Stochastic Neural Dynamics as a Principle of Perception

Gustavo Deco and Ranulfo Romo

Abstract Typically, the neuronal firing activity underlying brain functions exhibits a high degree of variability both within and between trials. The key question is: are these fluctuations just a concomitant effect of the neuronal substrate without playing any computational role or do they have a functional relevance? In this chapter, we first review the theoretical framework of stochastic neurodynamics that allows us to investigate the roles of noise and neurodynamics in the computation of probabilistic behavior. The relevance of this framework for neuroscience will be demonstrated by focusing on the simplest type of perceptual task, namely sensory detection. We focus on the following remarkable observation in a somatosensory task: when a near-threshold vibrotactile stimulus is presented, a sensory percept may or may not be produced. These perceptual judgments are believed to be determined by the fluctuation in activity of early sensory cortices. We show, however, that the behavioral outcomes associated with near-threshold stimuli depend of the neuronal fluctuations of more central areas to early somatosensory cortices. The theoretical analysis of the behavioral and neuronal correlates of sensation will show how variability at the neuronal level in those central areas can give rise to probabilistic behavior at the network level and how these fluctuations influence network dynamics.

Introduction

In this chapter, we consider how the noise contributed by the probabilistic spiking times of neurons (spiking noise) plays an important and advantageous role in brain function. We go beyond the deterministic noiseless description of the dynamics of cortical networks, and show how the properties of the system are influenced by the spiking noise. We show that the spiking noise has a significant contribution to the outcome that is reached, in that this noise is a factor in a network with a finite

G. Deco (✉)
Institució Catalana de Recerca i Estudis Avançats (ICREA), Universitat Pompeu Fabra, Passeig de Circumval.lació, 8 08003 Barcelona, Spain
e-mail: gustavo.deco@upf.edu

K. Josić et al. (eds.), *Coherent Behavior in Neuronal Networks*, Springer Series in Computational Neuroscience 3, DOI 10.1007/978-1-4419-0389-1_13,

(i.e., limited) number of neurons. This spiking noise can be described by introducing statistical fluctuations into the finite-size system. It is important that the outcome that is reached, and not just its time course, is influenced on each trial by these statistical fluctuations.

In particular, we will use integrate-and-fire models with spiking neurons to model the actual neuronal data that are obtained from neurophysiological experiments. The integrate-and-fire simulations capture the stochastic nature of the computations. However, we show that to understand analytically (mathematically) the stable points of the network, for example what decisions may be reached, it is helpful to incorporate a mean field approach that is consistent with the integrate-and-fire model. The mean field approach allows one to determine, for example, the synaptic strengths of the interconnected neurons that will lead to stable states of the network, each of which might correspond to a different decision, or no decision at all. The spiking simulations then examine which fixed points (or decisions) are reached on individual trials, and how the probabilistic spiking of the neurons influences these outcomes.

More specifically, we will show that both neurodynamics and stochastic fluctuations matter, in the sense that both have an essential computational role for a complete explanation of perception. To this purpose, we will take as a prototypical example the most elemental and historical task of perceptual detection. By constructing and analyzing computational models, we will establish the link that accounts for measurements both at the cellular and behavioral level. In particular, we show that the behavioral correlate of perceptual detection is essentially given by a noise driven transition in a multistable neurodynamical system. Thus, neuronal fluctuations can be an advantage for brain processing, as they lead to probabilistic behavior in decision-making in this and other sensory tasks. For example, decisions may be difficult without noise. In the choice dilemma described in the medieval Duns Scotus paradox, a donkey who could not decide between two equidistant food rewards might suffer the consequences of the indecision. The problem raised is that with a deterministic system, there is nothing to break the symmetry, and the system can become deadlocked. In this situation, the addition of noise can produce probabilistic choice, which is advantageous, as will be described in this paper.

Brain Dynamics: From Spiking Neurons to Reduced Rate-Models

The computation underlying brain functions emerges from the collective dynamics of spiking neuronal networks. A spiking neuron transforms a large set of incoming input spike trains, coming from different neurons, into an output spike train. Thus, at the microscopic level, neuronal circuits of the brain encode and process information by spatiotemporal spike patterns. We assume that the transient (nonstationary) dynamics of spiking neurons is properly captured by one-compartment, point-like models of neurons, such as the leaky integrate-and-fire (LIF) model [38]. In the LIF model, each neuron i can be fully described in terms of a single internal variable,

namely the depolarization $V_i(t)$ of the neural membrane. The basic circuit of a LIF model consists of a capacitor C in parallel with a resistor R driven by a synaptic current (excitatory or inhibitory postsynaptic potential, EPSP or IPSP, respectively). When the voltage across the capacitor reaches a threshold θ, the circuit is shunted to a reset potential V_{reset}, and a δ-pulse (spike) is generated and transmitted to other neurons. The subthreshold membrane potential of each neuron evolves as a simple RC-circuit, with a time constant $\tau = RC$ given by the following equation:

$$\tau \frac{\mathrm{d}V_i(t)}{\mathrm{d}t} = -[V_i(t) - V_{\mathrm{L}}] + \tau \sum_{j=1}^{N} J_{ij} \sum_k \delta(t - t_j^{(k)}), \tag{1}$$

where V_{L} is the leak potential of the cell in the absence of external afferent inputs and the total synaptic current flow into cell i is given by the sum of the contributions of δ-spikes produced at presynaptic neurons, with J_{ij} the efficacy of synapse j and $t_j^{(k)}$ the emission time of the kth spike from the jth presynaptic neuron.

In the brain, local neuronal networks comprise a large number of neurons which are massively interconnected. The set of coupled differential equations (1) above describe the underlying dynamics of such networks. Direct simulations of these equations yield a complex spatiotemporal pattern, covering the individual trajectory of the internal state of each neuron in the network. This type of direct simulation is computationally expensive, making it very difficult to analyze how the underlying connectivity relates to various dynamics. One way to overcome these difficulties is by adopting the population density approach, using the Fokker–Planck formalism [21, 22, 28]. We will follow here a derivation done by Stefano Fusi (private communication). In this approach, individual integrate-and-fire neurons are grouped together into populations of statistically similar neurons. A statistical description of each population is given by a probability density function that expresses the distribution of neuronal states (i.e., membrane potential) over the population. In general, neurons with the same state $V(t)$ at a given time t have a different history because of random fluctuations in the input current $I(t)$. The main source of randomness is from fluctuations in the currents. The key assumption in the population density approach is that the afferent input currents impinging on neurons in one population are uncorrelated. Thus, neurons sharing the same state $V(t)$ in a population are indistinguishable. The population density $p(\upsilon, t)$ expresses the fraction of neurons at time t that have a membrane potential $V(t)$ in the interval $[\upsilon, \upsilon + d\upsilon]$. The evolution of the population density is given by the Chapman–Kolmogorov equation

$$p(\upsilon, t + \mathrm{d}t) = \int_{-\infty}^{+\infty} p(\upsilon - \varepsilon, t)\rho(\varepsilon|\upsilon - \varepsilon)\mathrm{d}\varepsilon, \tag{2}$$

where $\rho(\varepsilon|\upsilon) = \mathrm{Prob}\{V(t + \mathrm{d}t) = \upsilon + \varepsilon | V(t) = \upsilon\}$ is the conditional probability that generates an infinitesimal change $\varepsilon = V(t+\mathrm{d}t) - V(t)$ in the infinitesimal interval $\mathrm{d}t$. The temporal evolution of the population density can be reduced to a simpler differential equation by the mean-field approximation. In this approximation, the

currents impinging on each neuron in a population have the same statistics, because, as mentioned above, the history of these currents is uncorrelated. The mean-field approximation entails replacing the time-averaged discharge rate of individual cells with a common time-dependent population activity (ensemble average). This assumes ergodicity for all neurons in the population. The mean-field technique allows us to discard the index denoting the identity of any single neuron. The resulting differential equation describing the temporal evolution of the population density is called the Fokker–Planck equation, and reads

$$\frac{\partial p(\upsilon,t)}{\partial t} = \frac{1}{2\tau}\sigma^2(t)\frac{\partial^2 p(\upsilon,t)}{\partial \upsilon^2} + \frac{\partial}{\partial \upsilon}\left[\left(\frac{\upsilon - V_{\mathrm{L}} - \mu(t)}{\tau}\right)p(\upsilon,t)\right]. \tag{3}$$

In the particular case that the drift is linear and the diffusion coefficient $\sigma^2(t)$ is given by a constant, the Fokker–Planck equation describes a well-known stochastic process called the Ornstein–Uhlenbeck process [31]. The Ornstein–Uhlenbeck process describes the temporal evolution of the membrane potential $V(t)$ when the input afferent current is $\mu(t) + \sigma\sqrt{\tau}w(t)$, with $w(t)$ a white noise process. This can be interpreted, by means of the Central Limit Theorem, as the case in which the sum of many Poisson processes becomes a normal random variable with mean $\mu(t)$ and variance σ^2.

The nonstationary solutions of the Fokker–Planck equation (3) describe the dynamical behavior of the network. However, these simulations, as the direct simulation of the original network of spiking neurons (1), are computationally expensive and their results probabilistic, which makes them unsuitable for systematic explorations of parameter space. On the other hand, the stationary solutions of the Fokker–Planck equation (3) represent the stationary solutions of the original integrate-and-fire neuronal system. The stationary solution of the Fokker–Planck equation satisfying specific boundary conditions (see [3, 23, 30]) yields the *population transfer function* of Ricciardi (ϕ):

$$\nu = \left[t_{\mathrm{ref}} + \tau\sqrt{\pi}\int_{\frac{V_{\mathrm{reset}}-V_{\mathrm{L}}-\mu}{\sigma}}^{\frac{\theta-V_{\mathrm{L}}-\mu}{\sigma}} \mathrm{e}^{x^2}\{1+\mathrm{erf}(x)\}\mathrm{d}x\right]^{-1} = \phi(\mu,\sigma), \tag{4}$$

where $\mathrm{erf}(x) = 2/\sqrt{\pi}\int_0^x \mathrm{e}^{y^2}\mathrm{d}y$. In last equation t_{ref} is the refractory time.

The population transfer function gives the average population firing rate as a function of the average input current. For more than one population, the network is partitioned into populations of neurons whose input currents share the same statistical properties and fire spikes independently at the same rate. The set of stationary, self-reproducing rates ν_i for different populations i in the network can be found by solving a set of coupled self-consistency equations, given by:

$$\nu_i = \phi(\mu_i, \sigma_i), \tag{5}$$

This reduced system of equations allows a thorough investigation of the parameters. In particular, one can construct bifurcation diagrams to understand the nonlinear mechanisms underlying equilibrium dynamics and in this way solve the "inverse problem," i.e., the selection of the parameters that generate the attractors (steady states) that are consistent with the experimental evidence. This is the crucial role of the mean-field approximation: to simplify analyses through the stationary solutions of the Fokker–Planck equation for a population density under the diffusion approximation (Ornstein–Uhlenbeck process) in a self-consistent form. After that, one can perform full nonstationary simulations using these parameters in the integrate-and-fire scheme to generate *true dynamics*. The mean field approach ensures that these dynamics will converge to a stationary attractor that is consistent with the steady-state dynamics we require [3, 16]. The stochastic (random) firing times of neurons introduces noise into neuronal networks, and it is the consequences of this randomness expressed in a finite (limited) sized network of such neurons with which we are concerned in this review. We show that the noise in such systems not only helps us to understand many aspects of decision-making as implemented in the brain, but is in fact beneficial to the operation of decision-making processes.

The mean-field approach has been applied to model single neuronal responses, fMRI activation patterns, psychophysical measurements, and the effects of pharmacological agents and of local cortical lesions [4, 5, 8, 10–15, 32, 37].

Perceptual Detection and Stochastic Dynamics

Neurophysiology

The detection of sensory stimuli is among the simplest perceptual experiences and is a prerequisite for any further sensory processing. A fundamental problem posed by the sensory detection tasks is that repeated presentation of a near-threshold stimulus might unpredictably fail or succeed in producing a sensory percept. Where in the brain are the neuronal correlates of these varying perceptual judgments? Pioneering studies on the neuronal correlates of sensory detection showed that, in the case of vibrotactile stimuli, the responses of S1 neurons account for the measured psychophysical accuracy [27]. However direct comparisons between S1 responses and detection performance were not directly addressed and, therefore, it is not clear whether the activity of S1 accounts for the variability of the behavioral responses. Psychophysical performance was measured in human observers and S1 recordings were made in anesthetized monkeys.

This problem has been recently addressed [6, 7]. These authors trained monkeys to perform a detection task. In each trial, the animal had to report whether the tip of a mechanical stimulator vibrated or not. Stimuli were sinusoidal, had a fixed frequency of 20 Hz and were delivered to the glabrous skin of one fingertip. Crucially, they varied in amplitude across trials. Stimulus-present trials were interleaved

with an equal number of stimulus-absent trials in which no mechanical vibrations were delivered. Depending on the monkeys' responses, trials could be classified into four types of responses: hits and misses in the stimulus-present condition, and correct rejections and false alarms in the stimulus-absent condition. Stimulus detection thresholds were calculated from the behavioral responses. Thus an important issue in this and similar tasks is to determine the neuronal correlates that account for these behavioral reports.

De Lafuente and Romo [6] simultaneously characterized the activity of S1 neurons (areas 3b and 1) and the monkey's psychophysical performance by recording the extracellular spike potentials of single S1 neurons while the monkeys performed the detection task. Figure 1 shows the experimental design and main results. To test whether the responses of S1 neurons accounted for the monkey's psychophysical performance, [6] calculated neurometric detection curves and compared them with the psychometric curves. The proportion of "yes" responses for neurometric curves was defined, for a given amplitude, as the proportion of trials in which the neuron's firing rate reached or exceeded a criterion value [6, 18]. For each neuron, this criterion was chosen to maximize the number of correct responses. Pairwise comparisons of detection thresholds obtained from logistic fits to the simultaneously obtained neurometric and psychometric data showed that the detection thresholds of individual S1 neurons were not significantly different from the animals' psychophysical thresholds, and the two thresholds measures highly covaried. In addition, the shape

Fig. 1 The detection task. (**a**) Drawing of monkey working in the detection task. (**b**) The sequence of events during the detection trials. Trials began when the stimulation probe indented the skin of one fingertip of the left, retrained hand (probe down, PD). The monkey then placed its right, free hand on an immovable key (key down, KD). On half of the randomly selected trials, after a variable prestimulus period (Prestim, 1.5–3.5 s), a vibratory stimulus (Stim, 20 Hz, 0.5 s) was presented. Then, after a fixed delay period (Delay, 3 s), the stimulator probe moved up (probe up, UP), indicating to the monkey that it could make the response movement (MT) to one of the two buttons. The button pressed indicated whether or not the monkey felt the stimulus. Henceforth referred to as yes and no responses, respectively. (**c**) Depending on whether the stimulus was present or absent and on the behavioral response, trial outcome was classified as a hit, miss, false alarm (FA), or correct reject (CR). Trials were pseudo-randomly chosen: 90 trials were stimulus absent (amplitude 0), and 90 trials were stimulus present with varying amplitudes (9 amplitudes with 10 repetitions each). (**d**) Classical psychometric detection curve obtained by plotting the proportion of yes responses as a function of the stimulus amplitude. (**e**) Mean firing rate of hit trials for S1 ($n = 59$) and MPC ($n = 50$) neurons. (**f**) Comparison of normalized neuronal population activity of S1 neurons during hits and misses for near-threshold stimuli, and during correct rejections and false alarms in stimulus-absent trials. Normalized activity was calculated as a function of time, using a 200 ms window displaced every 50 ms. This was calculated by substractng the mean activity and dividing by the standard deviation of the activity from a 200 ms window of the prestimulus period. *Lower panels* show the choice probability index as a function of time. This quantity measures the overlap between two response distributions: in this case, between hits and misses and between correct rejection and false alarm trials. *Dotted lines* mark significance levels. (**g**) The same as in **f**, but for a neuronal population activity of MPC neurons. Adapted with permission from De LaFuente and Romo, 2006

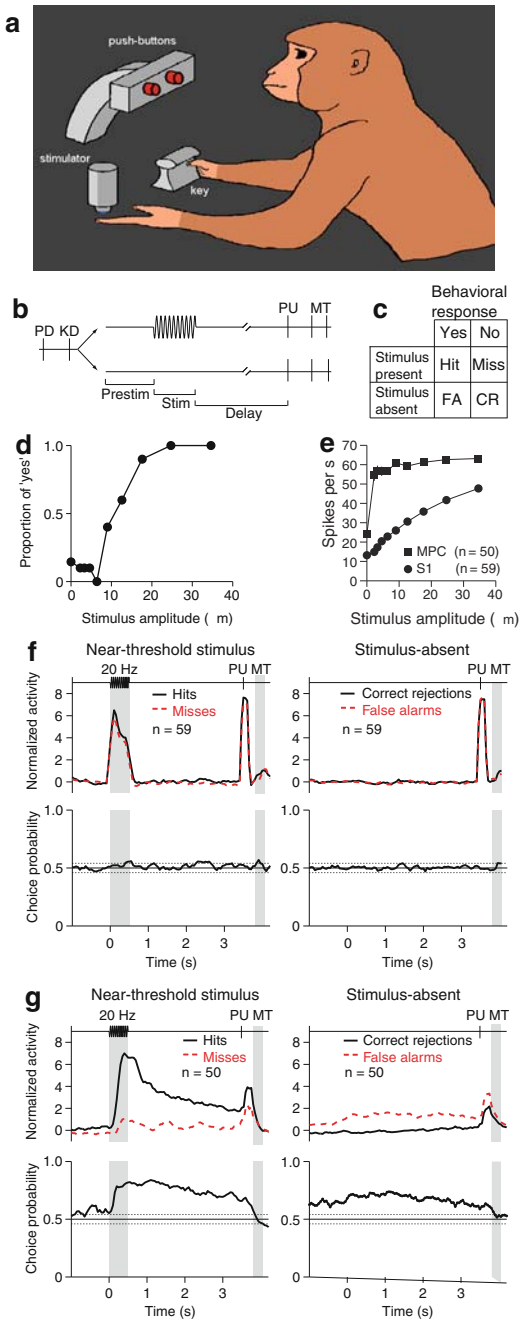
a
push-buttons
stimulator
key
b
PD KD
PU MT
Prestim
Stim
Delay
c
Behavioral response
Yes No
Stimulus present Hit Miss
Stimulus absent FA CR
d
Proportion of 'yes'
Stimulus amplitude (m)
e
Spikes per s
MPC (n = 50)
S1 (n = 59)
Stimulus amplitude (m)
f
Near-threshold stimulus
Stimulus-absent
20 Hz
PU MT
Hits
Misses
n = 59
Correct rejections
False alarms
Normalized activity
Choice probability
Time (s)
g
Near-threshold stimulus
Stimulus-absent
n = 50

of the mean neurometric curve resulting from the activity of the S1 neurons showed close correspondence with the shape of the mean psychometric curve.

An important question addressed in this study is whether the activity of S1 neurons covaries with the perceptual "yes"–"no" judgments that the monkeys made on a trial-by-trial basis [6]. To test this, these authors compared the activity during hit and miss trials for the near-threshold stimulus, as well as for the corresponding activity in correct reject and false alarm trials in the stimulus-absent condition. They found no significant differences in the activity of S1 neurons between hits and misses, nor between correct rejections and false alarms. This indicated that activity of individual S1 neurons did not predict the monkey's behavior. To further quantify this, [6] calculated a choice probability index, which estimates the probability with which the behavioral outcome can be predicted from the neuronal responses [2, 20]. Again they found no significant differences between hits and misses, or between correct rejections and false alarm trials.

The low choice probability values are consistent with a detection model in which the activity of S1 serves as input to an additional processing stage(s) that determines whether a stimulus has occurred or not. Under this hypothesis, the correlation between S1 activity and the final decision about the stimulus presence or absence is highly dependent on the amount of correlated-noise among sensory neurons [39]. Indeed, [6] found that the mean noise correlation coefficient across pairs of S1 neurons was 0.16 ± 0.02. This amount of correlated-noise is similar to that reported in previous studies [1, 35, 39], and is also consistent with the near chance choice probability values reported in the study of [6]. These results further support a detection model in which a central area(s) must be observing the activity of S1 neurons to judge about the stimulus presence or absence. Therefore, the functional role of S1 in this and other perceptual tasks may be mainly to generate a neural representation of the sensory stimulus for further processing in areas central to it [19, 33–36]. However, a previous study found that fMRI signals in primary visual cortex (V1) reflected the percepts of human subjects, rather than the encoded stimulus features [29]. This result suggests that, in V1, top-down signals (nonsensory inputs delivered to visual cortex via feedback projections) can be combined with bottom-up (sensory) information and contribute to sensory percepts [29]. S1 data did not show evidence for this type of neural interaction; rather, it indicated that S1 represents the physical properties of stimuli and contributes little to near-threshold percepts [6]. The discrepancy could be due to fundamentally different organizations across sensory cortices, or to differences between species. Another possibility to consider is that the modulation revealed through fMRI may have an effect that is invisible from the point of view of single neurons. This would happen if, for instance, such modulation acted only to synchronize the spikes of multiple target neurons [17].

To test whether the neuronal correlates of the perceptual decisions associated with detection might reside outside S1, [6] recorded neurons from the medial premotor cortex (MPC), a frontal cortical area known to be involved in the evaluation of sensory information and in decision-making processes [20, 24]. They found that, in contrast to the graded dependence on stimulus amplitude observed in S1, MPC neurons responded in an all-or-none manner that was only weakly modulated by

the stimulus amplitude, but that closely correlated with "yes" and "no" behavioral responses. The mean normalized activity across the MPC neurons was strong and sustained, and with near-threshold stimuli it was clearly different for hit and miss trials. Moreover, almost 70% of the false alarm trials were predicted from increases in neuronal activity in stimulus-absent trials. de Lafuente and Romo [6] also found that the MPC activity preceding stimulus onset was higher during hits than during misses. These early increases in activity predicted detection success significantly above chance levels, as is shown by the choice probability plots. Although de Lafuente and Romo (2005, 2006) do not know the role of this increased prestimulus activity, they speculate that it might be associated with trial history during a run. To support this conjecture, [6] wondered about the behavioral responses on trials previous to false alarm responses. They found that the probability of an "yes" response was increased in those trials preceding a false alarm, supporting the notion that monkeys were biased toward "yes" responses. de Lafuente and Romo [6] speculated that given that "yes" responses to the three subthreshold amplitudes were rewarded, this could have encouraged the monkeys to respond "yes" in the next trial, producing a false alarm response. The results indicate that the responses of all MPC neurons studied were associated with stimulus presence or with false alarms; that is, with "yes" responses. They did not find neurons whose increases in their activities were associated with "no" responses. [6] do not know the reason for this but they speculate that "no" is the default response that is installed from trial beginning and that the stimulus presentation overrides this default response.

The close association between neuronal responses and behavioral responses, and the weak relationship between neuronal activity and stimulus amplitude, supported the interpretation that MPC neurons do not code the physical attributes of stimuli, but rather represent perceptual judgments about their presence of absence. As the monkeys reported their decisions by a motor act, a key question needed to be answered: was the MPC activity truly related to stimulus perception, or was it simply reflecting the different motor actions associated with the two response buttons? To test this, [6] designed a control task in which the correct response button was illuminated at the beginning of every trial. In this variant of the detection task, the monkeys simply had to wait until the end of the trial to push the illuminated button, without the need to attend to the presence or absence of the mechanical vibration. It is important to note that in this test condition all-or-none activity was still observed in response to the near-threshold stimulus, and the probability of activation depended on the stimulus amplitude, similar to that observed in the standard detection task. Given that in the control test the monkeys did not have to choose a response button based on the vibratory stimulus, the results are consistent with the interpretation that the activity of MPC neurons is related to the subjective perception of sensory stimuli, rather than to the selection of the motor plan. These results therefore favor the hypothesis that this MPC population reflects the failure or success of the near-threshold stimulus in triggering a sensory percept.

A Computational Model of Probabilistic Detection

An aim of this chapter is to show how stochastic dynamics helps to understand the computational mechanisms involved in perceptual detection. The computational analysis of detection focuses on the paradigm and experimental results of [6] described above. In summary, they used a behavioral task where trained awake monkeys report the presence or absence of a mechanical vibration applied to their fingertips by pressing one of two pushbuttons. They found that the activity of MPC neurons was only weakly modulated by the stimulus amplitude, and covaried with the monkeys' trial-by-trial reports. On the contrary, S1 neurons did not covary with the animals' perceptual reports, but their firing rate did show a monotonically increasing graded dependence on the stimulus amplitude (see Fig. 1d and e). The fact that MPC neurons correlate with the behavioral performance, with a high firing rate for an "yes" report and a low firing rate for a "no" report, suggests an underlying bistable dynamic in an attractor framework.

A minimal network model is now described that captures the computation involved in perceptual detection and is consistent with the neurophysiological and behavioral evidence described [9]. The main idea of the model is to establish a neurodynamics that shows two possible *bistable* decision states associated with the two possible behavioral responses: "stimulus detection" and "no stimulus detection." The computation underlying perceptual detection is then understood as a fluctuation-driven, probabilistic transition to one of the two possible bistable decision states.

A patch of MPC neurons in the frontal lobe is modeled by a network of interacting neurons organized into a discrete set of populations. Populations are defined as groups of excitatory or inhibitory neurons sharing the same inputs and connectivities. Some of the excitatory population of neurons have a selective response, which reflects the sensitivity to an external applied vibrotactile stimulus (note that for simplicity in 1A only one selective population is shown for the single specific vibrotactile frequency utilized in the experiment). All other excitatory neurons are grouped in a "Non-selective" population. There is also one inhibitory population of local inhibitory neurons that regulates the overall activity by implementing competition in the network. Neurons in the networks are connected via three types of receptors that mediate the synaptic currents flowing into them: AMPA and NMDA glutamate receptors, and GABA receptors. Neurons within a specific excitatory population are mutually coupled with a strong weight ω_+. Neurons between two different selective populations have anticorrelated activity, which results in weaker connections ω_-.

In this model, activity in a selective excitatory population corresponds to the detection of a percept associated with an external applied vibrotactile stimulus. The strength of the input (λ) impinging on that excitatory population is proportional to the strength of the presented vibrotactile stimulus (as for example encoded in S1, i.e., the input to MPC is transmitted from S1). When a stimulus is presented, there is just one population sensitive to it. To model this characteristic, we use a network composed of two selective populations, but only one will be selective to

the stimuli applied. The relevant bistability is therefore given by the state where the excitatory populations have low activity (corresponding to no detection of a percept, i.e., a "no" response), and the state where the excitatory population sensitive to the presented vibrotactile stimulus is highly activated (corresponding to the detection of the percept, i.e., an "yes" response). We refer to this model as "Non-Competing Yes-Neurons" (NCYN) (Fig. 2a). Just the selective population sensitive to the applied vibrotactile stimulation used in the experiment is represented by a specific excitatory population. (A full specification of the whole connectivity is provided in [9].)

The characteristics of the network in the stationary conditions were studied with the mean-field approach reviewed above. Using this approximation, the relevant parameter space given by the population cohesion ω_+ vs. the external input λ was scanned. The mean-field results for the NCYN-model are illustrated in a phase diagram (Fig. 2b) that shows different regimes of the network. For small values of λ and for a weak population cohesion, the network has one stable state where all populations are firing at a weak level (spontaneous state). This spontaneous state encodes the "no" response in the NYCN model. For higher population cohesion and higher values of λ, a state corresponding to strong activation of the selective population sensitive to the applied vibrotactile stimulation emerges. We call this excited state encoding the "yes" response, the "yes" state. Between these two regions, there is a bistable region where the state corresponding to weak ("no" response) or strong ("yes" response) activation states of the selective population sensitive to the applied vibrotactile stimulation are both stable.

To study the probabilistic behavior of the neuronal dynamics of the network, the spiking simulations of the configurations corresponding to the region of bistability were analyzed with methods similar to those used for the neurophysiological data by [6].

The results are presented of the nonstationary probabilistic analysis calculated by means of the full spiking simulations averaged over several trials. In all cases, the aim was to model the behavior of the MPC neurons which are shown in Fig. 1d and e, which reflect detection of the percept [6]. It is proposed that the perceptual response results from a neurodynamical bistability [9]. In this framework, each of the stable states corresponds to one possible perceptual response: "stimulus detected" or "stimulus not detected." The probability of detecting the stimulus is given by the transitions between these two states. In fact, the probabilistic character of the system results from the stochastic nature of the networks. The source of this stochasticity is the approximately random spiking of the neurons in the finite-size network. We note that there are two sources of noise in such spiking networks: the randomly arriving external Poissonian spike trains and the fluctuations due to the finite size of the network. Here we refer to finite-size effects due to the fact that the populations are described by a finite number N of neurons. In the mean-field framework, (see [25,26]) "incoherent" fluctuations due to quenched randomness in the neurons' connectivity and/or to external input are already taken into account in the variance, and "coherent" fluctuations give rise to new phenomena. In fact, the number of spikes emitted by the network in a time interval $[t, t + \mathrm{d}t)$ is a Poisson variable with mean and variance $N\nu(t)\mathrm{d}t$. The estimate of $\nu(t)$, is then a stochastic process $\nu_N(t)$, well

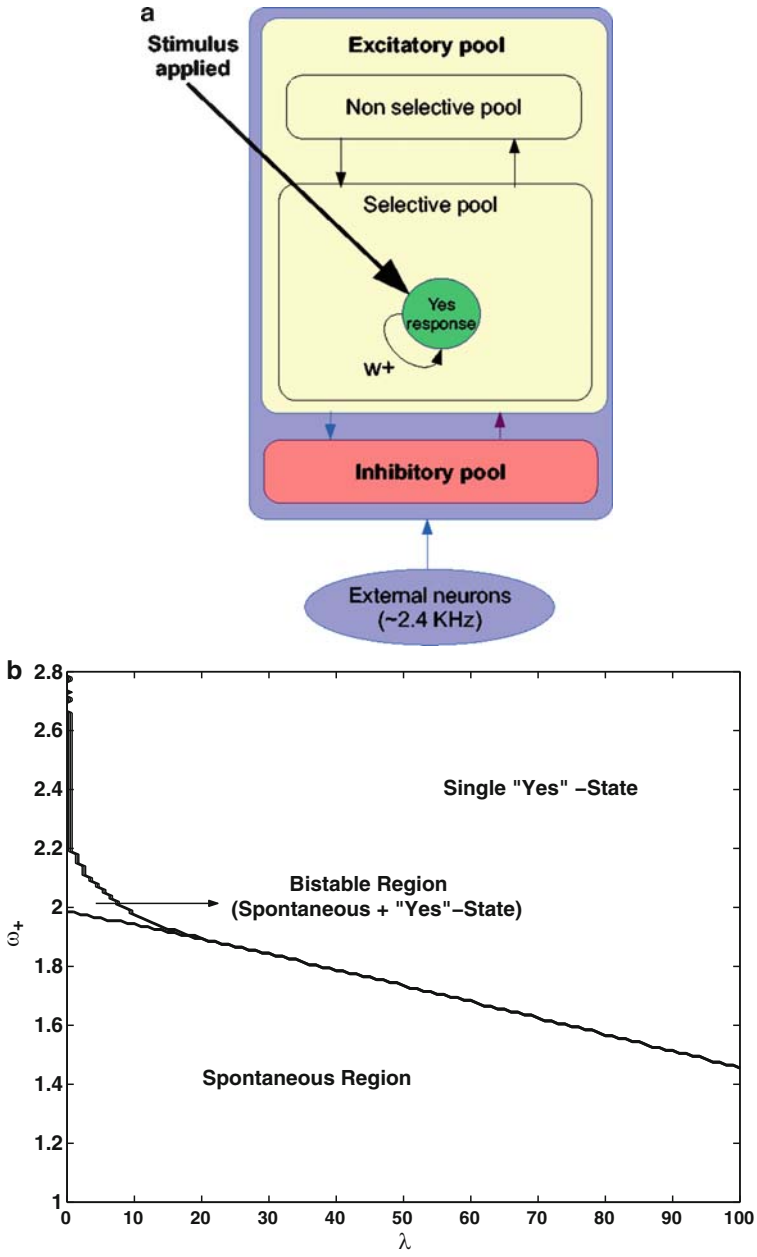

Fig. 2 (**a**) The perceptual detection model (NCYN) has excitatory populations selective to the applied vibrotactile stimulation. A "no" response is given when the selective population has low activity and an "yes" response when it has high activity. The *arrows* indicate the recurrent connections between the different neurons in a pool. (See text for more details). (**b**) Phase diagrams for the NYCN model for parameter exploration of the attractor states (stationary states) of the underlying dynamical system. The diagrams show the different attractor regions as a function of the stimulus input (λ) and the level of coupling ω_+ within the neurons of the same selective population (cohesion)

described in the limit of large $N\nu$ by $\nu_N(t) \simeq \nu(t) + \sqrt{\nu(t)/N}\xi(t)$, where $\xi(t)$ is Gaussian white noise with zero mean and unit variance, and $\nu(t)$ is the probability of emitting a spike per unit time in the infinite-size network. Such finite-N fluctuations, which affect the global activity ν_N, are coherently felt by all neurons in the network and lead to an additive Gaussian noise corrections in the mean-field equations.

To compare the theoretical results with the experimental results, the characteristics of the bistable neurodynamical model NCYN were studied. The behavior of the relevant populations encoding the different bistable states corresponding to the two alternative choices is shown in Fig. 3. Figure 3a plots the proportion of "yes" responses as a function of the intensity of the applied vibrotactile stimulation, i.e., as a function of the strength λ of the stimulus presented. The figure shows that the proportion of "yes" responses (hits) increases as the intensity of the stimulus applied grows. The model is consistent with the experimental results of Lafuente and Romo shown in Fig. 1. Hence, the model shows a probabilistic behavior that emulates the real behavior of subjects detecting a vibrotactile stimulus [6]. Let us now concentrate on the level of firing activity observed in MPC neurons that covary with the behavioral responses. Figure 3b shows the activity of the neurons encoding the "yes" response (selective excitatory population sensitive to the applied vibrotactile stimulus) averaged over trials that reported a percept (hits). In the model, the mean firing activity is almost constant and is not linearly related to the stimulus amplitude, as reflected in the experimental results. The fact that neurons encoding the "yes" response present a relatively constant level of activation on trials that report a detected percept, whereas on trials that fail to detect a percept these neurons have low activity (spontaneous level), is consistent with an attractor network. Therefore the transitions driven by the spiking-related statistical fluctuations are consistent with the behavioral data.

Deco et al. [9] studied also a second different bistable network model called "Competing Yes-No-Neurons" (CYNN). Both models (NCYN and CYNN) are consistent with the existing single cell recordings, but they involve different types of bistable decision states, and consequently different types of computation and neurodynamics. By analyzing the temporal evolution of the firing rate activity of neurons on trials associated with the two different behavioral responses, they were able to produce evidence in favor of the CYNN model. Specifically, the CYNN model predicts the existence of some neurons that encode the "no" response, and other neurons that encode the "yes" response. The first set of neurons slightly decrease their activity at the end of the trial, whereas the second group of neurons increase their firing activity when a stimulus is presented. Thus in this case, the simulations indicate that the CYNN model fits the experimental data better than the NCYN model.

In conclusion, computational stochastic neurodynamical models provide a deeper understanding of the fundamental mechanisms underlying perceptual detection and how these are related to experimental neuroscience data. We argue that addressing such a task is a prerequisite for grounding empirical neuroscience in a cogent theoretical framework.

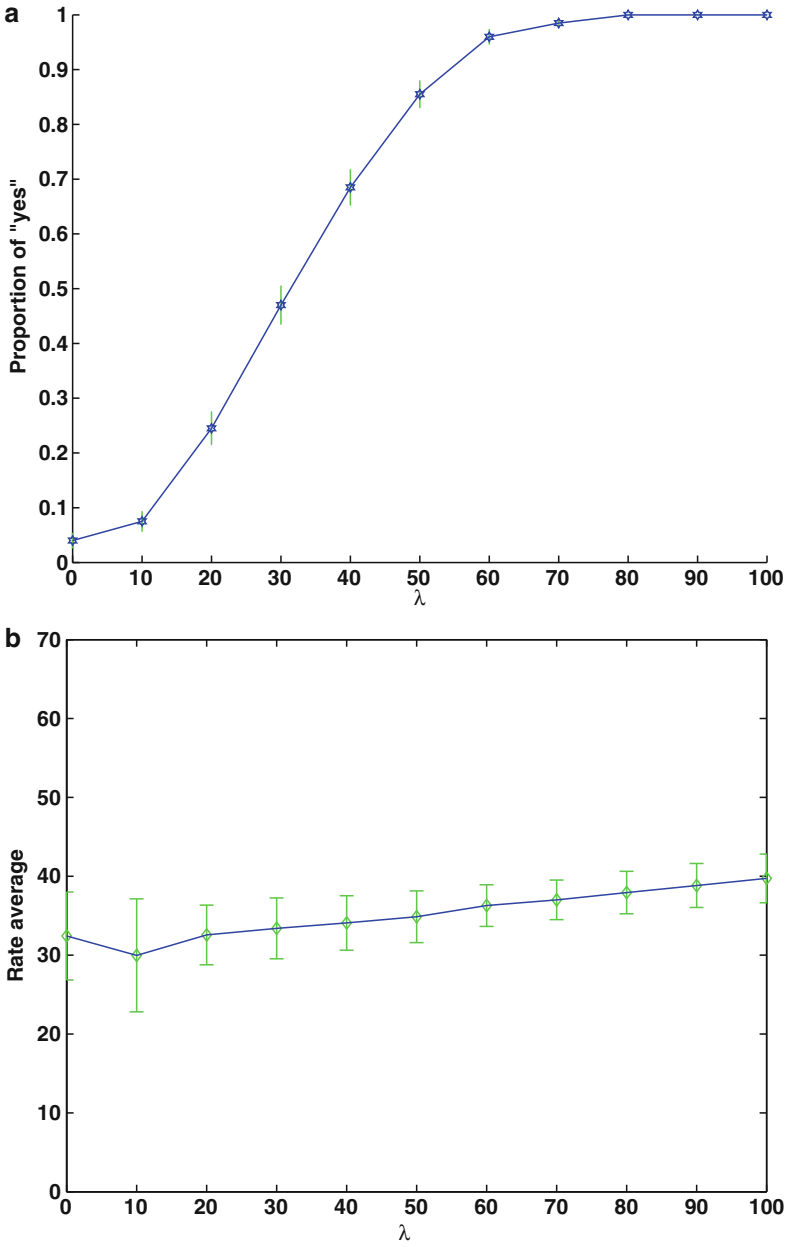

Fig. 3 Simulated results plotting the detection curves resulting from 200 trials (overall performance) and the mean rate activity of hit trials at a function of the input strength λ for the MPC neurons for the experimental design of de Lafuente and Romo (2005). (**a**) Probability of an "yes" response (hit). (**b**) Mean firing rate activity of neurons in the "yes" population on "yes" trials. The simulations of the nonstationary and probabilistic behavior of the neurodynamical activity were performed by a full spiking and synaptic simulation of the network

References

1. W. Bair, E. Zohary, and W. Newsome. Correlated firing in macaque in visual mt: Time scales and relationship to behavior. *Journal of Neuroscience*, 21:1676–1697, 2001.
2. K. Britten, W. Newsome, M. Shadlen, S. Celebrini, and J. Movshon. A relationship between behavioral choice and the visual responses of neurons in macaque mt. *Visual Neuroscience*, 13:87–100, 1996.
3. N. Brunel and X. J. Wang. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11:63–85, 2001.
4. S. Corchs and G. Deco. Large-scale neural model for visual attention: integration of experimental single cell and fMRI data. *Cerebral Cortex*, 12:339–348, 2002.
5. S. Corchs and G. Deco. Feature based attention in human visual cortex: simulation of fMRI data. *Neuroimage*, 21, 36–45, 2004.
6. V. de Lafuente and R. Romo. Neuronal correlates of subjective sensory experience. *Nature Neuroscience*, 8:1698–1703, 2005.
7. V. de Lafuente and R. Romo. Neural correlates of subjective sensory experience gradually builds up accross cortical areas. *Proceedings of the National Academy of Science USA*, 103:14266–14271, 2006.
8. G. Deco and T. S. Lee. A unified model of spatial and object attention based on inter-cortical biased competition. *Neurocomputing*, 44–46:775–781, 2002.
9. G. Deco, M. Perez-Sanagustin, V. de Lafuente, and R. Romo. Perceptual detection as a dynamical bistability phenomenon: a neurocomputational correlate of sensation. *Proceedings of the National Academy of Sciences USA*, 104:20073–20077, 2007.
10. G. Deco, O. Pollatos, and J. Zihl. The time course of selective visual attention: theory and experiments. *Vision Research*, 42:2925–2945, 2002.
11. G. Deco and E. Rolls. Neurodynamics of biased competition and cooperation for attention: a model with spiking neurons. *Journal of Neurophysiology*, 94:295–313, 2005.
12. G. Deco and E. T. Rolls. Object-based visual neglect: a computational hypothesis. *European Journal of Neuroscience*, 16:1994–2000, 2002.
13. G. Deco and E. T. Rolls. Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex. *European Journal of Neuroscience*, 18:2374–2390, 2003.
14. G. Deco and E. T. Rolls. A neurodynamical cortical model of visual attention and invariant object recognition. *Vision Research*, 44:621–644, 2004.
15. G. Deco, E. T. Rolls, and B. Horwitz. 'What' and 'where' in visual working memory: a computational neurodynamical perspective for integrating fMRI and single-neuron data. *Journal of Cognitive Neuroscience*, 16:683–701, 2004.
16. P. Del Giudice, S. Fusi, and M. Mattia. Modeling the formation of working memory with networks of integrate-and-fire neurons connected by plastic synapses. *Journal of Physiology Paris*, 97:659–681, 2003.
17. P. Fries, J. Reynolds, A. Rorie, and R. Desimone. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291:1560–1563, 2001.
18. D.M. Green and J.A. Swets. *Signal Detection Theory and Psychophysics*. Wiley, New York, 1966.
19. A. Hernandez, A. Zainos, and R. Romo. Neuronal correlates of sensory discrimination in somatosensory cortex. *Proceedings of the National Academy of Sciences USA*, 97:6191–6196, 2000.
20. A. Hernandez, A. Zainos, and R. Romo. Temporal evolution of a decision-making process in medial premotor cortex. *Neuron*, 33:959–972, 2002.
21. B. Knight. Dynamics of encoding in neuron populations: some general mathematical features. *Neural Computation*, 12(3):473–518, 2000.
22. B.W. Knight, D. Manin, and L. Sirovich. Dynamical models of interacting neuron populations. In E.C. Gerf, ed. *Symposium on Robotics and Cybernetics: Computational Engineering in Systems Applications*. Cite Scientifique, Lille, France, 1996.

23. P. Lansky, L. Sacerdote, and F. Tomassetti. On the comparison of Feller and Ornstein-Uhlenbeck models for neural activity. *Biological Cybernetics*, 73:457–465, 1995.
24. L. Lemus, A. Hernandez, R. Luna, A. Zainos, V. Nacher, and R. Romo. Neural correlates of a postponed decision report. *Proceedings of the National Academy of Sciences USA*, 104:17174–17179, 2007.
25. M. Mattia and P.D. Giudice. Population dynamics of interacting spiking neurons. *Phys Rev E*, 66(5):051917, 2002.
26. M. Mattia and P.D. Giudice. Finite-size dynamics of inhibitory and excitatory interacting spiking neurons. *Phys Rev E*, 70:052903, 2004.
27. V. Mountcastle, W. Talbot, H. Sakata, and J. Hyvarinen. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. neuronal periodicity and frequency discrimination. *Journal of Neurophysiology*, 32:452–484, 1969.
28. A. Omurtag, B.W. Knight, and L. Sirovich. On the simulation of large populations of neurons. *Journal of Computational Neuroscience*, 8:51–53, 2000.
29. D. Ress and D. Heeger. Neuronal correlates of perception in early visual cortex. *Nature Neuroscience*, 6:414–420, 2003.
30. L. Ricciardi and L. Sacerdote. The Ornstein-Uhlenbeck process as a model for neuronal activity. I. mean and variance of the firing time. *Biological Cybernetics*, 35:1–9, 1979.
31. H. Risken. *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer Verlag, Berlin, 1996.
32. E. T. Rolls and G. Deco. *Computational Neuroscience of Vision*. Oxford University Press, Oxford, 2002.
33. R. Romo, C. Brody, A. Hernandez, and L. Lemus. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, 339:470–473, 1999.
34. R. Romo, A. Hernandez, A. Zainos, L. Lemus, and C. Brody. Neural correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, 5:1217–1225, 2002.
35. R. Romo, A. Hernandez, A. Zainos, and E. Salinas. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron*, 38:649–657, 2003.
36. E. Salinas, A. Hernandez, A. Zainos, and R. Romo. Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *Journal of Neuroscience*, 20:5503–5515, 2000.
37. M. Szabo, G. Deco, S. Fusi, P. Del Giudice, M. Mattia, and M. Stetter. Learning to attend: Modelling the shaping of selectivity in inferotemporal cortex in a categorization task. *Biological Cybernetics*, 94:351–365, 2006.
38. H. Tuckwell. *Introduction to Theoretical Neurobiology*. Cambridge University Press, Cambridge, 1988.
39. E. Zohary, M. N. Shadlen, and W. T. Newsome. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370:140–143, 1994.