Article

Neuron

Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty

Highlights

- A template-matching algorithm detects neural correlates of false alarm events
- The subject's response criterion modulates over the course of a trial
- The response criterion is represented by the dynamics of a neural population
- A trained recurrent network unveils a mechanism for flexible response criterion

Authors

Federico Carnevale, Victor de Lafuente, Ranulfo Romo, Omri Barak, Néstor Parga

Correspondence

rromo@ifc.unam.mx

In Brief

Carnevale et al. explore how monkeys exploit previous knowledge to cope with temporal uncertainty in a perceptual detection task. The study demonstrates a neural mechanism by which prior information is intrinsically encoded in the dynamics of a neural population.

Carnevale et al., 2015, Neuron *86*, 1067–1077
May 20, 2015
http://dx.doi.org/10.1016/j.neuron.2015.04.014

CellPress

Neuron
Article

Dynamic Control of Response Criterion in Premotor Cortex during Perceptual Detection under Temporal Uncertainty

Federico Carnevale,[1] Victor de Lafuente,[2] Ranulfo Romo,[3,4,*] Omri Barak,[5,6] and Néstor Parga[1,6]
[1]Departmento de Física Teórica, Universidad Autónoma de Madrid, Cantoblanco 28049, Madrid, Spain
[2]Instituto de Neurobiología, Universidad Nacional Autónoma de México, Querétaro 76230, México
[3]El Colegio Nacional, 06020 México DF, México
[4]Instituto de Fisiología Celular-Neurociencias, Universidad Nacional Autónoma de México, 04510 México DF, México
[5]Faculty of Medicine, Technion–Israel Institute of Technology, Haifa 32000, Israel
[6]Co-senior author
*Correspondence: rromo@ifc.unam.mx
http://dx.doi.org/10.1016/j.neuron.2015.04.014

SUMMARY

Under uncertainty, the brain uses previous knowledge to transform sensory inputs into the percepts on which decisions are based. When the uncertainty lies in the timing of sensory evidence, however, the mechanism underlying the use of previously acquired temporal information remains unknown. We study this issue in monkeys performing a detection task with variable stimulation times. We use the neural correlates of false alarms to infer the subject's response criterion and find that it modulates over the course of a trial. Analysis of premotor cortex activity shows that this modulation is represented by the dynamics of population responses. A trained recurrent network model reproduces the experimental findings and demonstrates a neural mechanism to benefit from temporal expectations in perceptual detection. Previous knowledge about the probability of stimulation over time can be intrinsically encoded in the neural population dynamics, allowing a flexible control of the response criterion over time.

INTRODUCTION

One of the main challenges of cognitive neuroscience is to understand how external sensory stimuli and internal brain states interact to give rise to perception (Romo and de Lafuente, 2013). Internal states are believed to reflect acquired experience that can be used for making the best sense of sensory inputs (Gilbert and Sigman, 2007). During perceptual decisions, for example, the brain uses previous knowledge to transform noisy sensory evidence into percepts on which decisions are based (Forstmann et al., 2010; Hanks et al., 2011; Rao et al., 2012; Ratcliff and McKoon, 2008; Simen et al., 2009; Summerfield and Koechlin, 2008). In this study, we explore the dynamic nature of these internal states by asking how previous information about the timing of sensory evidence is incorporated in the decision-making process. We combine computational modeling with neurophysiological and behavioral data recorded while monkeys performed a somatosensory detection task (de Lafuente and Romo, 2005, 2006).

Subjects performing a decision-making task can benefit from the use of temporal expectations (Coull and Nobre, 2008) at multiple stages of the sensorimotor transformation (Nobre et al., 2007): (1) perception can be enhanced by increasing sensory accuracy at the relevant times (Correa et al., 2005; Ghose and Bearl, 2010; Ghose and Maunsell, 2002; Jaramillo and Zador, 2011; Rohenkohl et al., 2012); (2) the response criterion—the subject's internal rule to decide whether or not to report a stimulus—can be modulated to incorporate prior information without changes in the sensory representation (Katzner et al., 2012); and (3) motor readiness can be heightened, increasing the response speed in reaction-time tasks (Nobre, 2001; Scheibe et al., 2009). These studies have found neurophysiological evidence for the use of temporal information in the sensory and motor stages. However, little is known about the neural mechanisms that underlie the use of timing at intermediate stages of the sensorimotor transformation.

We address this intermediate step by analyzing recordings of premotor cortex neurons from monkeys performing a detection task with variable stimulus onset times (Figure 1A; see de Lafuente and Romo 2005, 2006). The task's temporal structure dictated that the stimulus only arrived within a 2 s temporal window but not before or after (Figure 1B). We asked whether monkeys can infer and take advantage of this temporal structure to increase performance. One possible way to incorporate this knowledge is to modulate the response criterion (the amount of sensory evidence required to produce a stimulus-present response) over the time course of the trial (Figure 1C). An efficient modulation of the criterion is to raise it outside the possible stimulation window to avoid false positive outcomes, and lower it within the window to allow correct detections. The exact shape of the response criterion within the possible stimulation window depends on the animal's inference about the underlying distribution of stimulus onset times (the subjective hazard function; Janssen and Shadlen, 2005; Luce, 1986; see Discussion).

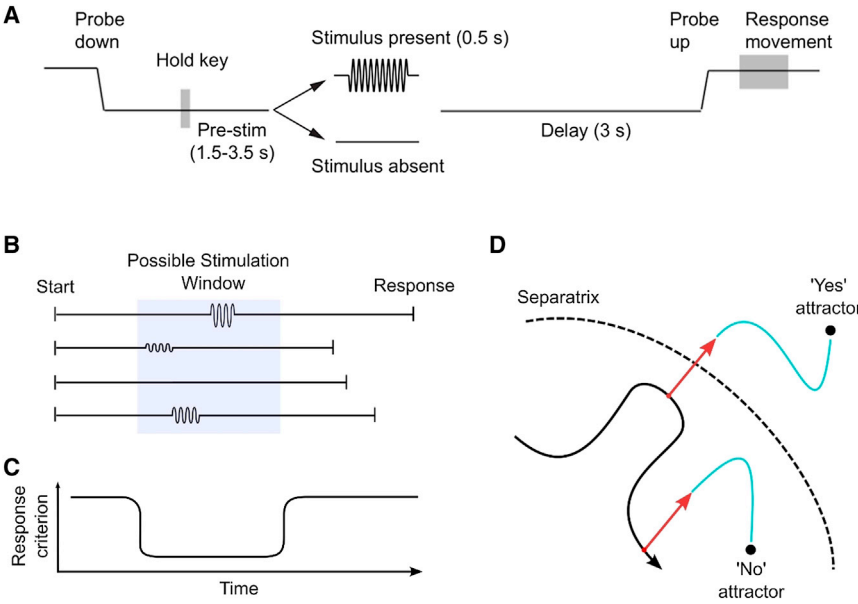

Figure 1. Detection Task and Dynamical Response Criterion

(A) Behavioral task represented by the vertical position of the mechanical probe during a trial. The stimulator indented the skin of one fingertip of the restrained hand ("probe down"), and the monkey reacted by placing its free hand on an immovable key ("hold key"). After a variable prestimulus period (1.5–3.5 s), a vibratory 0.5 s stimulus was presented on half of the trials. At the end of a fixed delay period (3 s), the stimulator moved up ("probe up"), instructing the monkey to make a response movement to one of two push buttons. The pressed button indicated whether or not the monkey felt the stimulus.

(B) The variability in stimulus onset times and the fixed delay period defined a 2 s temporal window of possible stimulation. No stimulus was delivered before 1.5 s or after 3.5 s from the "hold key" event. The window of possible stimulation was not explicitly cued to the animal.

(C) A possible mechanism to efficiently solve the task requires modulating the response criterion (the strength of sensory evidence required to produce a stimulus-present response) over time. Outside the possible stimulation window, the response criterion is high to avoid false positives. Within the window, the response criterion decreases to allow correct detections.

(D) The mechanism described in (C) could be dynamically implemented by a separatrix in the neural space, dividing the basins of attraction of two attractors. The black trace is a trajectory of a correct rejection trial. The blue traces represent a hit (ending in the "yes" attractor) or a miss trial (ending in the "no" attractor). The distance from the current neural state to the separatrix at each point in time represents the response criterion.

How could a population of neurons implement such a mechanism? If we consider the abstract high dimensional space of neural activity, the threshold to commit to a decision can be pictured as a boundary, that once crossed, triggers perceptual detection (Figure 1D; the dynamical systems term for the boundary is "separatrix"). The response criterion is given by the distance from the current state of the network to that boundary. Temporal expectations can then be manifested via the trajectory of neural dynamics while the monkey is waiting for a stimulus—drawing closer to the boundary when the stimulus is expected and vice versa (Figure 1D).

In this work, we present experimental and modeling evidence in favor of this dynamical mechanism. We start by using the timing of false alarms to infer the dynamics of the response criterion. To obtain these times, we develop a method to detect neural correlates of false alarms and find that indeed their probability increases during the period of possible stimulation. Following the intuition outlined above, we analyze the dynamics of correct rejection trials—as these encapsulate the "waiting for a stimulus" condition. We show that in these trials the neural trajectory is modulated precisely during the period of possible stimulation. Finally, we derive a model by training a recurrent network to perform an analogous detection task. We find that the model is able to infer the task's temporal structure, and using it, we unveil the explicit dynamical implementation of the proposed neural mechanism.

RESULTS

Monkeys were trained to detect a weak mechanical vibration of variable amplitude applied to one of their fingertips. Reward was provided for correctly reporting the presence (hit trials) or absence (correct rejection trials, CR) of the stimulus. In contrast, no reward was delivered during incorrect trials, which arose either from missing a stimulus (miss trials) or reporting a false positive (false alarm trials, FA). The stimulus onset time varied from trial to trial between 1.5 s and 3.5 s after the "hold key" event (Figure 1A). Following stimulation (or absence thereof), monkeys had to wait for a 3 s delay period until a cue indicated to report their decision. Because of this temporal structure, we expect subjects to modulate their response criterion to benefit from the fact that no stimulus arrived before 1.5 s or after 3.5 s (Figures 1B and 1C).

A modulation in the response criterion has predictable consequences on behavior. A higher response criterion leads to an increase in the frequency of stimulus-absent responses, while a lower response criterion implies an increase in the frequency of stimulus-present ones. Therefore, evidence of change in response criterion over time could be obtained by estimating the frequency of stimulus-present responses as a function of time. However, in a delayed-response task there is no behavioral information about the exact time at which the subject reached a decision and, therefore, it is not possible to estimate a time-varying response criterion from behavioral data.

Nonetheless, in any two-alternative forced choice task, a decision represents a commitment to one of the two possible alternatives. Thus, we hypothesized that information about the timing of the subject's decision could be found in the neural activity. Premotor cortex (PMc) activity was previously shown to correlate more with the subject's perceptual decision than with the physical properties of the stimulus (de Lafuente and Romo, 2005). In

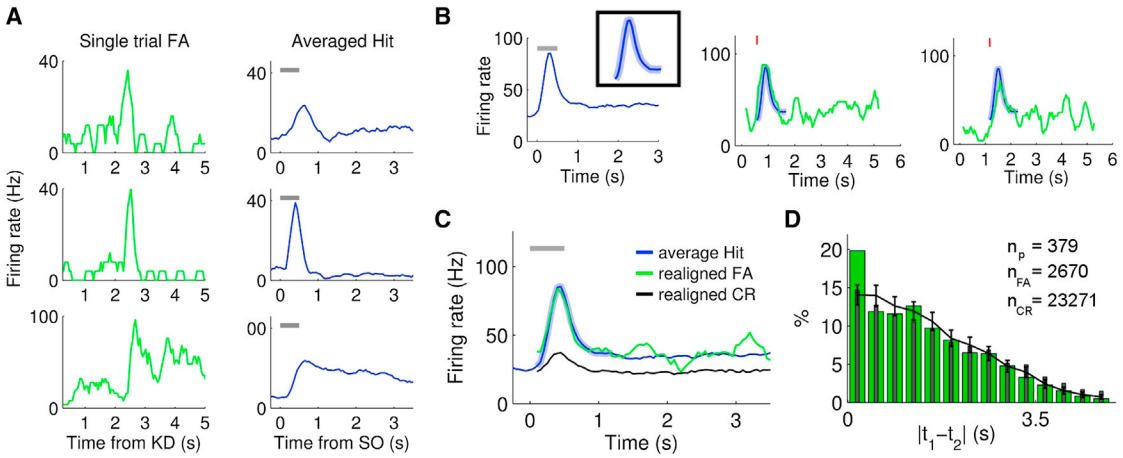

Figure 2. Detection of False Alarm Events by Template Matching

(A) Firing rate of three simultaneously recorded neurons during a single FA trial (green traces) and average response of the same neurons during hit trials (blue traces). The shaded bar indicates the stimulation period.

(B) We used a 1 s segment of the averaged activity during hit trials (left, blue trace) as a template (inset) to detect FA events in single FA trials. FA events were identified in single FA trials (middle and right, green traces) on the basis of the mean squared error between the single FA trial firing rate and the template. Red lines indicate the start of the template.

(C) The average activity over FA trials realigned according to the times of detected events (green trace) matches the average over hit trials (blue trace) even outside of the period used as template. In contrast, the same method applied to CR trials produces a much weaker match.

(D) Histogram of differences in the detected FA times from pairs of simultaneously recorded neurons. A significant fraction of FA trials was detected at the same time compared to CR trials (black bars) and chance level (black line, $p < 0.001$). Chance level was obtained by shuffling the trials to disrupt the correspondence between detected FA events in simultaneously recorded neurons. Error bars indicate 95% confidence intervals. n_p is the number of neural pairs, and n_{FA} and n_{CR} are the number of FA and CR trials, respectively.

fact, it was shown that the subject's decision can be unambiguously decoded from a population of neurons in this cortical area (Carnevale et al., 2013). Moreover, premotor cortex activity was previously shown to reflect an internal component of the decision process (Carnevale et al., 2012). Therefore, we set to find information about the subject's response criterion from the firing activity of PMc neurons. We analyzed an experimental data set from two earlier studies (de Lafuente and Romo, 2005, 2006) which included single-neurons and small sets of simultaneous neuronal recordings (up to 6 cells with a median of 2), summing to a total of 384 extracellularly recorded neurons (see Experimental Procedures).

False Alarms as a Window onto the Response Criterion's Dynamic

As stated above, if monkeys modulate their response criterion during the time course of a trial, this should be reflected in the probability of producing a false alarm over time. Here we set to find this information from single trial neural activity. We assume that the decision process carried out in every trial led to a commitment to one of the two possible alternatives. Therefore, we expect some neurons to present irreversible and stereotypical activity profiles reflecting this commitment. Indeed, we noticed that premotor cortex neurons during single FA trials show temporally localized fluctuations in their firing rate (Figure 2A, left). The profile of these events resembles the neuron's responses evoked by the vibrotactile sensory stimulation (Figure 2A, right). Furthermore, they occur at the same time in simultaneously recorded neurons, suggesting that they correspond to the same perceptual event. Taken together, this suggests that these fluctuations, which we called FA events, are neural correlates of false alarms.

We devised a method to detect the time of production of FA events from single FA trials (Figure 2B; see Experimental Procedures). For each neuron, we used the average firing rate in hit trials to define a 1 s template representing the neuron's specific response to the external stimulation (Figure 2B, inset). Applying this template to each individual FA trial, we searched for similar firing profiles, providing putative FA event times (Figure 2B). By realigning the FA trials according to the detected times, we obtained an average response resembling that of the hit trials (Figure 2C, blue and green traces). Importantly, this is true not only during the 1 s period used as a template, but also during the remaining 2 s of the delay period, consistent with the idea that what we detected is a stereotypical activity profile equivalent to the one evoked by the external stimulus. Moreover, this applies to neurons with very diverse firing temporal profiles (see Figure S1 available online).

Due to the noisy nature of single trial data, our template-matching algorithm produces a large amount of false detections. In particular, simply observing the average of the realigned FA trials (Figure 2C) suffers from a circular logic—we may pick out events from noise and by definition get a similar waveform after averaging the realigned trials. To validate the significance of the detected events, we used the activity of simultaneously recorded neurons. If the FA events are neural correlates of FAs, they should occur at the same time on different neurons. For each trial, we compared the FA event times obtained independently from

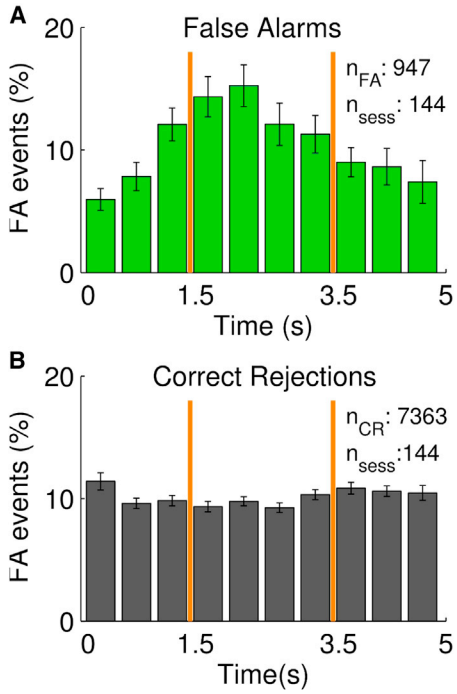

Figure 3. Probability of False Alarm over Time

(A) Mean relative frequency of detected FA events over time during the time course of the trial. The probability of a FA event increases during the period of possible stimulation (within orange lines). Relative frequency was calculated as the portion of FA events detected in each time bin relative to the number of FA trials in which a FA event was detected at any time bin. The mean histogram was obtained by averaging across n_{FA} = 947 FA trials distributed in n_{sess} = 144 sessions. Error bars represent SEM.

(B) Same as (A) for CR trials.

two different neurons. The histogram in Figure 2D (green bars) shows the frequency of FA event's time differences over the entire set of FA trials. A significant fraction of FA events were detected at the same time compared to chance level (compare first green time bin to black line). We applied the same template-matching algorithm to CR trials to further control for circular logic, as they presumably have the same noisy nature as FA trials but do not present FA events. CR trials revealed both a weaker agreement with the average hit firing rate (Figure 2C, black trace), as well as no significant number of simultaneous events (Figure 2D, black bars). Taken together, these results suggest that at least a subset of FA trials can be explained by an event that is localized in time and that triggers an irreversible and stereotypical neural activity pattern, equivalent to the one evoked by external stimulation.

Under the hypothesis that these FA events generate a false percept, the estimation of the times at which these events are produced allows to obtain the probability of a false alarm over the time course of the trial. By using those trials in which FA events were detected in two or more neurons simultaneously (first bin of Figure 2D), we computed the frequency of detected events across time. The resulting probability is not uniform (Figure 3A, green bars) but reaches a maximum during the period of possible stimulation (Figure 3A, orange lines). In contrast, the same quantity obtained for CR trials, used as a control, revealed no modulation during this period (Figure 3B, black bars). The increase in the probability of FA during the possible stimulation window is consistent with a decrease in the subject's response criterion when the stimulus is more likely to arrive. Figure 3A suggests that monkeys are able to infer the task's temporal structure and make use of this knowledge to modulate their response criterion according to the stimulation probability. Notice that the focus of our analysis is the change in the animal's response criterion between outside and inside the possible stimulation window, information that cannot be obtained from behavioral data in our task (see Figure S2).

Premotor Cortex Dynamics Suggests a Neural Mechanism for Modulating the Response Criterion

So far we showed that monkeys used previously acquired temporal information when performing the vibrotactile detection task. In the previous section we used FA trials to show that the subject's response criterion was modulated over the time course of the trial. In this section we aim to find signatures of the dynamics of this modulation in the neural activity. We analyze CR trials to show that the activity in this condition reflects the subject's expectations about the probability of stimulation over time.

The activity of neurons in PMc was previously shown to reflect the subject's perceptual judgment about the presence or absence of the stimulus during the vibrotactile detection task (de Lafuente and Romo, 2005). Several pieces of evidence were examined in that study. First, PMc neurons responded in an all-or-none manner, only weakly modulated by the amplitude of the stimulus. Second, when presented with a fixed near-threshold stimulus, PMc activity strongly correlated with the subject's choice. Third, reversing the direction of the arm movements in control experiments did not change the activity of PMc. Fourth, when PMc was electrically microstimulated, the probability of stimulus-present responses was higher than when only the mechanical stimuli was presented. In fact, another study showed that the subject's decision could be unambiguously predicted from the activity of populations of PMc neurons (Carnevale et al., 2013).

If PMc activity represents the subject's perceptual judgments, we expect that a modulation in the subject's response criterion will be reflected in the firing rate of PMc neurons. However, the response of PMc is quite heterogeneous across neurons (de Lafuente and Romo, 2006). When presented with a suprathreshold stimulus, some neurons increased their firing rate while others tended to decrease it. Moreover, the temporal profile of PMc neural responses was also diverse. Some neurons responded only during stimulus presentation while others showed persistent activity or even ramping profiles during the delay period of the task (see also Figure 3A). In face of this heterogeneity, it is not trivial to predict how a modulated response criterion would be reflected in each single neuron's activity.

State-space analysis was shown to be a useful tool to study the neural dynamics at the population level (Cunningham and Yu, 2014; Mante et al., 2013; Shenoy et al., 2013; Stokes et al., 2013). In this framework, the activity of a population of N neurons at each point in time is represented as an N-dimensional point in the space spanned by each neuron's activity. The population

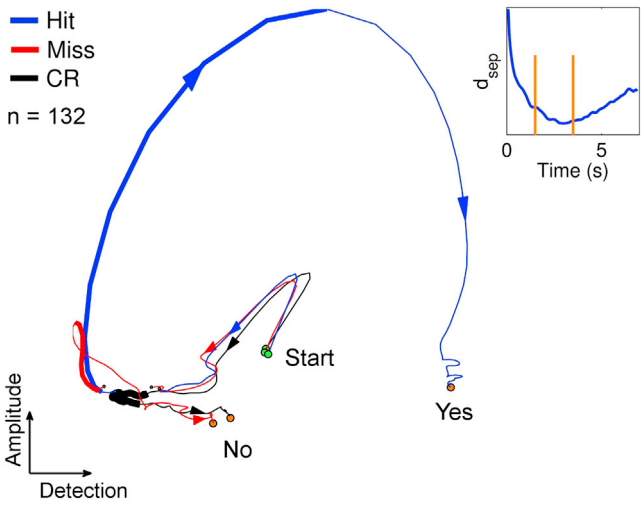

Figure 4. Two-Dimensional Projection of the Population Dynamics

Average neural trajectories during hit (blue), miss (red), and CR (black) trials projected onto two task-related axes (stimulus amplitude and stimulus detection). The trajectories are plotted from the beginning of the trial (green circles) to end of the delay period (orange circles). Stimulus-present conditions are plotted until 1.5 s and realigned at the stimulus onset time. Thick blue and red traces indicate the period of stimulation. The thick black line denotes the possible stimulation window (1.5 s to 3.5 s). Units are arbitrary. The inset is the N-dimensional Euclidean distance between the CR neural trajectory and the neural state at the stimulus offset time during the miss condition (end of the thick red trace), as an estimate of the distance to the separatrix over time. See Figure S4 for the same analysis performed for each subject separately.

activity across time defines a trajectory within this space. The set of neural trajectories often occupies a low-dimensional subspace within the space of possible activities, and various methods can be used to visualize it. We generated 132-dimensional trajectories by combining neural data mostly recorded separately (the number of neurons was limited by the need to match conditions between different recording sessions, see Experimental Procedures). Then, we projected these trajectories onto two task related axes—tuning to stimulus amplitude and tuning to a detection event. The former was defined by regressing each neuron's trial-to-trial response to the stimulus amplitude. The latter represented the direction in which the network evolves during a detection. We did not use the regression to define the "detection" axis to avoid colineality problems, as stronger stimuli are more likely to be detected. In general, however, the direction in which the stimulus drives the population activity, and the direction in which the population activity evolves during the formation of the decision are different. To capture this, we defined the "detection" axis as the vector connecting the network state just before the application to the network state at the end of the delay period, during hit trials. With this definition, the angle between detection and amplitude axes is 87°.

Figure 4 shows the average neural trajectories during hit, miss, and CR trials projected onto these axes. We omit here the FA condition because, as we showed before, the neural activity during each individual FA trial is, to a large extent, equivalent to the one during a hit trial and thus the averaged FA trajectory is actually composed by many individual FA events misaligned and distributed over time (see, however, the realigned FA trajectory in Figure S3). For the hit condition, we discarded weak amplitude trials to avoid the inclusion of possible FA events, while for the miss condition, we dismissed strong amplitude trials to avoid errors due to lapsus. We plotted each trajectory from the beginning of the trial to the end of the delay period (Figure 4, green and orange circles, respectively). Due to the variable stimulus onset times, stimulus present conditions (hit and miss) are aligned first to the "hold key" time and then to the "stimulus onset" time.

As expected, the neural trajectories for the three conditions overlap at the beginning of the trials. In hit and miss conditions (blue and red traces) the application of the stimulus drives the network in the stimulus direction. The stimulus on hit trials moves the network further than in the miss condition. Afterward, during the delay period, the network evolves into two different trajectories. The endpoints of these trajectories (orange circles) represent the final states of the network, presumably corresponding to the stimulus-present and stimulus-absent choices (compare to Figure 1D).

We used the same axes to project the population activity during CR trials. In this condition the monkey waited for a stimulus that never came, and then correctly reported its absence. Thus, the monkey's knowledge about the task's temporal structure and the resulting expectations about the stimulus arrival should be reflected in this condition. The neural trajectory during CR trials is shown in Figure 4 (black trace). Interestingly, it presents a modulation precisely during the period of possible stimulation (thick black line). After 3.5 s from the beginning of the trial, the network state evolves to the same final state as the trajectory in the miss condition (stimulus-absent state). Importantly, the projection axes were not selected 'ad hoc' to see this modulation. In fact, a similar modulation is observed when the neural trajectories are projected onto the two principal components of the data (Figure S5). In general, task-related axes give a more intuitive picture of the neural dynamics during the task, but the particular choice of axes does not critically affect our results.

Figure 4 is consistent with our proposed dynamical mechanism (Figure 1D). The modulation observed in CR trials during the possible stimulation window can be a signature of the network approaching a separatrix beyond which the dynamics leads to a stimulus-present response. During the possible stimulation window, this distance should decrease, lowering the response criterion when the stimulus is more likely to come and then it should increase again. While the location of such a separatrix cannot be obtained from the recorded neural activity, the network state just after the offset of the stimulus during a miss trial should be below and close to it. We can then estimate the distance to the separatrix as the Euclidean distance in the high-dimensional space between this state (the average neural activity during miss trials at stimulus offset) and the neural trajectory during CR trials. This measure indeed decreases during the period of possible stimulation (Figure 4, inset).

Summarizing, the state of the network while the monkey is waiting for a stimulus is not stationary. In contrast, Figure 4 suggests that the neural trajectory could intrinsically encode the temporal information about the probability of stimulation over time. While the subject is waiting for the stimulus, the neural trajectory is determined solely by the internal neural dynamics.

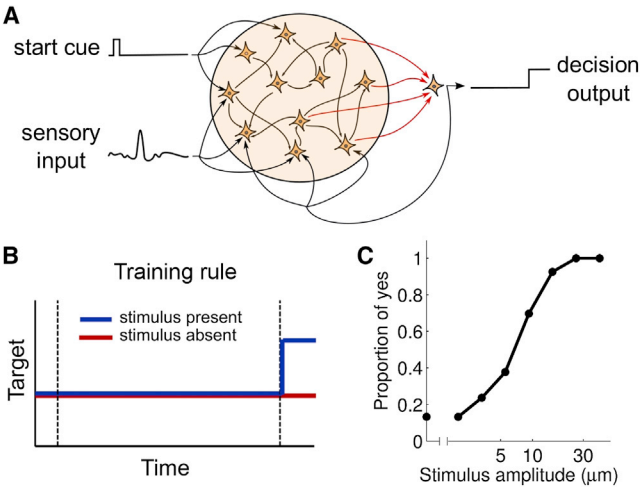

Figure 5. A Recurrent Network Model Learns to Solve the Task

(A) Recurrent network model of rate units provided with a start cue input, a noisy sensory channel, and a decision output. The start cue indicates the beginning of a new trial. The stimulus is modeled as a pulse corrupted by noise. The decision is extracted from a linear combination of rates after the delay period. We trained the initially random network by changing the output connections. Because of the feedback loop, this effectively alters the recurrent dynamics of the network.

(B) Target signal of the FORCE algorithm. The information provided during training was restricted to the behavioral outcome on each trial. Thus, no information about the probability of stimulation over time was given during training.

(C) "Psychometric" function of the trained model obtained as the frequency of stimulus-present responses as a function of stimulus amplitude.

Therefore, the temporal expectations that the subject built during training, might be internally stored by the dynamics of the neural population.

A Recurrent Network Unveils the Dynamical Implementation of Response Criterion Modulation

What dynamical mechanism supports the use of prior temporal information during perceptual detection? We used a recurrent neural network model to answer this question. Starting with a random recurrent network, we trained it to perform a simplified version of the experimental task. After verifying that the model is able to solve the task, we analyzed the solution achieved. We were especially interested in whether the developed solution makes use of temporal information. We asked if the network is able to benefit from temporal information acquired during the training phase and, if so, what are the dynamical mechanisms by which this information is integrated with the sensory evidence to detect the presence of a stimulus.

Our model is a recurrent neural network of rate units, provided with two inputs and one output (Figure 5A; see Experimental Procedures). The first input is used to signal the start of a new trial, while the second one represents the sensory channel via which the stimulus is applied. The stimulus is modeled as a pulse proportional to the vibration's amplitude, embedded in a noisy background. In each trial the decision about the presence or absence of the stimulus is indicated by the value of the output during a readout interval. Trials begin with a start cue applied through one of the input channels. After a variable pre-stimulus time, on half of the trials, a stimulus is presented through the sensory channel. Then, after a fixed delay period, the decision is extracted from the network's output through a linear read out. Trials are simulated in a continuous manner, without any reset between them—the start cue input provides the time reference within a trial.

After randomly initializing all synaptic weights, we trained the recurrent network to solve the task. During the training phase we used the FORCE algorithm to change the output connections (Sussillo and Abbott, 2009). Although these are the only plastic weights, because of the feedback loop this change results in a rank-one perturbation to the effective recurrent weight matrix, therefore changing the dynamics of the network (Sussillo and Abbott, 2012; see Experimental Procedures). Training is controlled by a teaching signal representing the desired output in each trial. Since we want to find out whether the strategy developed during training makes use of the timing of the sensory evidence, we provided no explicit information about the probability of stimulation over time. The teaching signal was restricted to the behavioral outcome on each trial (Figure 5B) —an analogous information to the one that the monkeys receive in the experimental setup.

The resulting network learns to solve the task. Performance is controlled by the amount of noise in the sensory channel, so, once trained, we calibrated the noise amplitude to approximately reproduce the averaged experimental psychometric function (Figure 5B; compare to de Lafuente and Romo 2005, Figure 1C). Then, we asked if the network is able to infer the task's temporal structure and use this information to perform the task. Because this is a model, we are able to test the network's behavior on a large number of stimuli without inducing any learning. Thus, we systematically probed the trained model with variable amplitude stimuli applied at different times from the beginning of the trial. We followed a bisection protocol to find the lowest stimulus amplitude which drives the network to a stimulus-present response (see Experimental Procedures). This quantity, which represents the model's response criterion, is not fixed but decreases during the period of possible stimulation used during training (Figure 6A). To verify the dependence of this measure on the statistics upon which the network was trained, we repeated this procedure for different possible stimulation windows and observed that the response criterion modulated accordingly (Figure 6A, inset).

The modulation of the response criterion was also revealed when we applied the same template-matching algorithm that we used on the experimental data to the FA trials produced by the simulation. False alarms in the model occur due to the noisy nature of the sensory input: in a stimulus-absent trial, a random fluctuation in the sensory signal is detected by the network as an external stimulus and induces a positive response. When we obtain the probability of false alarms as a function of time, we find that it increases during the possible stimulation window as it is expected from a decrease in the response criterion at that time (compare Figures 6B and 3A).

Finally, we set to understand what dynamical mechanism, developed during training, supports the modulation of the

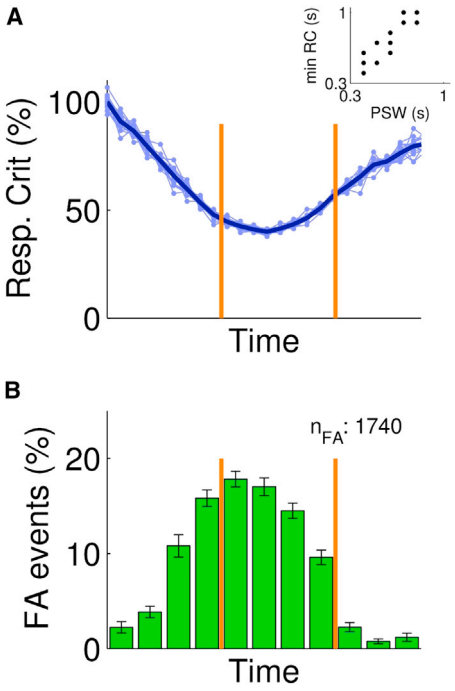

Figure 6. The Network Infers the Window of Possible Stimulation

(A) The response criterion, defined as the lowest stimulus amplitude that drives the network to a stimulus-present response, decreases during the period of possible stimulation (within orange lines). The response criterion was obtained by systematically probing the network with a bisection protocol at each time to find "borderline" stimulus amplitudes. Thin lines represent single realizations of this protocol. Thick line is the mean of n = 10 realizations. The response criterion was normalized with its maximum value during the trial. Inset shows the results of training networks with different possible stimulation windows. PSW is the center of the possible stimulation window used during training; min RC is the time in which the response criterion reaches its minimum value.

(B) Mean relative frequency of detected FA events over time in the model obtained by the same template-matching algorithm used for the experimental data. The probability of producing a FA increases during the period of possible stimulation (within orange lines). Relative frequency is defined as in Figure 3. The mean histogram was obtained by averaging across sessions. Error bars represent SEM. n_{fa}, number of FA trials.

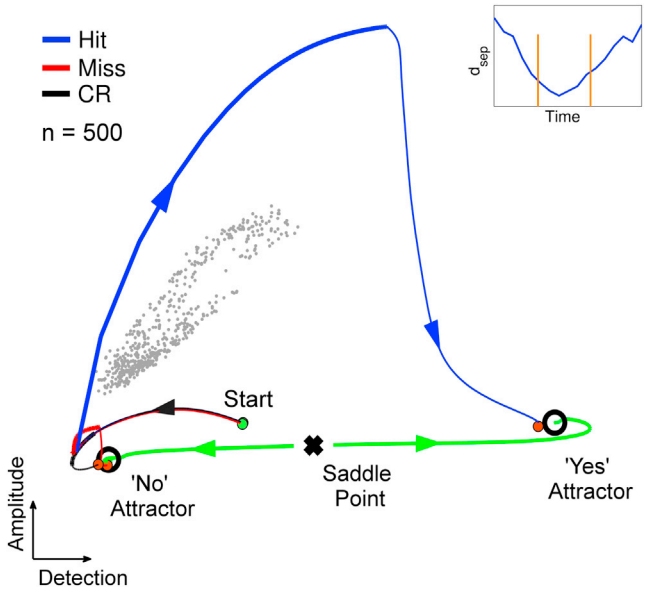

Figure 7. Neural Trajectories of the Recurrent Neural Network Model

Neural trajectories during a hit (blue), a miss (red), and a CR (black) trial projected in the same axes as in Figure 4. The three trajectories overlap during the beginning of the trial. The stimulus is applied (in the hit and miss conditions) at the middle of the possible stimulation window (thick black line in CR). The hit trajectory evolves to the "yes" attractor, while the miss and CR trajectories end in the "no" attractor. The gray dots are points close to the separatrix, estimated as the states achieved during "borderline" stimuli. Inset shows the distance between the network state during CR trials and the separatrix. Note that distance is measured in the high-dimensional space and therefore cannot be inferred from the 2D plot. The fixed-points analysis of the trained network revealed a saddle point mediating the decision between the two stable fixed points. The green traces represent the trajectories starting near the saddle point following its unstable direction. For better visualization of this figure, the simulations were run without noise in the sensory inputs, but the effects do not change under noisy stimuli.

response criterion during the time course of the trial. To do so, we reverse-engineered the network by looking for slow and fixed points, and analyzing the linear dynamics around them (Sussillo and Barak, 2013). We found that the network's dynamics is governed by three fixed points, two stable and one unstable (Figure 7). The two attractors correspond to each of the possible decision outcomes ("yes" and "no" attractors). The third fixed point presents only one unstable eigenvector, and therefore is a saddle point. This saddle point defines a separatrix between the basins of attraction of the two stable fixed points (Figure 7, gray dots; see also Figure S6). At any point in time, strong enough sensory input can make the network cross the separatrix and travel to the stimulus-present attractor, resulting in a hit trial (Figure 7, blue trace). In contrast, a weak sensory input will fail to drive the network across the separatrix, resulting in a missed stimulus (Figure 7, red trace). The distance between the network state and the separatrix controls how strong the sensory input must be to produce a stimulus-present response (Figure 7, inset; compare with Figure 6A). Therefore, the network's response criterion can be modulated by controlling the state of the network relative to the separatrix.

During training, the FORCE algorithm changes the network's synaptic weights, sculpting the dynamics of the model according to the examples' stimulation times. This can happen either by a modification in the relaxation process that takes place after the "start trial" signal is removed and before the stimulus arrives, or by sculpting the shape of separatrix in the state space, according to the stimulation times presented during training. In any of these cases, the neural trajectory that the network develops after the training phase relative to the separatrix, controls the response criterion at each point in time and incorporates the acquired knowledge about the timing of sensory evidence.

DISCUSSION

Under temporal uncertainty, the detection of a sensory stimulus embedded in a noisy background can be improved by previous

knowledge about the probability of stimulus arrival over time. This improvement can arise from a dynamical modulation of the response criterion over time. We used analysis methods to extract the response criterion from neural data, demonstrating that it indeed is modulated according to the learned temporal structure of the task. Furthermore, we proposed a dynamical mechanism for the modulation of the response criterion and showed that PMc activity is consistent with it.

Previously, we analyzed the firing activity and correlated variability during the same experimental task to show that the decision-making process starts before the arrival of the sensory evidence (Carnevale et al., 2012). Our results suggested an internal signal that is combined with the external stimulus to drive the animal's decision. In this work, we show that this internal process reflects the temporal expectations of the subject. First, we demonstrate that the response criterion modulates during the course of the trial. Second, we find that this modulation is represented by the dynamics of a PMc neural population. We reinterpret the internal signal proposed before as the evolution of the network state preceding the application of the stimulus. In both views, the decision results from the combination of an internal phenomenon and the sensory evidence. Our current view, however, shows that this internal component of the decision process arises from the subject's prior information (temporal expectations) and that the incorporation of this prior information can be done by the recurrent dynamics of a neural network.

The focus of our analysis is the modulation of response criterion during the course of the trial. That is, we study the change in response criterion between periods in which the stimulus is never presented and periods in which the stimulus is likely to arrive. A related—but different—issue is the variability of the response criterion across different trials. Trial-to-trial variability in the network state before the arrival of the stimulus could contribute to explain behavioral variability. This should be reflected in the average neural trajectories during hit and miss trials—the network state should be closer to the separatrix in hit than in miss trials. While Figure 4 suggests that the trial-to-trial variability is weak, a direct analysis of variability in the network state requires population recordings in single trials. Thus, future studies are needed to clarify to what extent response criterion variability predicts behavior.

In the framework of signal detection theory (Green and Swets, 1966), the response criterion is expected to vary with the hazard rate—the probability of observing the stimulus in the next instant, knowing that it has not arrived up till now. For the uniform distribution of onset times used in this experiment, the hazard rate is increasing within the possible stimulation window. This means that the optimal response criterion, assuming complete knowledge of the distribution of stimulation times, should decrease within this window. Although our estimation of the probability of false alarm over time from the experimental data is noisy (Figure 3), its profile does not seem to be consistent with a decreasing response criterion within the possible stimulation window, as predicted from the hazard rate. One possible explanation for this deviation from optimality is that monkeys might be identifying periods in which the stimulus is more likely, without estimating the exact distribution of stimulation times. Indeed, in our task, changing the response criterion from high, outside, to low, inside the possible stimulation window, leads to an increase in performance (by reducing the number of false alarms when the stimulus is unlikely) that is greater than the one obtained by precisely modulating the response criterion according to the stimulus hazard rate within this window. Finally, the response criterion developed by the recurrent network model is also suboptimal (Figure 6). This might be due to a difficulty in learning the distribution of stimulus onset times, an imperfect track of time within trials, or a limitation in changing the dynamics very rapidly.

The neural mechanism that we propose for modulating the response criterion over time relies on the network's recurrent dynamics. Because the subject is presented with many trials of different stimulus onset times, we speculate that information about the timing of the sensory evidence might be incorporated in the decision process by plastic changes in the internal synaptic connections of the ''decision'' network (Janssen and Shadlen, 2005; Karmarkar and Buonomano, 2007; Leon and Shadlen, 2003). This builds on the framework of state-dependent or reservoir computing (Buonomano and Maass, 2009; Jaeger and Haas, 2004; Maass et al., 2007) in which computations arise from the interaction between external stimuli and the internally generated dynamics produced by the network's recurrent connectivity. This framework is particularly suitable to explain neural mechanisms that require time—like the effect of temporal expectations—because arbitrary functions of time can be intrinsically encoded in reproducible neural trajectories (Laje and Buonomano, 2013). Notably, in our implementation, the modulation of the response criterion according to the time-varying probability of stimulation does not need to be explicitly trained. It arises from the multiple presentations of many trials of different stimulus onset times together with a target signal indicating the presence or absence of the stimulus. While the training algorithm is far from being biologically realistic (but see Hoerzer et al., 2014), it is important to note that the information used during the online supervised learning was analogous to the one that the monkeys receive in the experimental setup.

The combination of previous knowledge about the stimulus probability with incoming sensory evidence was extensively studied in two-alternative forced choice discrimination tasks (Forstmann et al., 2010; Hanks et al., 2011; Rao et al., 2012; Ratcliff and McKoon, 2008; Simen et al., 2009; Summerfield and Koechlin, 2008). These studies suggest that stationary priors are incorporated into the decision process as a shift in the amount of evidence needed to reach a decision. In this work, we extended this question to the temporal domain. We used a detection task to ask how subjects can use prior information about the timing of stimulus arrival. Our results suggest a neural mechanism that supports the incorporation of a time-varying prior probability into the decision process.

How do our results extend to a discrimination task? We speculate that temporal expectations could dynamically shift the amount of required evidence for each choice according to their time-varying prior probability. In order to test this, one possibility would be to use a combined detection-discrimination task (Jaramillo and Zador, 2011) in which subjects must both detect and

discriminate stimuli to receive reward. In this task, the relative frequency of each alternative could be manipulated so that it changes within the time course a trial. We anticipate that individuals will infer the task's temporal structure and use their temporal expectations to dynamically modulate their bias for each alternative across time.

Note that the dynamic control of response criterion could be combined with the effect of other dynamic biases. In the reaction time version of the random dot task, for example, it has been shown that prior probabilities of the stimulus, even when stationary, could be incorporated as a dynamic bias signal, increasing the relative weight of priors over evidence as decision time increases (Hanks et al., 2011). Thus, several effects could interact to shape the time-to-time amount of evidence required to reach a decision.

We devised a method to extract the timing of false alarm events from the neural activity based on template-matching of activity patterns. While a similar approach was previously used in the context of memory-trace replay during sleep (Louie and Wilson, 2001), here we apply it to extract decision-related information from neural activity. Our method is useful to provide timing information about the subject's decision in situations when this is not immediately reported by its behavior (i.e., when short-term memory of the chosen alternative is required). Although we developed it to infer a time-varying response criterion, our technique has a broader applicability. It could, for example, provide valuable insight into the sources of false alarms. If the activity of neurons in sensory cortices are recorded simultaneously with PMc neurons, our method could be used to provide the relevant times for building a "false-alarm triggered averaged" of sensory activity. This could potentially disambiguate multiple possible origins of false alarms trials and contribute to the understanding of the role of noise on behavioral variability (Renart and Machens, 2014).

Subjects performing a decision-making task can benefit from the use of temporal expectations at multiple stages of the sensorimotor transformation. In this work, we showed that temporal information can be used to modulate the subject's response criterion across time. However, our experimental paradigm is not able to rule out other possibilities. For example, the sensory representation of stimuli could be changing over time (Correa et al., 2005; Ghose and Bearl, 2010; Ghose and Maunsell, 2002; Jaramillo and Zador, 2011; Rohenkohl et al., 2012). In the periods of higher expectations the signal to noise ratio of the sensory channel could be increased by mechanisms as synchronization (Steinmetz et al., 2000). During periods of lower expectation there could be gating mechanisms helping to avoid noise-induced false positives. Different experimental paradigms and further studies are needed to analyze the existence and coordination of different neural mechanisms for benefiting from temporal expectations.

EXPERIMENTAL PROCEDURES

Detection Task

Data for this analysis were obtained from two earlier studies (de Lafuente and Romo, 2005, 2006). Stimuli were delivered to the skin of the distal segment of one digit of the restrained hand, via a computer-controlled stimulator (BME Systems; 2 mm round tip). Initial probe indentation was 500 μm. Vibrotactile stimuli consisted of trains of 20 Hz mechanical sinusoids with amplitudes of 2.3–34.6 mm. These were interleaved with an equal number of trials where no mechanical vibrations were delivered to the skin (amplitude = 0). Animals pressed one of two buttons to indicate stimulus present (left button) or stimulus absent (right button). They were rewarded with a drop of liquid both types of correct responses, i.e, correct detections in stimulus-present trials and correct rejections in stimulus-absent trials. Animals were handled in accordance with standards of the National Institutes of Health and Society for Neuroscience. All protocols were approved by the Institutional Animal Care and Use Committee of the Instituto de Fisiología Celular.

Recordings

Neuronal recordings were obtained with an array of seven independent, movable micro-electrodes (23 M Ω) inserted in the ventral premotor cortex (VPc), dorsal premotor cortex (DPc) and in medial premotor cortex (MPc) in both hemispheres. Neurons were selected if they responded to any of the different components of the detection task. The locations of the electrode penetrations were confirmed with standard histological techniques. Cortical areas were identified based on cortical landmarks. Detailed description of the experimental techniques was described in de Lafuente and Romo 2005, 2006. The experimental data set included 144 recording sessions from two monkeys (47 from monkey R16 and 97 from R19). Each session contained a variable number of simultaneously recorded neurons. The maximum number of simultaneous units was 6 and the median across sessions was 2. The total number of neurons was 384 (117 of monkey R16 and 267 in monkey R19).

Data Analysis

FA Detection by Template Matching

For each neuron, we computed the firing rate using 250 ms sliding windows displaced every 50 ms. We considered the average over hit trials as the neuron-specific typical trajectory triggered by the vibratory stimulation. From this profile we selected a 1 s segment and used it as a template to find similar patterns in single FA trial. The template included the 0.5 s stimulation period and the first 0.5 s of delay period. We slid the template over single FA trials, computing, for each time, the mean squared error between the firing activity on the single trial and the template profile. Because of the 1 s width of the template, this error was defined from the beginning of the trial until 1 s before the end of the delay period. On each trial, a significant match was identified as a FA event if the error presented a minimum that exceeded 1.5 times the error's standard deviation over time. With this algorithm, we found that 347 out of the 384 recorded neurons had at least one FA trial with a FA event and in approximately 80% of the neurons more that 75% of the FA trials contained a FA event.

To test the significance of the detected events we used the activity of simultaneously recoded neurons. We independently detected events on each trial from the activity of the two different neurons. If an event corresponds to a false percept, it should be detected at the same time on simultaneously recorded neurons. We computed the frequency of differences in the detected times, and compared it to both chance level and CR trials. Chance level was obtained by shuffling the trials, keeping the same set of detected times but breaking the trial-to-trial correspondence between neurons. The significance of simultaneous detections (within 350 ms, first bin in the histogram of Figure 2D) was tested with a z test resulting in a $p < 0.001$.

The probability of producing a FA over time was estimated as the number of trials in which of FA event was detected in 500 ms temporal windows, normalized by the total number of FA trials. We corrected for the different trials durations by considering those trials that ended within a time bin as contributing as a fraction to the normalization term.

State-Space Analysis

We constructed pseudo-simultaneous population responses by combining neural data mostly recorded separately. Matching the conditions between different recording sessions resulted in n = 132 neurons from which we had data in every condition (hits and misses of several amplitudes, CR's and FA's). We projected the averaged activity of these neurons onto two task-related axes: stimulus amplitude ($\mathbf{a}_{amp}$) and stimulus detection ($\mathbf{a}_{det}$).

The stimulus amplitude axis, $\mathbf{a}_{amp}$, was obtained as the set of coefficients that best relate each neuron's trial-to-trial response to the stimulus amplitude. To find it, we used a multivariate regression analysis on the firing rate r of each neuron k following

$$r_i^k(t) = \beta_1^k(t) amp_i + \beta_2^k(t) choice_i + \beta_3^k(t) amp_i\, choice_i + \beta_4^k \qquad \text{(Equation 1)}$$

where amp_i and $choice_i$ denote the stimulus amplitude and the subject's choice in trial i, respectively. The stimulus amplitude axis $\mathbf{a}_{amp}$ was defined as the set of coefficients β_1^k for the N recorded neurons (k = 1...N) at the stimulus onset time. The firing rate r was calculated using bins of 100 ms, so they were large enough to include the effect of the stimulus.

$$\mathbf{a}_{amp} = \left[\beta_1^1(t_{SO})\ \beta_1^2(t_{SO}) \ldots \beta_1^N(t_{SO})\ \right]' \qquad \text{(Equation 2)}$$

This axis represents the direction in neural space in which the stimulus drives the network.

The stimulus detection axis, $\mathbf{a}_{det}$, was defined as the vector connecting the population activity just before the application of the stimulus (SO) to that at the end of the delay period (ED), during hit trials.

$$\mathbf{a}_{det} = \boldsymbol{r}_H(t_{ED}) - \boldsymbol{r}_H(t_{SO} - \Delta t) \qquad \text{(Equation 3)}$$

where $\mathbf{r}_H$ is the N-dimensional vector of neural activity averaged over hit trials. The stimulus detection axis, $\mathbf{a}_{det}$, represents the direction in which the network evolves when the subject detects a stimulus.

Recurrent Network Model

We used a recurrent network of n = 500 nonlinear firing-rate units. Each unit is described by an activation variable x_i evolving as,

$$\tau \frac{dx_i}{dt} = -x_i + g\sum_{j=1}^{N} J_{ij} r_j + w_i^{fb} z + w_i^{start} u_{start} + w_i^{stim} u_{stim} \qquad \text{(Equation 4)}$$

where $r_i = \tanh(x_i)$ is the 'firing rate' and $z = \sum_{i=1}^{N} w_i^{out} r_i$ is the network's output. The sparse matrix $\mathbf{J}$ stores the recurrent connection weights and had density p = 0.1, meaning that each element had probability $1 - p$ of being set to 0. The nonzero elements of $\mathbf{J}$ were drawn from a Gaussian distribution with mean zero and variance $1/Np$. The parameter g that scales the strengths of the recurrent connections was set to 1.2. The neuronal time constant is τ = 100 ms and the simulations were performed by Euler integration with a step of dt = 10 ms. The network received two external inputs, u_{start} and u_{stim}, representing the start cue and the sensory channel, respectively. Each neuron received the inputs through a randomly chosen synaptic strength, w_i^{start} and w_i^{stim}. The start cue, u_{start}, is a 500 ms pulse applied at the beginning of each trial. The sensory input, u_{stim}, is modeled as a 300 ms pulsed signal proportional to the stimulus amplitude and embedded in a noisy background. The sensory noise was produced by an Ornstein-Uhlenbeck process of variance σ^2 = 0.4 and correlation time τ = 0.3 s. The decision was extracted from a linear readout of the network activity, z, whose coefficients w^{out} are initially set to zero and then modified by the learning algorithm.

Trials begin with the start cue signal, and after a variable pre-stimulation period, on half of the trials, the stimulus is applied. Trials of different stimulus amplitudes were randomly interleaved with stimulus-absent trials. The pre-stimulation period varies from trial to trial, taking values homogeneously distributed between 0.6 s and 1 s. The stimulus has a duration of 300 ms in every trial, even when its amplitude is 0 (stimulus-absent trials). After the offset of the stimulus there is a delay period which has a fixed duration of 400 ms. The valid decision interval is a 500 s temporal window that starts after the end of delay period.

Training was performed using the FORCE algorithm (Sussillo and Abbott, 2009) to modify the output weights w^{out}. Although these are the only plastic weights, the feedback weights w^{fb} translate this into a rank-one perturbation to the effective recurrent weight matrix $\boldsymbol{J}_{eff} = (g\boldsymbol{J} + \boldsymbol{w}^{fb}\boldsymbol{w}^{out'})$ (Sussillo and Abbott, 2012). The desired output during training trials was zero at all times, except during the decision interval in which it was 1 if the stimulus was present or 0 if it was absent.

Once trained, we quantified the performance of the network through the "psychometric" function. We simulated the trained network for 2,000 trials and obtained the frequency of stimulus-present responses as a function of the stimulus amplitude. To estimate the network's response criterion, we systematically probed the network with variable amplitude stimuli at different times from the beginning of the trial. We followed a bisection protocol for the stimulus amplitudes to look for the lowest amplitude that led to a stimulus-present response. Measuring this borderline stimulus at different times from the start cue gives an estimate of the response criterion over time. We also obtained the rate of FA as a function of time by using the template-matching algorithm described above. We applied to the model the same algorithm that we used for the experimental data.

Fixed-Points Analysis

To analyze the dynamics of the trained network we used the technique developed in (Sussillo and Barak, 2013). The network defined in Equation 4 is a high-dimensional dynamical system. To understand its behavior, we looked for fixed and slow points of phase space by minimizing the function

$$q(\mathbf{X}) = \frac{1}{2}|\boldsymbol{f}(\boldsymbol{x})|^2 \qquad \text{(Equation 5)}$$

where

$$\boldsymbol{f}(\boldsymbol{x}) = -\boldsymbol{x} + \boldsymbol{J}^{\boldsymbol{eff}} \tanh(\boldsymbol{x}) \qquad \text{(Equation 6)}$$

and $\boldsymbol{J}_{eff} = (g\boldsymbol{J} + \boldsymbol{w}^{fb}\boldsymbol{w}^{out'})$. The vector function $\boldsymbol{f}(x)$ defines the nonlinear dynamical system $\boldsymbol{x}' = f(\boldsymbol{x})$, presented in Equation 4.

In order to find minimums of q(x), we simulated the model with several stimulus amplitudes and used the state of the network ($\mathbf{x}$) at different points in time as initial conditions for the minimization algorithm. This procedure systematically identified 3 relevant fixed points. Then, for each point $\boldsymbol{x}^*$ we defined the local linear approximation, $\delta\dot{\boldsymbol{x}} = \mathbf{M}\,\delta\boldsymbol{x}$, where

$$\mathrm{M}_{ij} = \frac{\partial f_i}{\partial x_j} = -\delta_{ij} + J_{ij}^{\mathrm{eff}}\left[1 - \tanh(x_j)^2\right]. \qquad \text{(Equation 7)}$$

By studying the eigenvalues of M, we analyzed the stability of each fixed point.

SUPPLEMENTAL INFORMATION

Supplemental Information includes six figures and can be found with this article at http://dx.doi.org/10.1016/j.neuron.2015.04.014.

AUTHOR CONTRIBUTIONS

N.P. and O.B. designed the study, F.C. and O.B. performed data analyses and modeling, R.R. designed the experiment and executed it with V.d.L., F.C. wrote the paper with input from all the other authors, and N.P., O.B., and R.R. supervised all stages of the project.

ACKNOWLEDGMENTS

We thank Dori Derdikman, Stefano Fusi, and Misha Tsodyks for their comments on the manuscript. Funding was provided by the Spanish grants FIS 2012-33388 (F.C. and N.P.), by an ERC FP7 CIG grant 2013-618543 (O.B.), by an International Research Scholars Award from the Howard Hughes Medical Institute (R.R.), and by grants from the Dirección General de Asuntos del Personal Académico de la Universidad Nacional Autónoma de México and the Consejo Nacional de Ciencia y Tecnología (R.R. and V.d.L.).

Received: September 16, 2014
Revised: January 6, 2015
Accepted: April 3, 2015
Published: May 7, 2015

REFERENCES

Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. Nat. Rev. Neurosci. *10*, 113–125.

Carnevale, F., de Lafuente, V., Romo, R., and Parga, N. (2012). Internal signal correlates neural populations and biases perceptual decision reports. Proc. Natl. Acad. Sci. USA *109*, 18938–18943.

Carnevale, F., de Lafuente, V., Romo, R., and Parga, N. (2013). An optimal decision population code that accounts for correlated variability unambiguously predicts a subject's choice. Neuron *80*, 1532–1543.

Correa, A., Lupiáñez, J., and Tudela, P. (2005). Attentional preparation based on temporal expectancy modulates processing at the perceptual level. Psychon. Bull. Rev. *12*, 328–334.

Coull, J.T., and Nobre, A.C. (2008). Dissociating explicit timing from temporal expectation with fMRI. Curr. Opin. Neurobiol. *18*, 137–144.

Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. Nat. Neurosci. *17*, 1500–1509.

de Lafuente, V., and Romo, R. (2005). Neuronal correlates of subjective sensory experience. Nat. Neurosci. *8*, 1698–1703.

de Lafuente, V., and Romo, R. (2006). Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc. Natl. Acad. Sci. USA *103*, 14266–14271.

Forstmann, B.U., Brown, S., Dutilh, G., Neumann, J., and Wagenmakers, E.J. (2010). The neural substrate of prior information in perceptual decision making: a model-based analysis. Front. Hum. Neurosci. http://dx.doi.org/10.3389/fnhum.2010.00040.

Ghose, G.M., and Bearl, D.W. (2010). Attention directed by expectations enhances receptive fields in cortical area MT. Vision Res. *50*, 441–451.

Ghose, G.M., and Maunsell, J.H. (2002). Attentional modulation in visual cortex depends on task timing. Nature *419*, 616–620.

Gilbert, C.D., and Sigman, M. (2007). Brain states: top-down influences in sensory processing. Neuron *54*, 677–696.

Green, D.M., and Swets, J.A. (1966). Signal Detection Theory and Psychophysics, *Volume 1974* (New York: Wiley).

Hanks, T.D., Mazurek, M.E., Kiani, R., Hopp, E., and Shadlen, M.N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. J. Neurosci. *31*, 6339–6352.

Hoerzer, G.M., Legenstein, R., and Maass, W. (2014). Emergence of complex computational structures from chaotic neural networks through reward-modulated Hebbian learning. Cereb. Cortex *24*, 677–690.

Jaeger, H., and Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science *304*, 78–80.

Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. Nat. Neurosci. *8*, 234–241.

Jaramillo, S., and Zador, A.M. (2011). The auditory cortex mediates the perceptual effects of acoustic temporal expectation. Nat. Neurosci. *14*, 246–251.

Karmarkar, U.R., and Buonomano, D.V. (2007). Timing in the absence of clocks: encoding time in neural network states. Neuron *53*, 427–438.

Katzner, S., Treue, S., and Busse, L. (2012). Improving behavioral performance under full attention by adjusting response criteria to changes in stimulus predictability. J. Vis. *12*, 1.

Laje, R., and Buonomano, D.V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. Nat. Neurosci. *16*, 925–933.

Leon, M.I., and Shadlen, M.N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron *38*, 317–327.

Louie, K., and Wilson, M.A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. Neuron *29*, 145–156.

Luce, R.D. (1986). Response Times: Their Role in Inferring Elementary Mental Organization (Oxford: University Press).

Maass, W., Joshi, P., and Sontag, E.D. (2007). Computational aspects of feedback in neural circuits. PLoS Comput. Biol. *3*, e165.

Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84.

Nobre, A.C. (2001). Orienting attention to instants in time. Neuropsychologia *39*, 1317–1328.

Nobre, A.C., Correa, A., and Coull, J.T. (2007). The hazards of time. Curr. Opin. Neurobiol. *17*, 465–470.

Rao, V., DeAngelis, G.C., and Snyder, L.H. (2012). Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. J. Neurosci. *32*, 10063–10074.

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. *20*, 873–922.

Renart, A., and Machens, C.K. (2014). Variability in neural activity and behavior. Curr. Opin. Neurobiol. *25*, 211–220.

Rohenkohl, G., Cravo, A.M., Wyart, V., and Nobre, A.C. (2012). Temporal expectation improves the quality of sensory information. J. Neurosci. *32*, 8424–8428.

Romo, R., and de Lafuente, V. (2013). Conversion of sensory signals into perceptual decisions. Prog. Neurobiol. *103*, 41–75.

Scheibe, C., Schubert, R., Sommer, W., and Heekeren, H.R. (2009). Electrophysiological evidence for the effect of prior probability on response preparation. Psychophysiology *46*, 758–770.

Shenoy, K.V., Sahani, M., and Churchland, M.M. (2013). Cortical control of arm movements: a dynamical systems perspective. Annu. Rev. Neurosci. *36*, 337–359.

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J.D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. J. Exp. Psychol. Hum. Percept. Perform. *35*, 1865–1897.

Steinmetz, P.N., Roy, A., Fitzgerald, P.J., Hsiao, S.S., Johnson, K.O., and Niebur, E. (2000). Attention modulates synchronized neuronal firing in primate somatosensory cortex. Nature *404*, 187–190.

Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. Neuron *78*, 364–375.

Summerfield, C., and Koechlin, E. (2008). A neural representation of prior information during perceptual inference. Neuron *59*, 336–347.

Sussillo, D., and Abbott, L.F. (2009). Generating coherent patterns of activity from chaotic neural networks. Neuron *63*, 544–557.

Sussillo, D., and Abbott, L.F. (2012). Transferring learning from external to internal weights in echo-state networks with sparse connectivity. PLoS ONE *7*, e37372.

Sussillo, D., and Barak, O. (2013). Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. Neural Comput. *25*, 626–649.

Supplemental Information

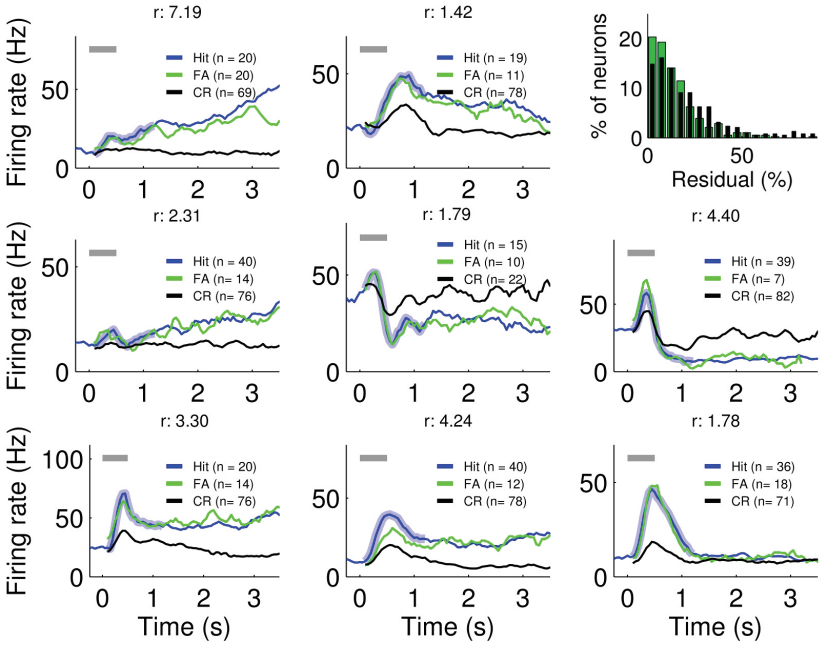

Figure S1: **Examples of realigned activity of FA and CR trials** (Related to Figure 2)

Examples of realigned averaged activity of FA (green) and CR (black) trials for 9 neurons with different firing temporal profiles (analogous to Figure 2C). Single trials in which a FA event was detected were realigned according to the detected time. Blue trace corresponds to the average over strong amplitude Hit trials and the shadow indicates the segment used as template. The number of trials for each condition is indicated by n. The top-right panel shows the histogram over neurons of the difference between realigned FA and Hit profiles (compared with the same measure for CR's realigned profile). The quantity, denoted r, was calculated as the sum of squared residuals, measured as a percent of the mean Hit activity. The distribution of this measure over the recorded neurons is shown for FA (green) and CR (black). Lower residuals are found for the realigned FA profiles than for the realigned CR profile, indicating a significant better match of the former. Indeed, 30% of the recorded neurons had residuals lower than 7% (for visual reference, the values for the example neurons is indicated in each panel).

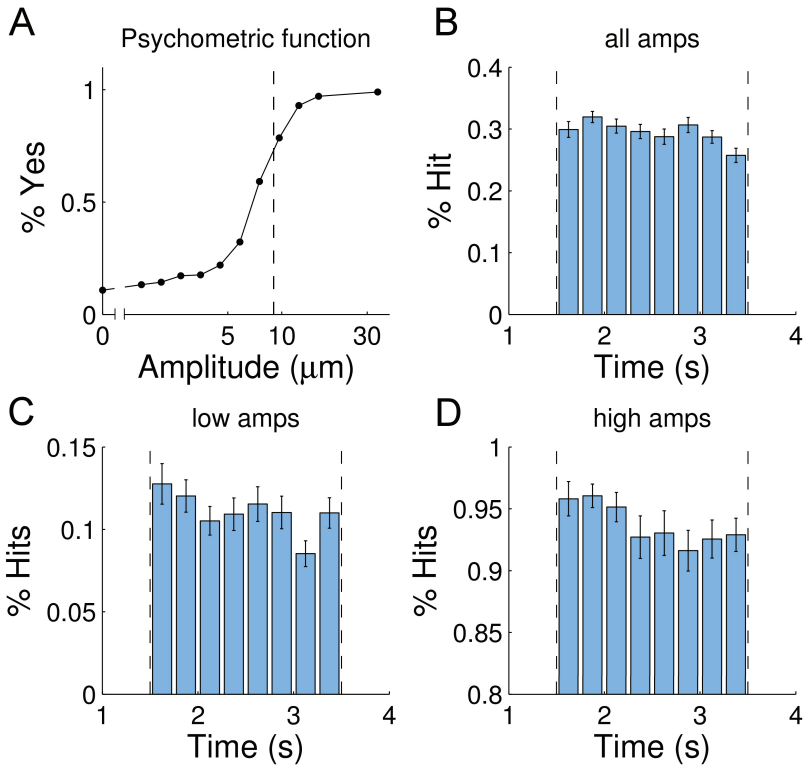

Figure S2: **Analysis of behavioral data.** (Related to Figure 2)
(A) Mean psychometric function averaged over all sessions. (B) Probability of detection as a function of time obtained as the number of correct detections over all stimuli presented in each time bin. Note that this only can be obtained during the period of possible stimulation (from 1.5s to 3.5s). (C) Same as (B) for amplitudes lower than 9μm. (D) Same as (B) for amplitudes higher than 9μm.

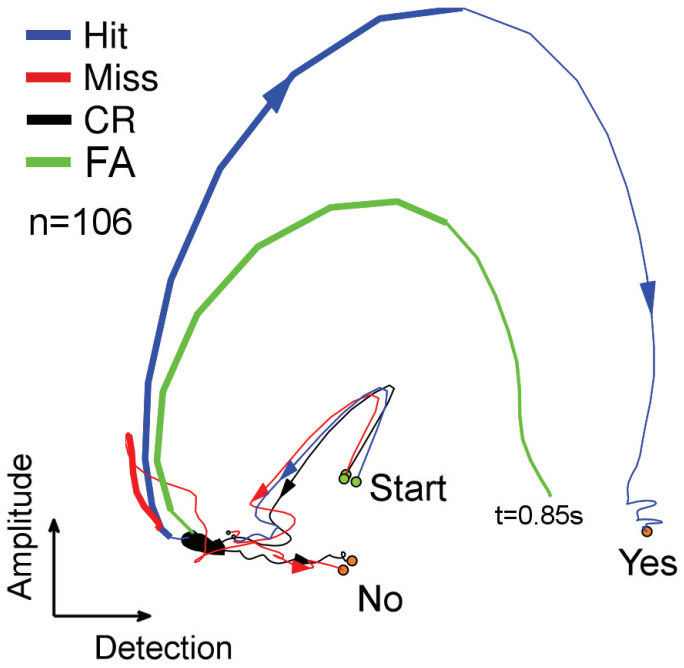

Figure S3: **Realigned neural trajectory during FA alarm trials** (Related to Figure 4)

Average neural trajectories including the realigned FA condition. The neural trajectory for FA trials was obtained by realigning them according to the FA events detected by the template-matching algorithm. Neural trajectories are different from those in Figure 4 because for some sessions there were not enough FA trials and therefore the number of neuron used for this plot is lower ($n = 106$). In addition, the longest time that can be plotted in the FA condition depends on how early an FA event is detected in every session. In our case, this constraint results in a maximum time of 0.85 s from the begining of the detection template.

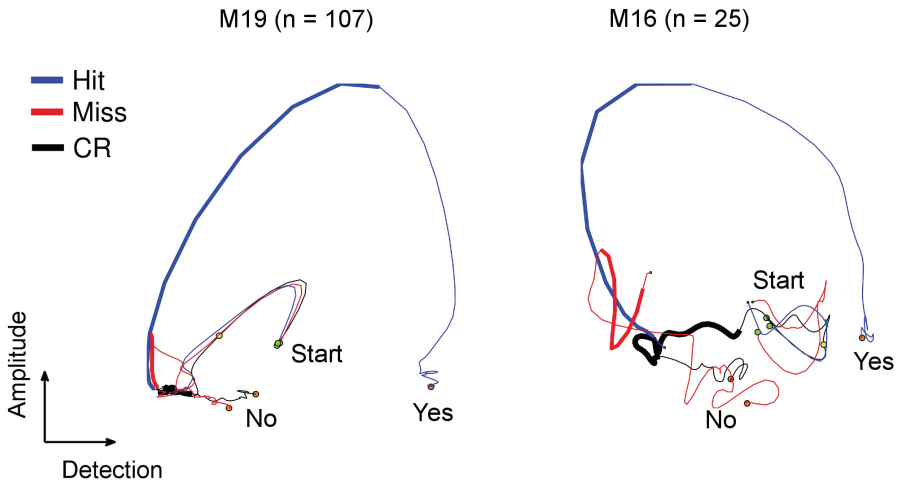

Figure S4: **Neural trajectories obtained separately for the two subjects** (Related to Figure 4)

Figure 4 in the main text was obtained by combining data from two monkeys (M16 and M19). Here we present the same analysis performed separately for each subject. Although the smaller number of neurons deteriorates the clarity of the neural trajectories -especially in subject M16- the same conclusions as in the main text can be reached. The average neural trajectories during Hit (blue), Miss (red) and CR (black) trials were projected onto two task-related axes (stimulus amplitude and stimulus detection). As in Figure 4 of the main text, the trajectories are plotted from the beginning of the trial (green circles) to end of the delay period (orange circles). Stimulus-present conditions are plotted until 1.5 s and realigned at the stimulus onset time. Thick blue and red traces indicate the period of stimulation. The thick black line denotes the possible stimulation window (1.5 s to 3.5s). Units are arbitrary.

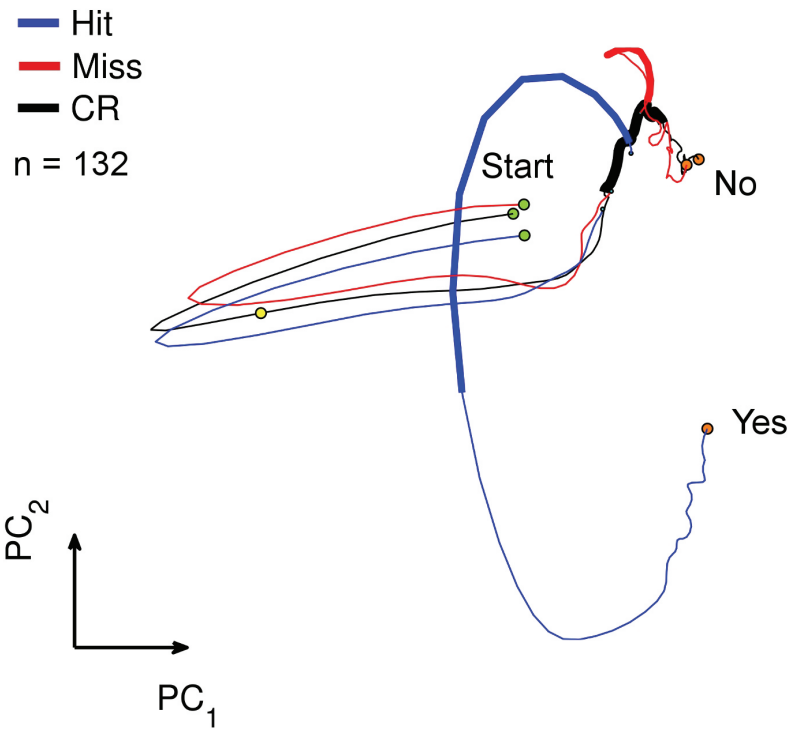

Figure S5: **Neural trajectories projected onto the the principal components** (Related to Figure 4)

Average neural trajectories during Hit (blue), Miss (red) and CR (black) trials projected onto the first two principal components of the data. The modulation in CR's trajectory during the window of possible stimulation (1.5-3.5 s, thick black line) is also visible here, and does not depend crucially on the selected axes (Figure 4, main text). The trajectories are plotted from the beginning of the trial (green circles) to end of the delay period (orange circles). Stimulus-present conditions are plotted until 1.5 s and realigned at the stimulus onset time. Thick blue and red traces indicate the period of stimulation. (1.5 s to 3.5s). Units are arbitrary.

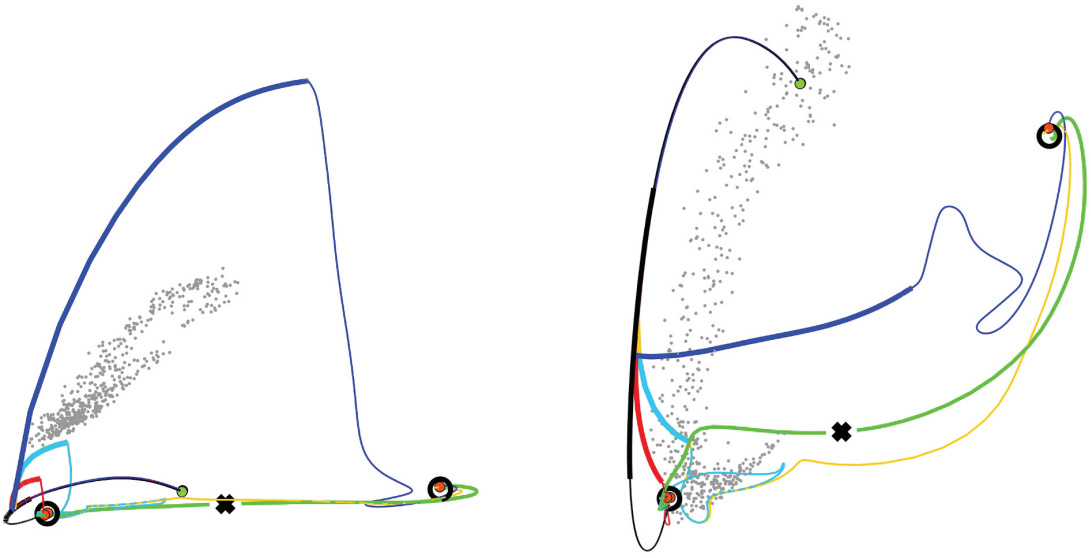

Figure S6: **Neural trajectories close to the separatrix**
(Related to Figure 7)

Same simulations as in Figure 7 but including two trajectories obtained by stimulating with 'borderline' amplitudes (cyan and yellow). (*Left*) Same projection as in Figure 7. (*Right*) A rotation in neural space to visualize how the two 'borderline' trajectories travel close to the separatrix and approach the saddle point (black cross). Afterward, each of them travels to a different attractor.