pi_i E_i(o_1)$ and $\varphi_1(i) = 0$,
where φ_t is a matrix of size $T \times M$ that stores the arguments that maximize $\delta_t(i)$.
- (2) Induction: $\delta_{t+1}(i) = \max_j \{\delta_t(j) A_{ji}\} \cdot E_i(o_{t+1})$ and
 $\varphi_{t+1}(i) = \text{argmax}_j \{\delta_t(j) A_{ji}\}$
- (3) Termination: $S_T^* = \text{argmax}_i \{\delta_T(i)\}$.

At $t = T$, we choose the highest probability argument, and then we backtrack from there to find the highest probability state sequence:

- (4) Backtracking: $S_t^* = \varphi_{t+1}(S_{t+1}^*)$, $t = T - 1, T - 2, \dots, 1$.

Selection of the number of states. For a given ensemble of neurons, the number of states is, in principle, determined by the number of degrees of freedom of the ensemble activity, the level of firing rate modulation, and the amount of available data. An important question concerns the optimal number of states of a given NE. The basic problem is that the model likelihood always increases with increasing number of states; however, a large number of parameters might impair the ability of the model to explain new data—thus, it would not be a “model” anymore. A classical approach used in model selection consists of penalizing the model likelihood by a measure of its complexity. The most popular approach uses the Bayesian information criterion (BIC) (Schwarz, 1978). The BIC for a model with M states is given as follows:

$$\text{BIC}(M) = \ln[P(O|\hat{\lambda}_M)] - \frac{\rho_M}{2} \ln(T),$$

where O is the observed data, $P(O|\hat{\lambda}_M)$ denotes the likelihood of the model with M states and estimated parameters $\hat{\lambda}_M$ using the EM algorithm, ρ_M is the number of free parameters of the model, and T the length of the observation data. For a model of M states and N neurons, the number of free parameters is equal to $M(M - 1) + MN$, the first and second terms being the contribution of the transition and emission matrices, respectively (notice that the prior do not add any free parameter, since $\pi = 1$ for state 1). For each NE, we calculated the BIC scores for

models with different number of hidden states and selected the model with maximum BIC score.

Generation of ideal observation sequences. We compared the distribution of transition durations obtained with the empirical dataset and the one obtained with a set of ideal observations generated with the HMM parameters that were estimated using the empirical dataset. We generated a long sequence of observations, $O = o_1 o_2 \dots o_T$, using the HMM parameters $\{E, A\}$, as follows: at $t = 1$, the model is initialized at state 1, o_1 is chosen according to the emission probabilities; next, the model transits to a new state according to the transition probabilities, and so on until $t = T$, where $T = 10^5$ ms. In these surrogate data, the model “neurons” change their firing rate as fast as possible and in a concomitant way; thus, the resulting distribution of transition durations represents a theoretical lower bound.

Surrogates of gradual state transitions. We further compared the empirical transition durations in the period containing the presentation of f_2 , calculated with HMMs with two states, to the ones obtained using ensembles of Poisson processes that linearly change from one state to another. Let λ be the two-state HMM that was estimated from the spiking activity of an ensemble of N recorded neurons using trials of the same trial type, during the period T . Here, the analysis period is delimited between $t(f_2) - 0.2$ s and $t(f_2) + 1$ s, where $t(f_2)$ is the onset time of f_2 . Let $\mathbf{R}_1 = \{r_1^1, \dots, r_N^1\}$ and $\mathbf{R}_2 = \{r_1^2, \dots, r_N^2\}$ be the two vectors containing the firing probabilities of each neuron in state 1 and 2 of model λ , respectively. We constructed a surrogate dataset of linearly ramping neurons by generating 100 realizations of nonhomogeneous Poisson processes with underlying rate functions defined as follows:

$$1 \leq i \leq N$$

$$f_i(t) = \begin{cases} r_i^1 & \text{if } 0 \leq t \leq \frac{T - \tau}{2} \\ \frac{r_i^2 - r_i^1}{\tau} \left[t - \frac{T - \tau}{2} \right] + r_i^1 & \text{if } \frac{T - \tau}{2} > t < \frac{T + \tau}{2} \\ r_i^2 & \text{if } \frac{T + \tau}{2} \leq t \leq T \end{cases}, \tag{7}$$

where τ is the period of ramping. Notice that all units reach the second state after τ , but their associated rate functions have different slopes, depending on $\mathbf{R}_2 - \mathbf{R}_1$.

Decoding significance. To assess statistical significance of the decoding performance, we calculated the probability of getting k hits by chance, which is given by the following: $\rho(k) = C_n^k p^k (1 - p)^{n-k}$, where n is the number of trials of the testing set and p is the probability of getting a hit by chance ($p = 1/2$). The p value can be derived as follows: $P_{\text{value}} = \sum_{i=k}^n \rho(i)$. Note that the p value is a function of both the number of hits and the number of tested trials, which can be different for different recording sessions.

Receiver operating curve analysis. We used the receiver operating curve (ROC) index to quantify the association between the activity of a neuron and the animal’s response during a given time window. The ROC index gives the probability that the neuronal response associated with one behavioral choice exceeds the neuronal response associated with the other choice. It is a measure of the experimenter’s ability to predict the monkey’s behavior from the neuronal response. It is estimated by the following equation:

$$\text{ROC index} = \int P(r_1 = x) P(r_2 < x) dx, \tag{8}$$

where $P(r_i)$ is the probability distribution of responses in condition i , where condition $f_1 < f_2$ is condition 1 and condition $f_1 > f_2$ is condition 2. This equation gives the proportion of trials in which the neuronal responses in condition 1 and 2 can be separated using a criterion x , over all possible values of x . This equation assumes that the variances of the responses are the same. We used a permutation test (Siegel and Castellan, 1988) (500 shuffles), in which the neuronal responses and the monkey’s

choices were randomly associated, to assess significant ROC indices that were different from 0.5, a value expected in the case in which the two response distributions are equal. Finally, time windows for which the mean firing rate of the neuron was < 5 Hz were left aside.

Results

Neuronal ensemble activity during the discrimination task

Two monkeys were trained to discriminate the difference in frequency between two consecutive mechanical vibrations delivered to one fingertip (Fig. 1A). The stimulus set used for the experiment (Fig. 1B) can be divided in three subsets: a first subset (set 1) in which the difference between the base frequency (f_1) and the comparison frequency (f_2) was large and equal to ± 8 Hz and two other subsets in which f_1 was constant while f_2 was variable (set 2) or vice versa (set 3). A trial type is defined by the values of the two applied frequencies. The evidence of a trial type is defined as the absolute difference between base and comparison frequencies, given by $\Delta f = |f_2 - f_1|$. The resulting psychometric curves are shown in Figure 1C. Monkeys were asked to report discrimination after a fixed delay period of 3 s, called the postponed decision period, between the end of f_2 presentation and the cue that triggered the motor report (PU). The reaction time is defined as the time between PU and the moment at which the animal releases the key (KU) to push one of the two response buttons. More than 99% of the RTs are > 200 ms for both monkeys (Fig. 1D), suggesting that the animals did not anticipate the appearance of the response signal PU. During performance of the task, we recorded the simultaneous single neuron activity from S1, S2, DPC, MPC, and M1 (Fig. 1E). Recordings in S1, S2, and DPC were made in the hemisphere contralateral to the stimulated hand, those in M1 were made in the other hemisphere, contralateral to the responding arm, and those in MPC were made either in the hemisphere contralateral or ipsilateral to the responding arm, noted MPCc and MPCi, respectively. We did not find qualitatively different results for MPCc and MPCi; thus, throughout this study, unless otherwise specified, these areas were analyzed together. A total of 264 NEs was analyzed ($n = 64, 58, 35, 49,$ and 58 from S1, S2, DPC, MPC, and M1, respectively). The number of simultaneously recorded neurons in each area was 3–15 neurons (median, 5 neurons).

In the present study, we first concentrated on the NE activity during the period between $t(f_2) - 0.5$ s and PU, where $t(f_2)$ is the onset time of f_2 . This period contains the entire f_2 presentation, where the decision-making process takes place, and the postponed decision period during which the monkey maintains the decision output until the movement is triggered. In Figure 1F, the activity of an example NE, composed of eight neurons from M1, is shown for two trials with $f_1 = 18$ Hz and $f_2 = 26$ Hz. The ensemble activity revealed a temporal pattern with some neurons increasing or decreasing their activity in a coordinated way. Interestingly, the coherent changes are appreciable not only in response to the stimulus but also during the postponed decision period. Nonetheless, the pattern is variable from trial to trial, indicating that NE activity ran with different time courses from trial to trial. To analyze these temporal patterns in a single-trial basis, we used the HMM.

Sequences of firing rate states in single trials

In the HMM, the NE activity is assumed to progress through a predetermined number of firing rate states. In each state, the neurons are assumed to discharge as stationary Poisson processes and the probability of going from one state to another is determined by a constant transition matrix (see Materials and Meth-




Figure 1. Discrimination task and neuronal ensemble activity. **A**, Timeline of the discrimination task. The mechanical probe is lowered, indenting the glabrous skin of one digit (PD); the monkey places its free hand on an immovable key (KD). The probe oscillates vertically, at the base frequency (f_1); after a fixed delay (3 s), a second mechanical vibration is delivered at the comparison frequency (f_2); after another fixed delay (3 s) between the end of f_2 and probe up (PU), the monkey releases the key (KU) and presses either the lateral or medial push button (PB) to indicate whether the comparison frequency was higher or lower than the base, respectively. **B**, Stimulus set during recordings. Each trial type is defined by a base/comparison frequency stimulus pair. The number inside the box indicates overall percentage of correct trials for each trial type, except for the ambiguous trial type, where $f_1 = f_2 = 22$ Hz, for which we plotted the percentage of trials the animal pressed the lateral push button. **C**, Psychophysical performance when f_1 was maintained fixed at 22 Hz and f_2 was variable. The black/gray points indicate the percentage of trials in which monkey 1/2 estimated f_2 higher than f_1 . Error bars indicate SE. **D**, Reaction time distributions for monkey 1/2 (black/gray traces). **E**, Recorded cortical areas: S1, Primary somatosensory cortex; S2, secondary somatosensory cortex; DPC, dorsal premotor cortex; MPC, medial premotor cortex; M1, primary motor cortex (notations used in all following figures). **F**, The series of black ticks represent the spike trains of 8 M1 neurons recorded simultaneously. The NE activity, during the period between $t(f_2) - 0.5$ s and PU is plotted for two trials of the same trial type ($f_1 = 18$ Hz and $f_2 = 26$ Hz). The gray area corresponds to the presentation of f_2 .

ods). We set the number of states to 3. The reason of this choice is justified below. Two HMMs were used to describe the activity of a given NE. The HMM parameters (transition matrix and firing rates in each state) of each of the two models were estimated using 15 randomly selected trials from condition $f_1 < f_2$ or $f_1 > f_2$, respectively, during the period between $t(f_2) - 0.5$ s and PU. The estimation of these parameters is called the training step, and the set of trials used for the estimation is called the training set. The set of remaining trials is called the testing set. After the training step, we obtained the sequence of states for all trials. This was done using the Viterbi algorithm that finds the most likely sequence of states given the model parameters and the spike trains (see Materials and Methods). Figures 2 and 3 show the state sequences for five trials from each condition and the associated firing rate states for one example NE from each sensory and frontal lobe cortices, respectively. The figures reveal many important features. First, the state sequence was robust across trials of the same experimental condition, but, in frontal lobe NEs, state sequences were different for different response conditions. Second,

in the somatosensory NEs, the activity went from a basal state to a state that was stable for the entire stimulus presentation period, whereas NEs of frontal lobes had a more complex, patterned, temporal dynamics with transitions occurring during the stimulus presentation and the postponed decision period. Importantly, the exact times of transitions varied more from trial to trial in frontal lobe cortices compared with transitions in sensory ones for which transitions are mainly determined by the stimulus onset and offset. In the next sections, we further investigate and evaluate these observations.

Response decoding using the neuronal ensemble activity in single trials

In this section, we investigated how specific NE dynamics are to the monkey’s decision (i.e., $f_1 < f_2$ or $f_1 > f_2$). For this, we measured the ability of the model to predict the monkey’s response in a single trial given the NE activity, by using the following jackknife cross-validation procedure. Let λ be the HMM that was trained using trials from condition $f_1 < f_2$ and λ' the HMM




Figure 2. State sequences of sensory NEs. For each NE, two HMMs were estimated using 15 trials from condition $f1 < f2$ and $f1 > f2$, respectively. One example NE activity of S1 (**A**) and S2 (**B**) are shown. For each condition, the ensemble spiking activity (black ticks) is shown together with the most likely sequence of states (each color represents a state visitation) for five trials. The histograms below the state sequences represent the three firing rate states of each of the two models. Each state is defined as a vector of N firing rates, one for each neuron of the ensemble (in these examples: $N = 7$ for both S1 and S2 NEs). Notice that, in **A** and **B**, light blue and yellow states appeared very transiently.

that was trained using trials from condition $f1 > f2$. Only correct trials were used and, to avoid variations in task difficulty, we analyzed the NE activity during trials from stimulus set 1. For each trial of the testing set, we compared the posterior probabilities that the NE activity was generated by either model λ , noted P_λ , or by model λ' , noted $P_{\lambda'}$. These probabilities were calculated using Equation 1 in Materials and Methods, with T being the length of the period between $t(f2) - 0.5$ s and PU. If $P_\lambda > P_{\lambda'}$, the predicted response was $f1 < f2$; otherwise, the predicted response was $f1 > f2$. The decoding performance was defined as the percentage of correct predictions (hits) and compared with the expected chance level (see Materials and Methods). A decoding performance that is significantly ($p < 0.05$) above chance indicates that HMM parameters were different for different response conditions. We calculated the decoding performances for all NEs of each area (Fig. 4A, blue distributions). We found that 4.7, 17.2, 57.1, 57.1, and 89.7% of the NEs from S1, S2, DPC, MPC, and M1, respectively, had a significant decoding performance, ranging between 62–78 and 62–100% for S2 and frontal lobe areas, respectively. Significant decoding was found in both hemi-

spheres of MPC in similar proportions (7 of 10 NEs from MPCc and 21 of 39 NEs from MPCi).

To preclude the possibility that the NE activity could represent movement parameters such as movement direction instead of the result of the comparison between $f1$ and $f2$, we did the same analysis (both training and testing) with trials from a control version of the task. In this control task, a light indicating the button to be pressed was illuminated at the beginning of the trial, whereas the same stimuli were applied (set 1) as in the discrimination task (see Materials and Methods). Thus, in this task, frequency comparison was not needed to perform the appropriate motor response. The resulting decoding performances (Fig. 4A, gray distributions) exceeded rarely the significance level ($\leq 2.3\%$). This suggests that NE activity carried information about the result of the frequency comparison and not about the motor output. The number of trials of the training set was the same for the control and discrimination tasks, but the testing set often contained fewer trials in the control task. To test whether the different decoding performances were produced by differences in the size of the testing sets, we decoded the discrimination trials with testing subsets of 20 randomly chosen trials (10 from each condition), equivalent to the minimal size of the testing set among all sessions of the control task. Overall, the use of fewer test trials reduced the percentages of significance but yielded the same qualitative differences between decoding performances in the discrimination and control task (Fig. 4A, see percentages in parenthesis). In addition, to test

whether the decoding performance depends on the number of states, we did the same analysis by using HMMs with four states (Fig. 4B). We found that the use of three or four states to decode the animal's response yielded qualitatively the same results.

Next, we asked whether error trials could be predicted by the HMM. For this, we selected the S2 and frontal lobe NEs that had significant response decoding performance and predicted the response in error trials using the same HMM parameters that were estimated with the training set composed of correct trials. The decoding scheme was similar to the previous one, but with an inverted rule: if $P_\lambda > P_{\lambda'}$ ($P_\lambda < P_{\lambda'}$), the predicted response was $f1 > f2$ ($f1 < f2$). Thus, with this scheme significant decoding indicates that the NE activity during error trials resembles the activity in correct trials, in which the animal pressed the opposite button. Only sessions with > 10 error trials were analyzed. The resulting decoding performance was compared with the one previously obtained with correct trials (Fig. 4C). We found that significant decoding of error trials, given significant decoding of




Figure 3. State sequences of frontal lobe NEs. For each NE, two HMMs were estimated using 15 trials from condition $f1 < f2$ and $f1 > f2$, respectively. In **A–C**, one example NE activity of each frontal lobe cortical area is shown. For each condition, the ensemble spiking activity (black ticks) is shown together with the most likely sequence of states (each color represents a state visitation) for five trials. The histograms below the state sequences represent the three firing rate states of each of the two models. Each state is defined as a vector of N firing rates, one for each neuron of the ensemble (in these examples: $N = 6, 11, 8$ for DPC, MPC, M1, respectively).

correct trials, is increasingly more likely for S2, DPC, MPC, and M1 NEs.

Comparison-selective sequences of states

As shown above, many ($>50\%$) NEs from frontal lobe areas progressed through different sequences of states depending on the actual computation of the difference between $f1$ and $f2$. To further investigate which features of single-neuron responses generate the selective sequences of states, we quantified how well the two response conditions could be discriminated given the activity of individual neurons. For this, we calculated the area under the receiver operating curve (ROC index), which measures the difference between the distributions of the neuronal responses in both conditions for both discrimination and control tasks (see Materials and Methods). Periods during which the ROC index was significant were defined as periods of differential activity. Using sliding windows of 0.5 s, we calculated the percentage of neurons with significant ROC indices (Fig. 5A, black traces), which increased rapidly after the onset of $f2$. During the entire postponed decision period, many (20–44%) neurons presented a differential activity. In contrast, substantially less differential activity was detected in the control task (Fig. 5A, gray traces).

To investigate the dynamics of differential activity, we selected those neurons that had at least one significant ROC index at any time window between $f2$ onset and PU. As shown in Figure 5B, differential activity appeared transiently in some neurons at the onset of $f2$, but this feature appeared later in the delay for many

other neurons. Note that only few neurons maintained a differential activity for a long period of time. Moreover, the heterogeneity was prominent in simultaneously recorded nearby neurons. The time-varying ROC indices of an M1 ensemble (composed of seven simultaneously recorded neurons) and an MPC ensemble (composed of nine simultaneously recorded neurons) are presented in Figure 5, **C** and **D**, respectively. Again, significant ROC values were reached at different moments for different neurons. Together, this indicates that differential activity was initiated in different groups of neurons at different times, a feature that necessarily leads to different paths of ensemble firing rates for the two different response conditions, thus giving rise to comparison-selective sequences of states.

Single-trial ensemble activity as a Markov chain

Our results demonstrate that the HMM can be used to decode the NE activity in single trials. However, this does not prove that the ensemble activity truly switched between different states. To test this, we first estimated the number of HMM states, and, second, we studied activity at the transitions.

Up to now, we have used HMMs with three states. An important question concerns the optimal number of states. We addressed this issue by using the BIC (see Materials and Methods), which penalizes the likelihood of a model with M states by a measure of its complexity (Fig. 6A). For each NE, we calculated the BIC scores for models with different number of hidden states, with M varying from $M = 1$ to $M = 6$, and selected the model with




Figure 4. Response decoding. **A**, The response decoding performances, using HMMs with three states, were calculated for all analyzed NEs from each cortical area in the discrimination task (blue distribution) and the control task (gray distribution). Red vertical line, Chance level (50%). The percentage of significant performances is indicated for both the discrimination (blue) and the control task (gray). Values in parentheses indicate the percentage of significant performances obtained using a reduced testing subset of discrimination trials (see text for details). **B**, Same as **A**, but using HMMs with four states. **C**, The different panels compare the decoding performance in error trials with the one obtained in correct trials (see text for details). Only S2 and frontal lobe NEs that had significant decoding performance in correct trials and were recorded in sessions with ≥ 10 error trials were used. Gray points, Significant decoding in error trials. The percentages of significant decoding on error trials are indicated.




Figure 5. Dynamics of differential activity. **A**, ROC indices were calculated in sliding windows of 500 ms, which were shifted in steps of 50 ms. The percentages of neurons with significant ROC indices are plotted for the frontal lobe cortices, in the discrimination task (black) and the control task (gray). **B**, For each area, we selected those neurons that had at least one significant ROC index between f2 onset (red vertical line) and PU, and sorted them by the time at which the first significant ROC index appeared after f2 onset. **C, D**, ROC indices of seven simultaneously recorded M1 neurons (**C**) and nine simultaneously recorded MPC neurons (**D**) (colored lines). The circles represent significant values. Time windows during which the mean firing rate of the neuron was < 5 Hz were left aside.

maximum penalized likelihood, in the observation window between $t(f2) - 0.5$ s and PU. The frequencies of selecting a number of states are shown in Table 1. The BIC procedure indicated that $> 85\%$ of the NEs were optimally modeled by two to four states, with three being the number of states most often selected (44.3%). This number of state was most often selected for frontal lobe and S2 NEs (37.9–57.1%), whereas most of the S1 NEs had BIC scores peaking at two states (40.6%).

We next tested whether the states were stable and represent different patterns of activity. For this, we constructed a three-state HMM for each trial type, using all correct trials for the estimation of the model parameters. We calculated, for each trial, the probability of each state at any moment in time using the corresponding HMM (see Materials and Methods). In this way, the evolution of the probability of each state can be followed in time. These probabilities are shown in Figure 6B for an example NE from M1. Most of the time one state dominated upon the others (i.e., one state had a high probability, whereas the two other states had a probability close to zero). We defined the periods during which none of the states had a probability > 0.8 as transition periods (Fig. 6B, inset). Only tran-

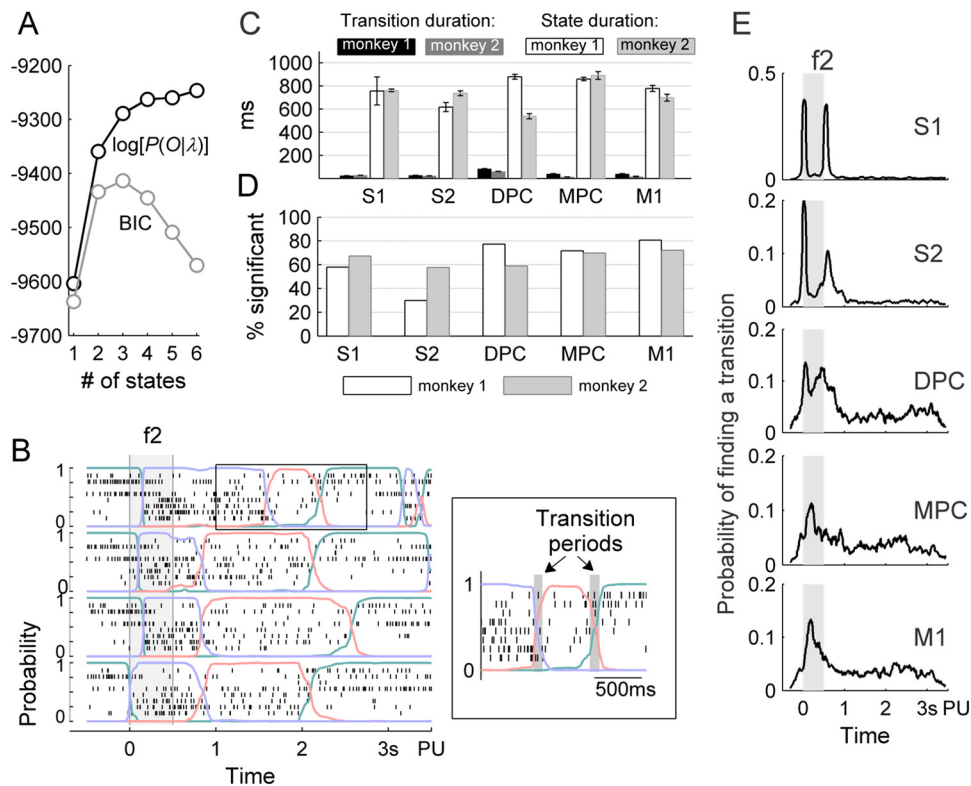


Figure 6. Ensemble activity as a Markov chain. **A**, The number of HMM states was estimated using the BIC procedure, which penalizes the likelihood of the model (black) by a measure of its complexity. The estimated number of states is the one for which the BIC score (gray) is maximal. Shown are data from an example ensemble of M1 neurons. **B**, The spiking activity of the example ensemble of M1 neurons is presented together with the likelihood of each state (colored lines) for four trials. The periods during which none of the states had likelihood larger than 0.8 were defined as transition periods (gray area). The inset corresponds to the epoch framed in **B**. **C**, The bars indicate the averaged duration of transition periods and averaged state duration for all analyzed NEs of a given cortical area and for each monkey. Error bars indicate SE. **D**, The Viterbi state sequences were determined for each trial with $f_1 = 26$ Hz and $f_2 = 34$ Hz, and the percentage of neurons with significantly different mean firing rates in different states ($p < 0.01$, ANOVA) was calculated (similar results were obtained using other trial types). **E**, For each cortical area, the probability of finding a transition was calculated using a sliding window of 100 ms (data from both monkeys). The shaded area represents the application of f_2 .

Table 1. Frequencies of selecting a number of hidden states using the BIC criterion for correct trials with $f_1 = 26$ Hz and $f_2 = 34$ Hz in the observation window between $t(f_2) - 0.5$ s and PU

	No. of states						No. of neurons (median)
	1 (%)	2 (%)	3 (%)	4 (%)	5 (%)	6 (%)	
S1 ($n = 64$)	0	40.6	31.3	17.2	4.7	6.2	5
S2 ($n = 58$)	0	24.1	37.9	20.7	13.8	3.5	6
DPC ($n = 35$)	14.3	28.6	45.7	11.4	0	0	5
MPC ($n = 49$)	6.1	18.4	57.1	14.3	2.0	2.0	5
M1 ($n = 58$)	5.1	13.8	53.5	13.8	10.3	3.5	5
All ($n = 264$)	4.2	25.4	44.3	15.9	6.8	3.4	5

Maximal frequencies are indicated in bold. The last column indicates the median number of neurons in the NEs of each cortical area. n , Number of NEs.

sition periods separating different states were analyzed; transitions flanked by the same state were discarded. The averaged transition period was calculated for all analyzed NEs (Fig. 6C, black and dark gray bars). Generally, the mean transition period was < 85 ms. This value was practically one order of magnitude smaller than the averaged lifetime of a state (Fig. 6C, white and light gray bars) that ranged between 538 and 890 ms. Hence, the transition periods were substantially shorter than the epochs during which the NE was in one state.

To test whether the firing rate states were different, we calculated the percentage of neurons that had significantly different firing rates in different states ($p < 0.01$, ANOVA comparing the three mean firing rates during each of the three states) in trials in which $f_1 = 26$ Hz and $f_2 = 34$ Hz (similar results were obtained

using other trial types). Indeed, the firing rate of many (30–80%) neurons significantly changed between states (Fig. 6D). Thus, individual neurons increased or decreased their activity in more than one state. These results confirm that the NE activity can be fairly well described by a sequence of distinct states separated by short transitions reflecting coordinated changes in the firing rates of many neurons.

Transition density

Next, we computed the probability of finding a transition at any moment in time in each cortical area. For all NEs, the state sequences were computed for all trials using the Viterbi algorithm. The exact times of transitions were stored and the probability of finding a transition was calculated. As shown in Figure 6E, transitions in somatosensory cortices were highly time-locked to the stimulus onset/offset, indicating that NE activity switched from a basal state to a state that was stable for the entire stimulus presentation period. When the stimulation was turned off, the NE activities of S1 and S2 behaved differently: that from S1 returned to the basal state, whereas that from S2 continued after stimulus offset. Transitions in frontal lobe NEs occurred both during the comparison stimulus and the postponed decision period. Generally, the transitions in frontal lobe NEs during f_2 presentation were more widely distributed than transitions in somatosensory NEs. In the next sections, we studied in more details the speed and the trial-by-trial variability of transitions during the comparison stimulus,

Table 2. Frequencies of selecting a number of hidden states using the BIC criterion in the observation window between $t(f_2) - 0.2$ s and $t(f_2) + 1$ s

	No. of states				Evidence
	1 (%)	2 (%)	3 (%)	4 (%)	Δf (Hz)
S1 ($n = 64$)	3.5	69.5	22.7	4.3	2
	2.3	62.9	26.9	7.8	4
	2.0	57.8	32.0	8.2	6
	0.8	55.8	31.4	12.0	8
S2 ($n = 58$)	10.8	60.3	26.3	2.6	2
	5.6	62.5	23.7	8.2	4
	4.3	53.0	30.6	12.0	6
	3.6	53.8	35.0	7.6	8
DPC ($n = 18$)	33.3	56.9	9.7	0	2
	19.4	68.1	12.5	0	4
	6.9	75.0	16.7	1.3	6
	16.7	67.2	15.0	1.1	8
MPC ($n = 20$)	37.5	53.8	7.5	1.3	2
	26.3	65.0	8.8	0	4
	18.8	70.0	11.2	0	6
	18.0	68.5	10.5	3.0	8
M1 ($n = 40$)	33.8	56.9	6.9	2.5	2
	30.0	56.9	10.0	3.1	4
	20.6	63.8	10.62	5.0	6
	19.8	63.0	14.3	3.0	8

Models were estimated using all correct trials from each trial type, for one to four states. The frequencies of selecting a number of states were calculated among trials of same evidence ($\Delta f = |f_2 - f_1|$). Maxima are indicated in bold. n , Number of NEs.

emphasizing the differences between frontal lobe and somatosensory cortices.

Fast coherent transitions

We concentrated on the first transition after the onset of the comparison stimulus. First, we selected all somatosensory NEs (64 and 58 NEs from S1 and S2, respectively) and those frontal lobe NEs that met the following criterion. Frontal lobe NEs were kept for the analysis if they presented a significant response classification, with the same HMM decoding scheme as previously but using the NE activity between $t(f_2) - 0.2$ s and $t(f_2) + 1$ s and using two states. This criterion ensures that the kept frontal lobe NEs underwent a decision-related transition in the analyzed period. A total of 78 frontal lobe NEs met the selection criterion (18 of 35, 20 of 49, and 40 of 58 NEs from DPC, MPC, and M1, respectively). Second, for all selected NEs, two-state HMMs were estimated for each trial type of the entire stimulus set, except those of the ambiguous trial type where $f_1 = f_2 = 22$ Hz (the actual button the monkeys pressed at the end of ambiguous trials ($f_1 = f_2$) was not stored. For this reason, ambiguous trials were not analyzed). Hence, for each NE, 22 HMMs were estimated. The choice of two states is supported by the BIC procedure, which most of the time selected models with two states (53–75%; Table 2). Figure 7A shows examples of the obtained state probabilities and state sequences. Before the onset of f_2 , the NE activity was in a first state, and then, after f_2 onset, it reached a second state. The fast transitions were produced by rapid coordinated changes in the activity of the neurons, as demonstrated in Figure 7, B and C, with an example of two neurons from the same MPC NE. These two neurons presented a complementary activity: one neuron increased its activity just after the transition (red ticks in each trial), whereas the activity of the other neuron rapidly decreased at the same moment. Importantly, the transition was calculated using the activity of all neurons from the ensemble (four in total).

To evaluate the speed and the coherency of the transition, we first computed the averaged transition durations for all selected

NEs in trials of same evidence, given by $\Delta f = |f_2 - f_1|$ (Fig. 8A, B, blue traces). The evidence is a measure of the difficulty of a given trial type. In principle, the duration of transition periods is determined by (1) the firing rate differences between states, (2) the number of neurons that changed their firing rate during transitions, and (3) whether or not the neurons changed their firing rate at the same time. To evaluate which factors determined the speed of the observed transitions, we compared the transition periods against the ones obtained with two sets of pseudodata.

We generated a first dataset of ideal observations by simulating Markov chains with the same HMM parameters that were estimated with the experimental spike trains (see Materials and Methods). The mean transition period obtained with this dataset gives the minimum expected value, given the parameters of the HMM. The second pseudodataset was obtained by combining randomly, for each trial type, the responses of the neurons, called shuffled dataset. This procedure keeps the same peristimulus time histogram (PSTH) profile for each trial type, but washes out the firing rate correlations that were not time-locked to the stimulus, as if neurons were no more simultaneously recorded. The HMM parameters were estimated in the shuffled dataset, and the mean transitions period was computed. We found that averaged empirical transition periods in frontal lobe NEs were not significantly different from the averaged ideal transition period ($p > 0.01$), whereas they were significantly shorter than the averaged transition periods of shuffled data ($p < 0.01$, differences in transition periods ranging between 11 and 19 ms; Fig. 8A); by contrast, in somatosensory NEs, the averaged transition periods were similar for all datasets, with differences of 1–2 ms (Fig. 8B). This indicates that the activity of the neurons “jumped” from one state to another very abruptly (given their firing rates) and at the same time; however, the jumps were time-locked to the stimulus in sensory NEs, whereas they occurred at variable times in frontal lobe NEs.

We further evaluated how fast were the state transitions with respect to gradual changes of the neuronal activity. For this, we simulated nonhomogeneous Poisson processes with underlying rates given by a piecewise function for which two stable configurations of rates are linked by a linear ramp (see Materials and Methods). The time of ramping was controlled by the parameter τ . The two rate configurations were the ones given by the HMM states that were estimated with the experimental spike trains, for each trial type. The state probabilities and transition periods were calculated as previously. By varying τ , we get the expected transition period for linear ramps, shown in Figure 8 (color plots). Hence, the empirical transition periods of the recorded NEs were shorter than the expected transition periods for ramps that last at least 50 ms. Longer ramps led to significantly longer transition durations ($p > 0.05$, t test).

Easy trials versus difficult trials

We then investigated the time at which these transitions occurred for the different degrees of task difficulty. For this, we calculated the transition latency, defined as the time elapsed between the onset of f_2 and the time of the transition detected using the Viterbi algorithm. Next, the mean latency and the SD of the transition latencies were computed. We found that both the mean latency and the variability in frontal lobe NEs decreased for increasing values of Δf (Fig. 9), indicating that the transitions were shorter and less variable for easy trials than for difficult trials. Note that changes in variability are not trivially explained by variability on the value of the comparison frequency; in fact, latency variability was minimal for trials of maximal evidence,

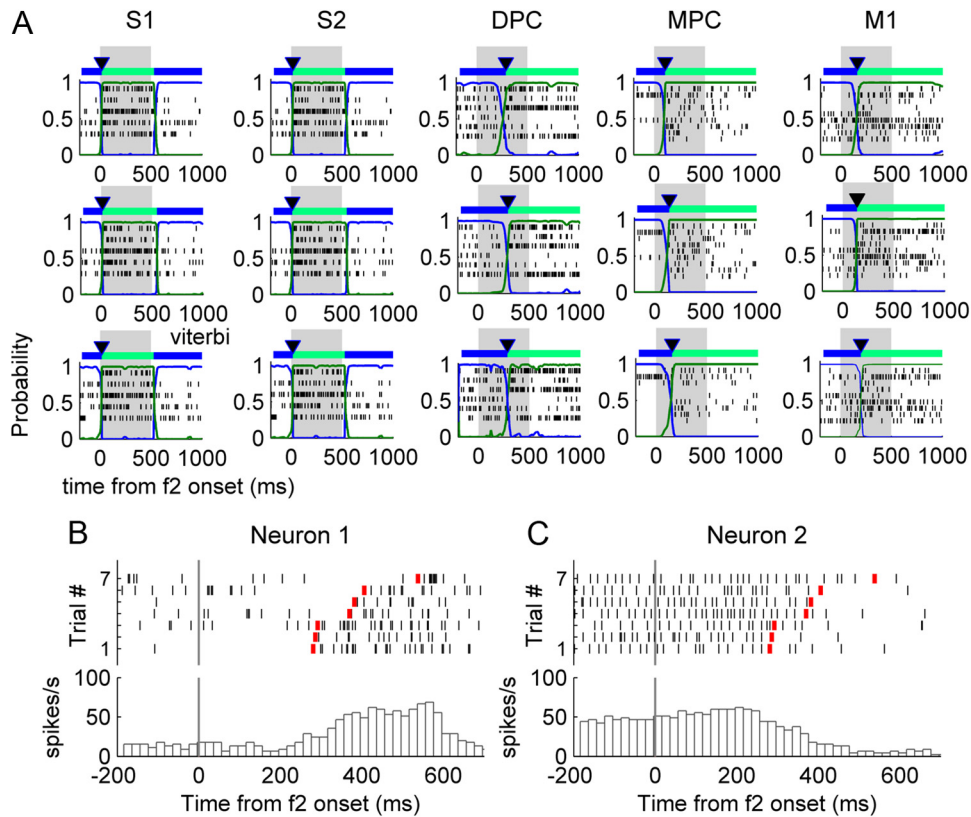


Figure 7. Coherent transitions. **A**, The spiking activity (black ticks), together with the probabilities of each state (blue and green traces) and the Viterbi sequence of states (top) are shown for one example NE from each cortical area, for three trials of the same trial type. The shaded gray region indicates the presentation of f_2 . **B**, **C**, The raster displays (top) represent the spiking activity of two neurons from the same MPC NE (consisting of 4 neurons recorded simultaneously), for seven correct trials. Each black tick corresponds to a spike, and the red ticks correspond to the time of the first transitions after f_2 onset (vertical gray line). Trials were rearranged according to increasing transition times. The PSTHs of the neurons are plotted in the bottom panels.

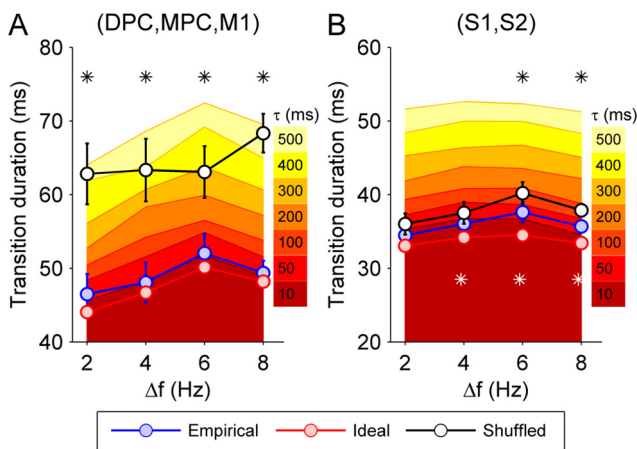


Figure 8. Transition duration for the empirical, ideal, and shuffled datasets. For each value of Δf , the averaged transition period between states was calculated for the empirical dataset and the two pseudodatasets, for sensory (**A**) and frontal lobe NEs (**B**). The three averaged transition periods were compared, for each Δf (ANOVA). The red stars indicate significant differences between empirical and ideal averaged transition periods ($p < 0.01$, Tukey’s test). The white stars indicate significant differences between empirical and shuffled averaged transition periods ($p < 0.01$, Tukey’s test). The color plot indicates the expected transition durations for artificial ramps with different periods of ramping τ . Artificial ramps were constructed using nonhomogeneous Poisson processes that linearly changed from a constant firing rate state 1 to a constant firing rate state 2. The two rate configurations were the ones given by the HMM states that were estimated with the experimental spike trains, for each NE and for each trial type.

$\Delta f = 8$ Hz, for which f_2 had the wider range of variation (from 10 to 34 Hz). Thus, variability reduction depends on the combination of past (f_1) and present (f_2) stimuli. In contrast, both mean latency and variability in somatosensory areas did not depend on Δf . Both the mean latency and the mean SD were significantly ($p < 0.05$, t test) larger for $\Delta f = 2$ Hz than for $\Delta f = 8$ Hz in the three frontal lobe cortices (except for the SDs in DPC), but not in the somatosensory areas.

Latency correlations

We next wondered whether or not the transition latencies in different cortical areas were correlated. Indeed, we may expect that the transitions in different cortical areas covary in single trials, because they are governed by the same overall decision process. Some of the previously selected NEs were recorded simultaneously allowing us to investigate the pairwise correlations between transition latencies of two cortical areas (190 pairs of NEs). For a given pair of NEs, we calculated the correlation using the set of trials for which $\Delta f = 8$ Hz. As shown in Table 3, the correlation seldom was significant ($p < 0.05$, Pearson’s linear correlation). We conclude that transition latencies from different cortical areas did not correlate.

Discussion

In this study, we demonstrated that the temporal dynamics of NE activity can be approximated by a sequence of distinct firing rate states separated by short transitions that can occur during both stimulus presentation and postponed decision period. The HMM

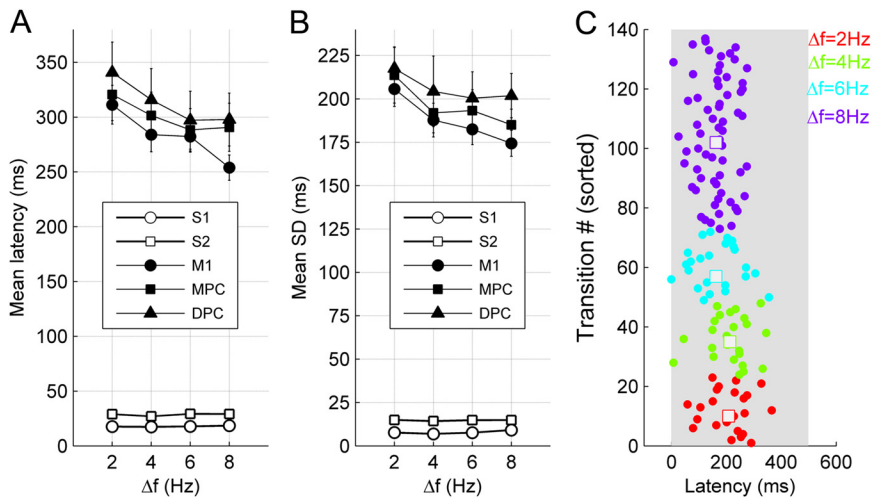


Figure 9. Transition statistics and task difficulty. The averaged transition latency (A) and transition variability (B) were calculated for the NEs of each cortical area and for each value of Δf . Error bars indicate SE. C, The transition latencies for an example NE from MPC were grouped according to the evidence (Δf , color code) of the trials in which they occurred. Each filled point represents a trial; the squares indicate the mean latency for each value of Δf . Note the decrease of both latency and variability for increasing values of Δf . The shaded area represents the application of f_2 .

Table 3. Number of significant ($p < 0.05$) pairwise correlations between transition latencies in simultaneously recorded NEs from different cortical areas

	S2	DPC	MPC	M1
S1	4 of 55	1 of 15	1 of 17	0 of 24
S2		2 of 16	1 of 17	4 of 25
DPC			0 of 5	2 of 5
MPC				0 of 11

assumes that the dynamics of NE activity is Markovian: state transitions depend only on the preceding state. However, cell features, such as the refractory period, bursting, and short-term plasticity, introduce a history dependency that makes the HMM assumption impossible to meet. Nevertheless, as pointed by Jones et al. (2007), temporal dependencies enhance the HMM performance (Antón-Haro et al., 2001), and, thus, the Markov assumption is rather conservative. Moreover, the HMM assumes that neurons are independent; however, dynamic changes in correlation between neurons in absence of firing rate differences have been reported previously (Vaadia et al., 1995; Riehle et al., 1997). A novel method for including correlations in HMMs was recently developed (Katahira et al., 2010) and offers new promises for investigating this fundamental question. Furthermore, the model clusters the data into a finite number of firing rate states. An important question concerns the optimal number of states of a given NE. In the present study, we used a penalized likelihood criterion, to estimate the number of statistically discernible states, given the number of recorded neurons, the firing rates, and the amount of data. However, the number of states should not be confused with the “real” number of states of collective activity in a given cortical area.

This study shows that distinct cortical areas contribute to the sequence of the processing linking decision making and action. One could argue that frontal lobe activity during the postponed decision period reflects other processes, such as movement preparation or time estimation. This seems unlikely, however, because (1) when the same movements were guided by visual cues the sequences of states were no more selective (Fig. 4A,B); (2) reaction time distributions did not show anticipations (Fig. 1D),

suggesting that the animals neither prepared the movement before the response signal PU nor anticipated its occurrence; and (3) the two response buttons were close to each other, so that neuronal activity related to the two arm movements should be similar. Single-neuron delay differential activity has been interpreted as related to the memory of the decision (Lemus et al., 2007; Hernández et al., 2010), even though frontal lobe cortical areas are not traditionally believed to be implicated in memory (but see Ohbayashi et al., 2003). Here, we showed that the storage of the decision is a collective dynamic process: neurons present transiently a differential delay activity, so that, as a result, the NE activity goes through two different paths depending on the animal’s decision and with variable time courses from trial to trial. Thus, the output of the decision might be maintained via a selective transient activation of groups of neurons. Hence, the mechanism sustaining the decision output could differ from the classical view of working memory (i.e., by means of a steady state of persistent activity). Consistent with our findings, there is empirical evidence that working memory can be stored by sequential activation of neural subpopulations (Baeg et al., 2003). In this sense, decision maintenance would be represented in a spatio-temporal pattern or “neural trajectory,” as proposed in the context of odor representation and decision making in invertebrates (Briggman et al., 2005; Mazor and Laurent, 2005), navigation-based decisions in mice (Harvey et al., 2012), and motor planning in monkeys (Yu et al., 2009; Afshar et al., 2011). Indeed, in our study, low-dimensional neural trajectories are coarsely described by sequences of firing rate states.

Martínez-García et al. (2011) recently proposed that information about the decision output is maintained during the postponed decision period in the synapses of MPC neurons, via synaptic facilitation. This model explains the behavior of those neurons for which differential activity vanishes during the delay. However, some MPC neurons do maintain the differential activity during the delay (Fig. 5). Thus, possibly both subthreshold and suprathreshold mechanisms contribute to the storage of the decision output during the postponed decision period.

The fact that the animal’s response could be decoded in error trials, using the same HMM parameters estimated with sets of correct trials, suggests that the NE activity fairly represents a commitment toward one response action; particularly in M1 that represents a final step in the decision-to-action conversion.

We showed that, after presentation of f_2 , frontal lobe NEs undergo an abrupt state transition occurring at variable times from trial to trial (Figs. 7, 8). Frontal lobe areas are engaged in a binary decision-making process, as revealed by our decoding scheme and by previous studies on single-neuron activity (Hernández et al., 2002, 2010). Theoretical studies have proposed different attractor-based models of decision making, in which two populations of neurons, each encoding one of the two possible decisions, compete in a winner-take-all fashion (Brunel and Wang, 2001; Wang, 2002; Wong and Wang, 2006; Wong et al., 2007). Competition leads the network to reach one of two decision states in which the activity of either one of the populations

exceeds the activity of the other. This model has two operational modes (Deco et al., 2007, 2009; Martí et al., 2008; Miller and Katz, 2010): a bistable mode, in which the external input destabilizes the basal state forcing the network to relax to one of the decision states, and a multistable mode, in which the stimulus does not destabilize the basal state but enhances the probability that large fluctuations induce a transition toward one of the decision states. Although these modes have been extensively studied, their predictions are difficult to test using the activity of single neurons. Here, we used the activity of simultaneously recorded neurons to detect the state transitions in single trials, via the HMM. This approach was used by Miller and Katz (2010) with simulated data, providing valuable predictions. Specifically, the multistable mode leads to noise-driven transitions that are characterized by abrupt changes in the firing rate of the neurons that occur at random times. By contrast, in the bistable mode, the firing rate of the neurons changes gradually due to the relaxation mechanism. Hence, the observation of sharp and variable transitions reported here favors the multistable scenario.

Smooth “ramping” activity is believed to implement an accumulation over time of sensory evidence for the different choices. Note that averaging abrupt firing rate changes can lead to an artifactual ramping activity. This is clear in Figure 7, *B* and *C*: the trial-averaged firing rate of a neuron can be either gradually increased or decreased; but this is only due to the summation of steps of firing rate occurring at variable times across trials, but concomitantly among neurons. Using surrogates, we showed that transitions following *f2* presentation were indeed indistinguishable from ideal jumps. Furthermore, surrogates presenting a linear ramping activity that last longer than 50 ms led to transition durations longer than the empirical ones. Our study, however, does not intend to exclude ramping activity in single trials, which has been reported in the lateral intraparietal area during random-dot direction discrimination (Huk and Shadlen, 2005), but suggests that, in the present task and for the analyzed cortical areas, single-trial dynamics are better represented by fast state transitions. Alternatively, one could argue that the transitions seen in frontal lobe areas after the presentation of *f2* (which have a mean latency of 280 ± 49 ms, for $\Delta f = 8$ Hz) may reflect the result of the comparison rather than the decision process itself.

In addition, we found that transition statistics in frontal lobe NEs depend on the difficulty of the task. It is known that single-neuron responses during motion discrimination initiate earlier on easy trials than on difficult trials (Kim and Shadlen, 1999). Our analysis confirmed this effect and extended it: both the latency and the variability of NE activity increased on difficult trials. These changes in variability are expected for a decision-making network operating in the multistable mode (Rolls et al., 2010). Thus, our results add evidence that the underlying mechanism of decision making is stochastic, a view supported by psychophysical and fMRI measures (Deco et al., 2007; Rolls et al., 2010).

Our study shows that decision processes and their related output are distributed over several cortical areas. The question arises of how different cortical areas communicate to solve the overall task. Anatomical data show that somatosensory and frontal lobe cortices are either directly or indirectly interconnected and share inputs (Pons et al., 1987; Luppino et al., 1993; Iwamura, 1998; Liu et al., 2002; Boussaoud et al., 2005). Thus, one possibility is that consistent decisions arise from correlated fluctuations among the cortical areas. We found, however, that transition latencies in pairs of NEs were practically uncorrelated. The lack of correlation might be due to the small number of neurons of the NEs used in

the present study. Hence, to address the question of how distributed sensory and decision networks communicate through large-scale interactions requires further theoretical and experimental investigations, possibly using other neural signals such as local field potentials.

In conclusion, our results suggest that decision-making processes and the subsequent maintenance of the decision output in frontal lobe cortices may be linked by a cascade of transient events. The cascade would be initiated by a noise-driven transition toward one of two spatiotemporal patterns of activity that maintains the output of the corresponding decision. In this view, transient dynamics are a key feature for unraveling the mechanism underlying the conversion of a decision into an action.

References

- Afshar A, Santhanam G, Yu BM, Ryu SI, Sahani M, Shenoy KV (2011) Single-trial neural correlates of arm movement preparation. *Neuron* 71:555–564.
- Antón-Haro C, Fonollosa JAR, Fauli C, Fonollosa JR (2001) On the inclusion of channel's time dependence in a hidden Markov model for blind channel estimation. *IEEE Trans Veh Technol* 50:867–873.
- Baeg EH, Kim YB, Huh K, Mook-Jung I, Kim HT, Jung MW (2003) Dynamics of population code for working memory in the prefrontal cortex. *Neuron* 40:177–188.
- Baum LE, Petrie T, Soules G, Weiss N (1970) A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann Math Stat* 41:164–171.
- Boussaoud D, Tanné-Gariépy J, Wannier T, Rouiller EM (2005) Callosal connections of dorsal versus ventral premotor areas in the macaque monkey: a multiple retrograde tracing study. *BMC Neurosci* 6:67.
- Brigman KL, Abarbanel HD, Kristan WB Jr (2005) Optical imaging of neuronal populations during decision-making. *Science* 307:896–901.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100.
- Brody CD, Hernández A, Zainos A, Romo R (2003) Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex* 13:1196–1207.
- Brunel N, Wang XJ (2001) Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J Comput Neurosci* 11:63–85.
- Deco G, Scarano L, Soto-Faraco S (2007) Weber's law in decision making: integrating behavioral data in humans with a neurophysiological model. *J Neurosci* 27:11192–11200.
- Deco G, Rolls ET, Romo R (2009) Stochastic dynamics as a principle of brain function. *Prog Neurobiol* 88:1–16.
- Dempster A, Laird N, Rubin D (1977) Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Ser B* 39:1–38.
- Harvey CD, Coen P, Tank DW (2012) Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484:62–68.
- Hernández A, Salinas E, García R, Romo R (1997) Discrimination in the sense of flutter: new psychophysical measurements in monkeys. *J Neurosci* 17:6391–6400.
- Hernández A, Zainos A, Romo R (2000) Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc Natl Acad Sci U S A* 97:6191–6196.
- Hernández A, Zainos A, Romo R (2002) Temporal evolution of a decision-making process in medial premotor cortex. *Neuron* 33:959–972.
- Hernández A, Nácher V, Luna R, Alvarez M, Zainos A, Cordero S, Camarillo L, Vázquez Y, Lemus L, Romo R (2008) Procedure for recording the simultaneous activity of single neurons distributed across cortical areas during sensory discrimination. *Proc Natl Acad Sci U S A* 105:16785–16790.
- Hernández A, Nácher V, Luna R, Zainos A, Lemus L, Alvarez M, Vázquez Y, Camarillo L, Romo R (2010) Decoding a perceptual decision process across cortex. *Neuron* 66:300–314.
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436.
- Iwamura Y (1998) Hierarchical somatosensory processing. *Curr Opin Neurobiol* 8:522–528.

- Jones LM, Fontanini A, Sadacca BF, Miller P, Katz DB (2007) Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A* 104:18772–18777.
- Katahira K, Nishikawa J, Okanoya K, Okada M (2010) Extracting state transition dynamics from multiple spike trains using hidden Markov models with correlated Poisson distribution. *Neural Comput* 22:2369–2389.
- Kemere C, Santhanam G, Yu BM, Afshar A, Ryu SI, Meng TH, Shenoy KV (2008) Detecting neural state transitions using hidden Markov models for motor cortical prostheses. *J Neurophysiol* 100:2441–2452.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Lemus L, Hernández A, Luna R, Zainos A, Náchter V, Romo R (2007) Neural correlates of a postponed decision report. *Proc Natl Acad Sci U S A* 104:17174–17179.
- Liu J, Morel A, Wannier T, Rouiller EM (2002) Origins of callosal projections to the supplementary motor area (SMA): a direct comparison between pre-SMA and SMA-proper in macaque monkeys. *J Comp Neurol* 443:71–85.
- Luna R, Hernández A, Brody CD, Romo R (2005) Neural codes for perceptual discrimination in primary somatosensory cortex. *Nat Neurosci* 8:1210–1219.
- Luppino G, Matelli M, Camarda R, Rizzolatti G (1993) Cortico-cortical connections of area F3 (SMA-proper) and area F6 (Pre-SMA) in the macaque monkey. *J Comp Neurol* 338:114–140.
- Martí D, Deco G, Mattia M, Gigante G, Del Giudice P (2008) A fluctuation driven mechanism for slow decision processes in reverberant networks. *PLoS One* 3:e2534.
- Martínez-García M, Rolls ET, Deco G, Romo R (2011) Neural and computational mechanisms of postponed decisions. *Proc Natl Acad Sci U S A* 108:11626–11631.
- Matsuzaka Y, Aizawa H, Tanji J (1992) A motor area rostral to the supplementary motor area (presupplementary motor area) in the monkey: neuronal activity during a learned motor task. *J Neurophysiol* 68:653–662.
- Mazor O, Laurent G (2005) Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48:661–673.
- Miller P, Katz DB (2010) Stochastic transitions between neural states in taste processing and decision making. *J Neurosci* 30:2559–2570.
- Mountcastle VB, Steinmetz MA, Romo R (1990) Frequency discrimination in the sense of flutter: psychophysical measurements correlated with post-central events in behaving monkeys. *J Neurosci* 10:3032–3044.
- Ohbayashi M, Ohki K, Miyashita Y (2003) Conversion of working memory to motor sequence in the monkey premotor cortex. *Science* 301:233–236.
- Pons TP, Garraghty PE, Friedman DP, Mishkin M (1987) Physiological evidence for serial processing in somatosensory cortex. *Science* 237:417–420.
- Rabiner L (1989) A tutorial on hidden markov models and selected applications in speech recognition. *Proc IEEE* 77:257–286.
- Renoult L, Roux S, Riehle A (2006) Time is a rubberband: neuronal activity in monkey motor cortex in relation to time estimation. *Eur J Neurosci* 23:3098–3108.
- Riehle A, Grün S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950–1953.
- Rizzolatti G, Luppino G (2001) The cortical motor system. *Neuron* 31:889–901.
- Rolls ET, Grabenhorst F, Deco G (2010) Choice, difficulty, and confidence in the brain. *Neuroimage* 53:694–706.
- Romo R, Salinas E (2003) Flutter discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4:203–218.
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473.
- Romo R, Hernández A, Zainos A, Lemus L, Brody CD (2002) Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat Neurosci* 5:1217–1225.
- Romo R, Hernández A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron* 41:165–173.
- Salinas E, Romo R (1998) Conversion of sensory signals into motor commands in primary motor cortex. *J Neurosci* 18:499–511.
- Salinas E, Hernández A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J Neurosci* 20:5503–5515.
- Schwarz G (1978) Estimating the dimension of a model. *Ann Stat* 6:461–464.
- Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996) Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci* 16:752–768.
- Siegel S, Castellan NJ (1988) Nonparametric statistics for behavioral sciences. New York: McGraw-Hill.
- Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A (1995) Dynamics of neuronal interactions in the monkey cortex in relation to behavioral events. *Nature* 373:515–518.
- Wang XJ (2002) Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36:955–968.
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26:1314–1328.
- Wong KF, Huk AC, Shadlen MN, Wang XJ (2007) Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front Comput Neurosci* 1:6.
- Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M (2009) Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J Neurophysiol* 102:614–635.