Progress in Neurobiology 103 (2013) 194–213

Contents lists available at SciVerse ScienceDirect

Progress in Neurobiology

journal homepage: www.elsevier.com/locate/pneurobio

Brain mechanisms for perceptual and reward-related decision-making

Gustavo Deco [a], Edmund T. Rolls [b,1,*], Larissa Albantakis [c], Ranulfo Romo [d]

[a] Institució Catalana de Recerca i Estudis Avançats (ICREA), Universitat Pompeu Fabra, Dept. of Technology, Computational Neuroscience, Roc Boronat 138, 08018 Barcelona, Spain
[b] Oxford Centre for Computational Neuroscience, Oxford, UK
[c] Universitat Pompeu Fabra, Dept. of Technology, Computational Neuroscience, Roc Boronat 138, 08018 Barcelona, Spain
[d] Instituto de Fisiología Celular-Neurociencias, Universidad Nacional Autónoma de México, 04510 México, D.F., Mexico

ARTICLE INFO

Article history:
Received 22 September 2011
Received in revised form 24 January 2012
Accepted 24 January 2012
Available online 2 February 2012

Keywords:
Decision making
Attractor network
Confidence
Dynamical neuropsychiatry
Noise in the brain
Vibrotactile decision-making
Reward value decision-making

ABSTRACT

Phenomenological models of decision-making, including the drift–diffusion and race models, are compared with mechanistic, biologically plausible models, such as integrate-and-fire attractor neuronal network models. The attractor network models show how decision confidence is an emergent property; and make testable predictions about the neural processes (including neuronal activity and fMRI signals) involved in decision-making which indicate that the medial prefrontal cortex is involved in reward value-based decision-making. Synaptic facilitation in these models can help to account for sequential vibrotactile decision-making, and for how postponed decision-related responses are made. The randomness in the neuronal spiking-related noise that makes the decision-making probabilistic is shown to be increased by the graded firing rate representations found in the brain, to be decreased by the diluted connectivity, and still to be significant in biologically large networks with thousands of synapses onto each neuron. The stability of these systems is shown to be influenced in different ways by glutamatergic and GABAergic efficacy, leading to a new field of dynamical neuropsychiatry with applications to understanding schizophrenia and obsessive–compulsive disorder. The noise in these systems is shown to be advantageous, and to apply to similar attractor networks involved in short-term memory, long-term memory, attention, and associative thought processes.

Contents

Abbreviations: 2AFC, 2-alternative forced-choice; AMPA, α-amino-3-hydroxy-5-methyl-4-isoazoleproprionic acid; DDM, drift diffusion model; dlPFC, dorso-lateral prefrontal cortex; EPSC, excitatory post-synaptic current; ER, error rate; FEF, frontal eye-field; FF, feedforward; FFI, feedforward inhibition; GABA, γ-aminobutyric acid; IPSC, inhibitory post-synaptic current; LCA, leaky competing accumulator; LIF, leaky integrate-and-fire; LIP, lateral intraparietal cortex; LR, likelihood ratio; MT, middle temporal area; NMDA, N-methyl-D-aspartate acid; O–U, Ornstein–Uhlenbeck; PDF, probability density function; PFC, prefrontal cortex; PPC, posterior parietal cortex; PRR, parietal reach region; RDM, random-dot motion; RF, response field; R-target, response target; RT, reaction time; SAT, speed-accuracy tradeoff; SC, superior colliculus; SDT, signal detection theory; SPRT, sequential probability ratio test.

* Corresponding author.
E-mail address: Edmund.Rolls@oxcns.org (E.T. Rolls).
[1] http://www.oxcns.org

0301-0082/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.pneurobio.2012.01.010

1. Introduction

One approach to understanding decision-making is incorporated in a class of phenomenological "sequential sampling" models, which analyze the decision process in terms of a decision variable that evolves in time until it reaches a decision threshold (e.g. Laming, 1968; Ratcliff and Smith, 2004; Stone, 1960; Vickers, 1970). Recordings of single neuron activity during decision-making in primates (reviewed by Gold and Shadlen, 2007; Hernandez et al., 2010; Opris and Bruce, 2005; Romo and Salinas, 2003) provided constraints for these models, and also stimulated the development of biologically plausible models of decision-making (Deco and Rolls, 2006; Rolls and Deco, 2010; Wang, 2002). In Section 2 of this paper we provide an overview of different approaches to decision-making, describing the evolution from phenomenological models to neurophysiologically plausible integrate-and-fire attractor models of decision-making. We particularly highlight differences between the models and approaches.

In Section 3 we describe the development of these biologically plausible models to incorporate synaptic facilitation. We show how synaptic facilitation can be an important component in sequential decision-making tasks, and in tasks in which the report of the decision must be postponed for a delay period.

In Section 4 we show how confidence in a decision before the outcome is known is an emergent property of the decision-making mechanism in the brain. We also show how a second decision based on the confidence in the first decision can be made using a second attractor network. This provides a mechanism for monitoring previous decisions.

In Section 5 we show how noise in the brain can be used to predict a decision even before the evidence for the decision has been provided.

In Section 6 we consider factors that influence the level of noise in a decision-making network. This treatment provides a crucial foundation for understanding how the proposed mechanisms operate in large decision-making cortical networks with graded firing rate representations and diluted connectivity in the human brain.

In Section 7 we describe the applications of this stochastic neurodynamics approach to neuropsychiatry, including schizophrenia and obsessive–compulsive disorder.

In Section 8, we consider some effects of oscillations in the brain on decision-making processes.

2. Overview of different models of decision-making

In the following, we will review the most common models of two-alternative forced choice (2AFC) decision-making and their theoretical origins. We will start with basic, linear, conceptual models, which successfully capture decision behavior, followed by attempts to implement these models in a physiologically plausible way. Then we will turn to nonlinear attractor models and describe a biophysically inspired implementation of an attractor model with spiking neurons (Deco and Rolls, 2006; Rolls and Deco, 2010; Wang, 2002, 2008). Our objective is to provide an intuitive overview. Consequently, we restrict our formal presentation to basic equations and characteristic model features and refer to the original publications for detailed theoretical analysis.

2.1. Sequential-sampling models

Present conceptual models of decision behavior considering noisy evidence build on signal detection theory (SDT), developed to describe categorical choices under uncertainty (Green and Swets, 1966; Tanner and Swets, 1954). SDT typically assumes fixed, short stimulus times that are out of the subject's control. The class of models summarized as 'sequential sampling models' forms the logical extension of SDT to temporally stretched streams of (noisy) data (Stone, 1960; Wald, 1947). In addition to the probability of correct responses, these models give predictions of subjects' reaction times in 'free response' 2AFC paradigms. To form a decision, evidence for each of the two alternatives is integrated over time. Whether an independent integration for each alternative (e.g. a race model), or an integration of the difference in evidence (e.g. a drift–diffusion model) gives a better account of experimental 2AFC data, is, however, still open to debate, although the latter seems to fit a wider set of experimental observations (Bogacz et al., 2006; Ratcliff et al., 2003; Ratcliff and Smith, 2004).

2.1.1. Signal detection theory and the sequential probability ratio test (SPRT)

In simple perceptual 2AFC tasks, subjects are often faced with problems such as: 'Has a dim light been flashed or not?' or: 'Which of two similar images has been presented?' Signal detection theory provides a prescription for these kinds of decisions, where one of two hypotheses has to be chosen on the basis of a single observation in the presence of uncertainty, or noise (Gold and Shadlen, 2007). If the sensory observation is informative about the hypotheses, it provides 'evidence' favoring one alternative. We will generally refer to information that is indicative of a choice as evidence e. The two hypotheses H_1 and H_2 stand for the two choice-alternatives. The conditional probability $p(e|H_1)$ denotes the probability of observing evidence e if H_1 is true.

Depending on the signal-to-noise ratio (μ/σ) and the similarity of the hypotheses (μ_1–μ_2), the probability density functions (PDFs) of the two alternatives overlap to some degree (Fig. 1).

The smaller the signal-to-noise ratio is, the higher is the overlap of the PDF. Likewise, the more distinguishable the stimuli are, the smaller is the overlap. In the case of sensory stimuli the PDFs are

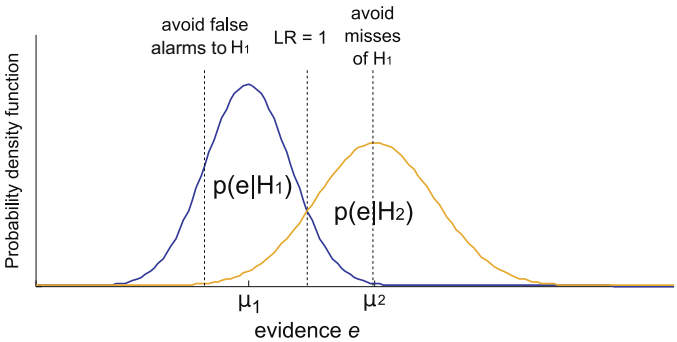

Fig. 1. Signal detection theory in two-alternative forced choice tasks. Because of uncertainty the probability density functions (PDFs) of the two alternative hypotheses overlap. A choice is made depending on the desired level of accuracy for one of the alternatives. Comparing the likelihood ratio (LR) to 1 minimizes the total number of errors.

often assumed to be normally distributed with means $\mu_1 \neq \mu_2$ and standard deviations $\sigma_1 = \sigma_2$.

The *a posteriori* probability $p(H_1|e)$ that hypothesis H_1 is true given the evidence e can be determined according to Bayes' theorem from the conditional probability $p(e|H_1)$, the *prior*, or *a priori* probability of the hypothesis $p(H_1)$, and the total probability of the evidence $p(e)$:

$$p(H_1|e) = \frac{p(e|H_1)\,p(H_1)}{p(e)}. \tag{1}$$

The prior $p(H_1)$ thereby denotes the probability that H_1 is true before any evidence has been obtained. If equal priors are assumed for both alternatives, H_1 is more likely to be correct than H_2 if the 'likelihood ratio' $LR(e) = p(e|H_1)/p(e|H_2)$ is larger than 1.

Choosing H_1 if $LR > 1$ is the optimal strategy, in the sense that it provides the lowest overall error rate. In the case of equal rewards or costs, it also indicates the optimal choice in terms of the highest reward.

For some decisions, however, the consequences of a false alarm, for example, are negligible compared to missing a signal. Because of the noise, mistakes are inevitable. Still, the type of errors, i.e. *false alarms* or *misses*, can be adjusted by the decision criterion (Fig. 1). Generally, the desired level of accuracy for one of the alternatives determines the decision threshold, or bound, B. For unequal prior probabilities, but identical rewards, H_1 should be chosen if $\mathrm{LR}(e) > B = p(H_2)/p(H_1)$ (Green and Swets, 1966). In sensorimotor tasks, unequal priors arise for instance if one stimulus is presented more often than the other.

If not just one, but multiple pieces of evidence $e_1 \ldots e_N$ are available over time, as for instance in the random-dot motion (RDM) task, in which subjects have to decide on the net direction of motion from a patch of moving dots, the likelihood ratio has to be updated with each new sample of evidence. With the simplifying assumption that all evidence samples $e_1 \ldots e_N$ are independent, the likelihood ratio becomes

$$\mathrm{LR}(e) = \frac{p(H_1|e_1 \ldots e_N)}{p(H_2|e_1 \ldots e_N)} = \prod_{n=1}^{N} \frac{p(e_n|H_1)}{p(e_n|H_2)}. \tag{2}$$

A decision bound $B = 1$ again minimizes the error rate, as it determines the most likely hypothesis (Neyman and Pearson, 1933). From the perspective of 2AFC problems, Eq. (2) applies to the 'interrogation' paradigm, where the decision is based on a fixed sample of evidence.

In the 'free response' paradigm, where the decision-maker is allowed to control the decision time autonomously, she or he is faced with the additional problem of *when* to end the evidence accumulation. Accordingly, optimality in free response tasks is often assessed as the strategy that yields the shortest expected reaction time (RT) for a given error rate.

The sequential probability ratio test provides a solution to this specific optimality problem (Wald, 1947). Here, the momentary likelihood ratio $\mathrm{LR}(e)$ is again calculated as in Eq. (2), but instead of one, there are now two decision bounds B_1 and B_2 and the sampling process continues as long as

$$B_2 < \mathrm{LR}(e) < B_1, \quad \text{with} \quad B_2 < B_1. \tag{3}$$

In other words, if B_1 is crossed, alternative 1 is selected, if B_2 is crossed, alternative 2 is selected, and while the evidence for both alternatives is insufficient, meaning below a certain level, the decision process continues. Interestingly, a decision rule equivalent to Eq. (3) can be obtained using any quantity that is monotonically related to the LR if B is scaled appropriately (Green and Swets, 1966). Hence, by taking the logarithm of Eqs. (2) and (3), the decision process can be written as a simple addition:

$$\log(B_2) < \log(\mathrm{LR}(e)) = \sum_{n=1}^{N} \log \frac{p(e_n|H_1)}{p(e_n|H_2)} < \log(B_1). \tag{4}$$

Moreover, the temporal evolution of the log-likelihood ratio (log LR) can be described as a discrete decision variable V, starting at $V_0 = 0$, which is subsequently updated at each time step, according to:

$$V^n = V^{n-1} + \log \frac{p(e_n|H_1)}{p(e_n|H_2)}. \tag{5}$$

Using the log-likelihood ratio to express the SPRT has the further advantage that evidence in favor of H_1 intuitively adds to V with a positive value, while evidence supporting H_2 contributes negatively. In that sense, the trajectory of the decision variable $V(t)$ for noisy evidence is analogous to a one-dimensional 'random walk' bounded by a positive and negative threshold.

In the limit of infinitesimally small time steps, equivalent to continuous sampling, the discrete SPRT/random walk model converges to the drift–diffusion model (DDM) described in the next section. A more detailed theoretical description of optimality, including the case of unequal priors, and the continuum limit of the SPRT, is provided elsewhere (Bogacz et al., 2006).

Before we turn to the drift–diffusion model, we briefly discuss how the theory presented above might relate to real neural computations during decision-making, and the random-dot motion task in particular. According to Shadlen and colleagues (Churchland et al., 2008; Gold and Shadlen, 2001, 2002, 2007; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001), decision-related neural activity in area LIP (lateral intraparietal cortex) is consistent with the notion of an 'accumulation-to-bound' process, while area MT (middle temporal cortical area) encodes the absolute amount of visual motion present in the random dot motion stimulus and might consequently provide evidence to LIP (Britten et al., 1992, 1993). Could LIP activity actually correspond to a neural decision variable in the mathematical sense of the sequential probability ratio test?

As the brain is unlikely to store the complete distribution of possible neural responses to every encountered stimulus, it probably has no access to the PDFs of the neural populations, which would be necessary to infer the likelihood ratio LR (Gold and Shadlen, 2001; but see Ma et al., 2006).

However, motivated by the apparent analogy between the trajectory of V and LIP firing rates, Gold and Shadlen (2001, 2002, 2007) argued that a quantity approximating the log LR could indeed be computed in the brain. More precisely, knowledge about the PDFs would not be explicitly necessary to implement a decision rule approximating the optimal sequential probability ratio test, if output firing rates of two antagonistic sensory neurons or neural

populations were used as evidence. One example would be the responses I_1 and I_2 of two populations of MT neurons, one selective for rightward, the other for leftward motion, which respond to their preferred and null direction with mean firing rates $\mu_1 > \mu_2$, and roughly equal variance σ. In that case, the optimal log LR decision rule will depend only on the firing rate difference $I_1 - I_2$, apart from a scaling factor.

This largely holds true for a variety of possible PDFs (Gold and Shadlen, 2001). In particular:

$$\log \mathrm{LR}_{\mathrm{left,right}} = \frac{\mu_1 - \mu_2}{\sigma^2}(I_1 - I_2), \tag{6}$$

if I_1 and I_2 are sampled from normal distributions, which is a plausible assumption for the average firing rate of a neural population. However, the responses of single neurons may better be described by a Poisson distribution. In that case

$$\log \mathrm{LR}_{\mathrm{left,right}} = \log\left(\frac{\mu_1}{\mu_2}\right)(I_1 - I_2). \tag{7}$$

Knowing the sign of the difference $I_1 - I_2$ in MT activity would hence be sufficient for downstream areas like LIP to elicit a left or right saccade according to a sequential probability ratio test optimal rule.

Furthermore, a study by Platt and Glimcher (1999) revealed that both the prior probability of getting a reward, and the expected magnitude of the reward, could modulate LIP activity, consistent with the suggestion that LIP activity might be a neural correlate of the decision-variable V (Eq. (5)).

As a final note on the sequential probability ratio test, the argument of Gold and Shadlen (2001, 2002, 2007) can also be extended to multiple alternatives, or neural populations, which results in a comparison between the neural population with the highest rate and the average rate of the other populations ('max-vs-average test') (Ditterich, 2010). However, contrary to the 2-alternative case, the resulting statistical test is not optimal. Interestingly, the optimal algorithm for decisions between more than two alternatives is still unknown (McMillen and Holmes, 2006). The multihypothesis SPRT was shown to approximate optimality for small error rates (Dragalin et al., 1999). For moderate error rates, the physiologically plausible max-vs-average test performs almost as well as the multihypothesis SPRT (Ditterich, 2010).

2.1.2. The drift–diffusion model

The continuum limit of the SPRT represents the most basic form of the DDM. A continuous decision variable $v(t)$ is accumulating the evidence difference between the two choice-alternatives, or hypotheses (Laming, 1968; Ratcliff, 1978; Stone, 1960). In the unbiased case with equal prior probabilities, $v(t)$ is integrated over time according to

$$dv(t) = \mu dt + \sigma dW, \quad v(0) = 0, \tag{8}$$

with symmetric decision bounds $b_1 = -b_2$, and the accumulation time interval dt (assumed to be very small). Eq. (8) is the continuous extension of Eq. (5). The right side of Eq. (8) denotes the new noisy evidence obtained during dt. It is composed of a constant drift μdt, with drift rate μ, and the diffusion term $\sigma\, dW$, which represents white noise drawn from a Gaussian distribution with mean 0 and variance $\sigma^2 dt$. (dW is proportional to $N(0,1)\cdot\sqrt{dt}$, as the variance of uncorrelated stochastic random variables is additive with successive

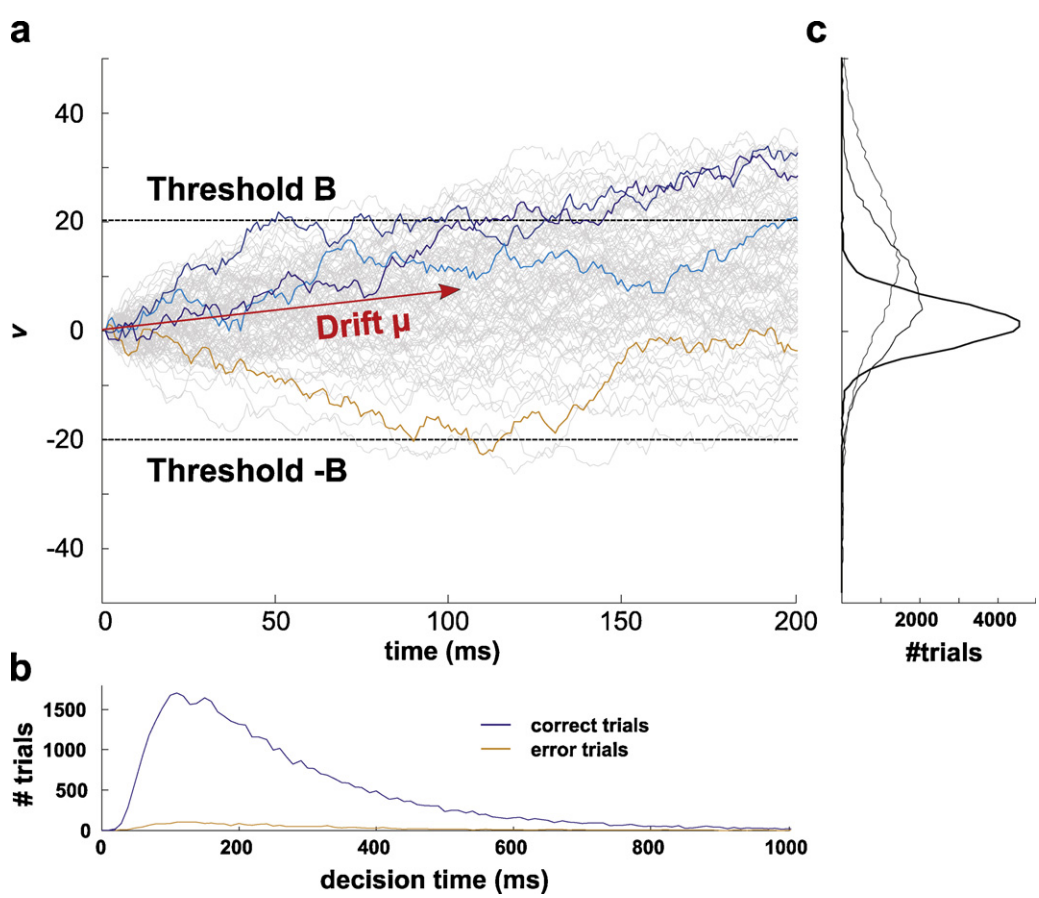

Fig. 2. Basic drift–diffusion model. (a) 100 example traces of the time evolution of $v(t)$. Three correct trials were labeled in blue, one error trial in orange. (b) Left-skewed RT Histograms of correct and error choices from 50,000 trials. (c) The variance of v increases with time. The distribution of v for 50,000 trials is given for t = 20 ms, 100 ms and 200 ms (bold to narrow lines). Model parameters: μ = 0.07, σ = 1, dt = 1 ms, B = 20.

time-steps, which leads to a square-root behavior for the standard deviation (Usher and McClelland, 2001).) The correct alternative is determined by μ, which, in the case of the random dot motion task, can be interpreted as the amount of coherent motion. Using the terminology of the sequential probability ratio test, if $\mu > 0$, H_1 is correct; if $\mu < 0$, H_2 is correct. Which alternative is eventually chosen by the DDM, however, is also subject to noise, depending on σ, and the height of the decision bounds b_1 and b_2. Still, the DDM, as a continuous implementation of the SPRT, solves two-alternative forced choice problems optimally: it will on average return a decision in the shortest possible time for a specified level of accuracy (Bogacz et al., 2006; Wald, 1947). (This description does not consider costs associated with behavior, or a cost function for gathering information that might change over time.)

Solutions of Eq. (8) are normally distributed with probability density $p(v,t) = N(\mu t, \sigma\sqrt{t})$, if the decision bounds are ignored (Gardiner, 1985). Consequently, the variance across trials of the temporal evolution of v increases linearly with t. This can be appreciated in Fig. 2, where several example trials of v and the variance are displayed.

Due to the threshold nonlinearity of the decision bounds, the reaction time distributions of correct and error trials are typically left-skewed, with equal means (Fig. 2b). Treated as a so called 'first-passage time problem', error rates (ER) and mean RTs of the basic DDM can be expressed as

$$\mathrm{ER} = \frac{1}{1 + e^{2\mu b/\sigma^2}}, \tag{9}$$

$$\mathrm{mRT} = \frac{b}{\mu}\tanh\left(\frac{\mu b}{\sigma^2}\right) + t_{\mathrm{ND}}, \tag{10}$$

where t_{ND} denotes the "non-decision" time (e.g. the time related to sensory and motor processes which add to RT).

Contrary to the theoretical predictions of the basic DDM, error responses in 2AFC tasks can have significantly different mean RTs than correct responses, depending on the experimental specifications, e.g. stressing accuracy or speed, or the difficulty of a condition (Luce, 1991; Pleskac and Busemeyer, 2010; Ratcliff and Smith, 2004). A more general, extended version of the DDM includes trial-to-trial variability in the drift rate and the starting point (Ratcliff and McKoon, 2008; Ratcliff and Rouder, 1998). A normally distributed drift rate with mean μ and standard deviation s_μ leads to longer RTs on error trials, as errors will occur more often in trials where μ is small. Drawing the starting point $v(0)$ from a uniform distribution ranging from $-s_v$ to s_v produces on average shorter error RTs, because errors are more likely for a bias towards bound b_2 and hence reach the threshold faster. Physiologically, this variability can be explained by trial-to-trial differences in attention, premature sampling, or other variable perceptual biases.

The extended DDM version is hardly tractable analytically. RT distributions and ERs can only be obtained numerically. Typically, the DDM is fitted to a particular set of behavioral data by minimizing the deviation between a simulated data set and the experimental data set (Vandekerckhove and Tuerlinckx, 2007).

A simplified deterministic version of the DDM has also been proposed by Reddi and Carpenter (2000). Their 'LATER' model produces variability in reaction times by varying the drift rate across trials, without any within-trial noise. For the random dot motion task, where the stimulus itself is explicitly stochastic, this might not seem a plausible model. Yet, trial-to-trial variability alone can be sufficient to fit the behavioral data of certain two-alternative forced choice tasks (Brown and Heathcote, 2008).

Nevertheless, the DDM has the advantage that reaction times and accuracy are directly related through the decision threshold. In experiments, more pressure for speed typically leads to faster RTs and lower accuracy. This negative correlation is known as the 'speed-accuracy tradeoff' (SAT). By adjusting the decision bounds, the DDM can reproduce the negative SAT correlation (Palmer et al., 2005). In the absence of noise, as for the LATER model, the threshold influences reaction times, but has no effect on accuracy.

The reverse conclusion of the SAT is that perfect accuracy could be achieved with unlimited processing time. Accordingly, the DDM implements perfect integration of the evidence difference in the sense that no information is 'forgotten' or overly emphasized. Nonetheless, participants' accuracy as a function of RT often reaches an asymptote, especially in difficult tasks, and the assumption of perfect integration might not hold true for neural systems. These discrepancies can be solved by introducing a 'leakage' term λv to the drift rate in Eq. (8). Information loss over time is modeled according to:

$$dv(t) = (\lambda v + \mu)dt + \sigma dW, \quad v(0) = 0, \tag{11}$$

with $\lambda < 0$, corresponding to a stable Ornstein–Uhlenbeck (O–U) process (Busemeyer and Townsend, 1993). This can be pictured by a diffusion of v in a curved instead of a flat landscape, where v approaches a stable fixed point $v^* = -\mu/\lambda$, where $dv = 0$. In the opposite case of $\lambda > 0$ both the mean and variance of v grow exponentially, as v is repelled from the now unstable fixed point (an unstable O–U process).

2.1.3. The race model

While in the DDM a single integrator accumulates the evidence difference between two alternatives, in the race model (Vickers, 1970, 1979) separate integrators v_1, v_2 are used for each alternative:

$$\begin{aligned} dv_1 &= I_1 dt + \sigma dW_1 \\ dv_2 &= I_2 dt + \sigma dW_2 \end{aligned}, \quad v_1(0) = v_2(0) = 0, \tag{12}$$

where I_1 and I_2 denote the average incoming evidence, respectively. As for the DDM, white noise is sampled from a normal distribution, $N(0, \sigma^2 dt)$. In the free-response mode, a decision is made as soon as one of the two integrators exceeds a threshold B. (We again assume equal prior probabilities. Therefore both integrators have the same decision bound B.) The two integrators thus perform a 'race to threshold' against each other.

Except for the case of perfectly anticorrelated noise in the two accumulators, the race model is not equivalent to the DDM and thus is not optimal (Bogacz et al., 2006). Nevertheless, in contrast to the DDM, the race model can easily be extended to multiple-choice decisions, simply by adding more integrators.

Moreover, in the race model, v can be interpreted as the population activity of two neural pools, receiving inputs from distinct populations of up-stream sensory neurons. For the DDM, however, it remains unclear how the difference in evidence might be computed physiologically.

This problem has been addressed in subsequent 'connectionist' models of two-alternative forced choice decision-making (Fig. 3). These abstract neural network models implement the diffusion process with inhibition between two integrators and will be reviewed in the following section.

2.2. Biologically motivated rate models

As we have seen, the drift–diffusion model is an intuitively appealing model of two-alternative forced choice decision-making and, moreover, achieves optimality according to the sequential probability ratio test. But, is there a way to implement this drift–diffusion concept in a physiologically plausible manner?

Several models have been proposed, which effectively compute evidence subtraction with different inhibitory mechanisms (Fig. 3). Although these models all exhibit two separate integrators just like the race model, dynamically they are more closely related to the DDM (Bogacz et al., 2006).

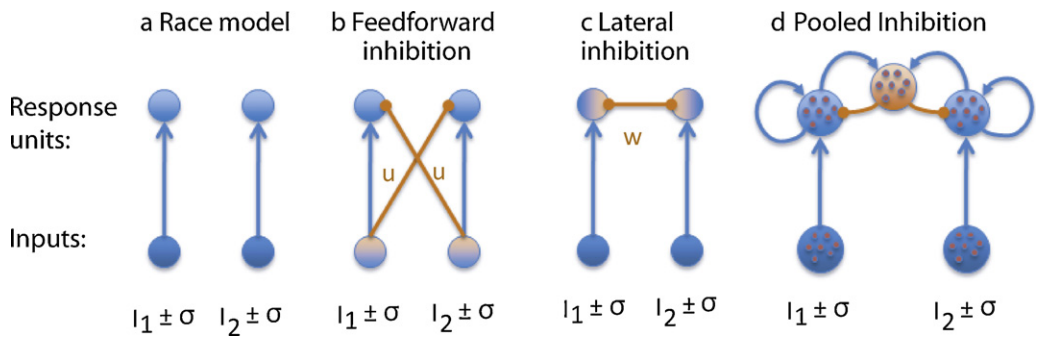

Fig. 3. Two-alternative forced choice (2AFC) decision models with two integrators. The blue color denotes excitatory connections and populations, and orange inhibitory. Adapted from Bogacz et al. (2006).

2.2.1. Feedforward inhibition (FFI)

Mazurek et al. (2003) proposed a feedforward inhibition model directly motivated by neural activity from MT and LIP during the random dot motion task (Britten et al., 1993; Shadlen and Newsome, 2001). The model instantiates the hypothesis that two populations of movement-sensitive MT neurons provide evidence in favor or against the opposed motion directions, which is then integrated in area LIP (Fig. 3b). A decision is made when the pooled activity reaches a threshold. The two LIP populations thus correspond to the response units, which receive excitatory inputs from one population of MT neurons, and inhibitory inputs from the other, according to:

$$\begin{aligned} dv_1 &= I_1 dt + \sigma dW_1 - u(I_2 dt + \sigma dW_2) \\ dv_2 &= I_2 dt + \sigma dW_2 - u(I_1 dt + \sigma dW_1) \end{aligned}, \quad v_1(0) = v_2(0) = 0, \qquad (13)$$

where u denotes the inhibitory FF connection weight. In this simple version of the FFI model, integration is perfect without leakage.

Just as with the basic DDM model, the FFI model cannot account for the slower error RTs found experimentally in the RDM task without further extensions. Ditterich (2006a,b) consequently suggested that a time-variant version of the FFI, including for example a within-trial increase of the input gain, could account quite well quantitatively for both the behavioral and neural data of Roitman and Shadlen (2002). An additional leak term further improved the fit to the neural data.

In addition, Niwa and Ditterich (2008) successfully applied a 3-alternative version of the FFI model to their random dot motion experiment with three possible directions of motion. Theoretically, the FFI model can be extended to any number of choice-alternatives, if the inhibitory weights are adapted accordingly. This assumption might be plausible in an experiment where trials with different numbers of alternatives are present in separate blocks of trials. By contrast, if the different trials are randomly interleaved, as in (Churchland et al., 2008), sufficient neural plasticity to adapt the connection weights is hardly practical in the short time between trials.

2.2.2. Lateral inhibition and the leaky competing accumulator

Apart from feedforward inhibition, also lateral inhibition between the two integrators, here the two LIP populations, could effectively support a diffusion process through competition within LIP (Fig. 3c). This picture is consistent with the established assessment that long range cortical connections are mostly excitatory, and inhibition thus acts locally within a cortical column (e.g. Lund et al., 2003). Usher and McClelland (2001) proposed a model called the 'leaky competing accumulator' (LCA) model, which in its simplest form can be written as:

$$\begin{aligned} dv_1 &= (-kv_1 - wv_2 + I_1)dt + \sigma dW_1 \\ dv_2 &= (-kv_2 - wv_1 + I_2)dt + \sigma dW_2 \end{aligned}, \quad v_1(0) = v_2(0) = 0 \qquad (14)$$

Here, $k > 0$ is the decay rate, or leak, equivalent to $\lambda < 0$ in the Ornstein–Uhlenbeck model (Eq. (11)), and w denotes the inhibitory connection strength between the integrator units. Usher and McClelland (2001) further incorporated nonlinearity in their model in the form of a threshold-linear activation function, which prevents the firing rates from dropping below zero. The LCA model accounts for correct and error RT distributions without the need of trial-to-trial, or within-trial, variability.

In addition, multi-choice decision-making was also addressed with the LCA model (Usher and McClelland, 2001), and the model captured the positive log-linear relation between the number of alternatives and reaction time, known as Hick's law (Hick, 1952).

To summarize, Bogacz et al. (2006) demonstrated that for a particular parameter range, namely large and balanced leakage and inhibition, the dynamics of the LCA and the FFI model approximate a one-dimensional linear diffusion process equivalent to the DDM with perfect integration. Moreover, the LCA and consequently also the basic DDM are approximated by a more physiologically plausible connectionist model with pooled inhibition and recurrent excitation, if self-excitation balances the decay and inhibition is strong (Fig. 3d). Yet, whether the brain actually works in a parameter regime of perfect integration has recently been called into question by a random dot motion study with time-varying evidence (Huk and Shadlen, 2005; Wong and Huk, 2008). For an accurate description of real neural dynamics, nonlinear attractor states that arise from strong recurrent connections are we believe important.

2.3. Attractor models

Neural networks with interconnected neurons, such as the pooled inhibition model displayed in Fig. 3d, form nonlinear dynamical systems, whose long-term behavior is determined by 'fixed points', or 'steady states'. These fixed points can be attractive, or repellent, and their existence depends on different parameters of the system, for example the recurrent connection weights, or the inputs to the neural network. A useful analogy of the system's trajectory through state space is a particle that moves on an energy landscape with valleys, corresponding to attractors or stable fixed points, and hills, corresponding to unstable, repellent fixed points (Fig. 4).

Decision-making can be implemented in an attractor network with at least two stable fixed points ('bistability'), representing the two choice-alternatives. The decision process is then regarded as the transition from an initial starting point towards one of the two attractors. Once such a decision-attractor is reached, this state will persist except for high levels of noise or perturbations, and can thus produce persistent neural activity. This is an advantage, as the decision state can be maintained in the attractor short-term

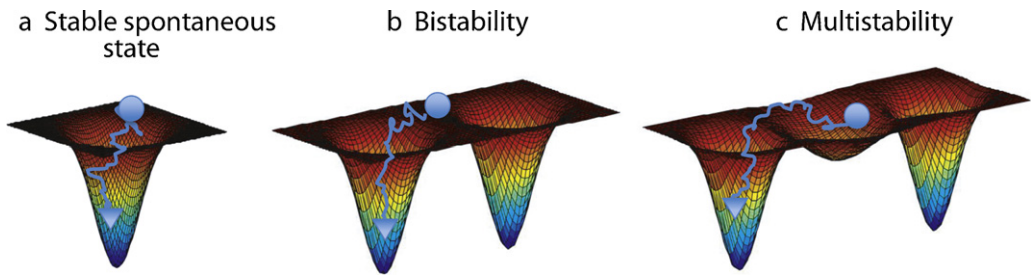

Fig. 4. Schematic presentation of possible attractor configurations in the attractor network of binary decision-making. Stable fixed points, or attractors, are illustrated as valleys in the 'energy landscape', and the hills correspond to unstable, repellent fixed points. Depending on the network parameters and inputs, (a) one, (b) two, or (c) three attractors can be simultaneously stable. (b) In the bistable case, the two decision attractors, but not the spontaneous state, are stable. The system (blue dot) will eventually move from the unstable spontaneous state into one of the two decision states. (c) In the multistable case, the spontaneous state as well as the two decision attractors are stable states. Transitions from the spontaneous state to a decision-attractor can happen due to noise fluctuations that are large enough to drive the system from the spontaneous state across the 'hill', or unstable fixed point, into the 'basin of attraction' of one of the decision-attractors.

memory that may provide the goal for an action while the action is being produced (Rolls, 2008).

In the pooled inhibition model, decision-attractors emerge through strong recurrent connections, which form positive feedback loops of activity in the excitatory neural populations. Runaway activity is averted through negative feedback from the inhibitory neural population.

The biophysically realistic attractor model that forms the basis of our theoretical work is a spiking-neuron implementation of the pooled inhibition model (Wang, 2002). Due to the nonlinear response properties of the spiking neurons, the full model can sustain a state of low spontaneous activity for a physiological range of background activity (Amit and Brunel, 1997). Therefore, depending on the magnitude of the (e.g. sensory) inputs and the strength of the recurrent connections, the model can work in three different dynamical regimes: (1) only the spontaneous state of low firing is stable (Fig. 4a), (2) a bistable regime of categorical decision-making (Fig. 4b), and (3) a 'multistable regime', where the spontaneous and the decision attractors are stable while the decision inputs are being applied (Fig. 4c). In the multistable regime, transitions from the spontaneous state to a decision-attractor can happen due to noise fluctuations that are large enough to drive the system across the 'hill', or unstable fixed point, into the 'basin of attraction' of one of the decision-attractors. This differs from the bistable scenario, in which the system will evolve towards one of the two choices even in the absence of fluctuations (Wong and Wang, 2006). In the multistable regime, fluctuations, perhaps noise-driven, are essential for decision making.

2.3.1. Biophysically realistic attractor model with spiking neurons

Although the connectionist models discussed above schematically describe neural processes in area MT and LIP, they still lack a direct connection between model variables and real neural parameters. This is different in the biophysically detailed implementation of the pooled inhibition model proposed by Wang (2002), where single neurons are modeled as leaky integrate-and-fire (LIF) neurons (Tuckwell, 1988) with conductance-based synaptic responses, described by realistic synaptic kinetics. The model was initially developed to account for (persistent) neural activity of PFC neurons during working memory tasks (Brunel and Wang, 2001). Its application to decision-making was inspired by the experimental observation that neurons, which exhibit ramping activity, characteristically show persistent neural firing in delayed memory or decision tasks (Gnadt and Andersen, 1988; Shadlen and Newsome, 2001). Wang (2002) successfully applied the spiking-neuron attractor model to behavioral and neurophysiological decision-making data measured from primates performing a binary random dot motion discrimination task (Roitman and Shadlen, 2002). In this context, the attractor network can be viewed as the representation of a local microcircuit in area LIP (or the posterior parietal cortex in general).

Physiological neuronal firing rates are obtained by averaging over the simulated action potentials, or output 'spikes', of distinct neural populations or 'pools' of LIF neurons in the network. Each LIF unit is characterized by its subthreshold membrane potential

$$C_{\mathrm{m}}\frac{dV(t)}{dt} = -g_{\mathrm{m}}(V(t) - V_{\mathrm{L}}) - I_{\mathrm{syn}}(t), \tag{15}$$

with resting potential V_{L}, membrane capacitance C_{m}, membrane leak conductance g_{m}. I_{syn} is the total synaptic current flowing into the cell. When the membrane potential V of a LIF neuron reaches the firing threshold V_{th}, it is reset to V_{reset} and a spike is emitted to all connected neurons with a subsequent absolute refractory period of τ_{ref}. Accordingly, LIF neurons do not explicitly model action potentials, but give a realistic account of the subthreshold membrane potential. Excitatory synaptic currents between LIF neurons are mediated by fast AMPA and slow NMDA glutamate receptors, and inhibitory synaptic currents by $\mathrm{GABA_A}$ receptors. The total synaptic current is given by the sum:

$$I_{\mathrm{syn}}(t) = I_{\mathrm{AMPA,rec}}(t) + I_{\mathrm{NMDA,rec}}(t) + I_{\mathrm{GABA}}(t) + I_{\mathrm{AMPA,ext}}(t). \tag{16}$$

The attractor network is organized into separate populations of LIF neurons, termed 'pools', which share common inputs and connectivities (Fig. 5a). As in the connectionist version (Fig. 3d), the spiking neuron model contains one homogenous pool of inhibitory neurons, globally connected to all neurons in the network. The two integrator units are implemented by two 'selective pools' of excitatory neurons, which are thought to respond *selectively* to one of the two possible directions of coherent motion and, hence, reflect the possible choice-alternatives in the random dot motion task (S1, S2 in Fig. 5a). Moreover, a third excitatory pool of 'nonselective' neurons represents the activity of surrounding LIP neurons that are not selective to either direction. More generally, there is one selective (or decision) pool for each possible choice alternative.

All neurons in the network receive an external background input simulated by uncorrelated, stochastic Poisson spike trains applied independently to the individual neurons. This background activity determines the spontaneous firing rate of all network neurons and is the main source of noise in the network together with finite-size effects due to the limited number of neurons (Rolls and Deco, 2010). Here, the term 'finite-size noise' describes the variance in the recurrent inputs, which tends to zero when the size of the network is increased.

In addition to the background input, the selective pools further receive additional time-dependent external inputs, corresponding to the decision evidence, such as, for instance, the decision sensory

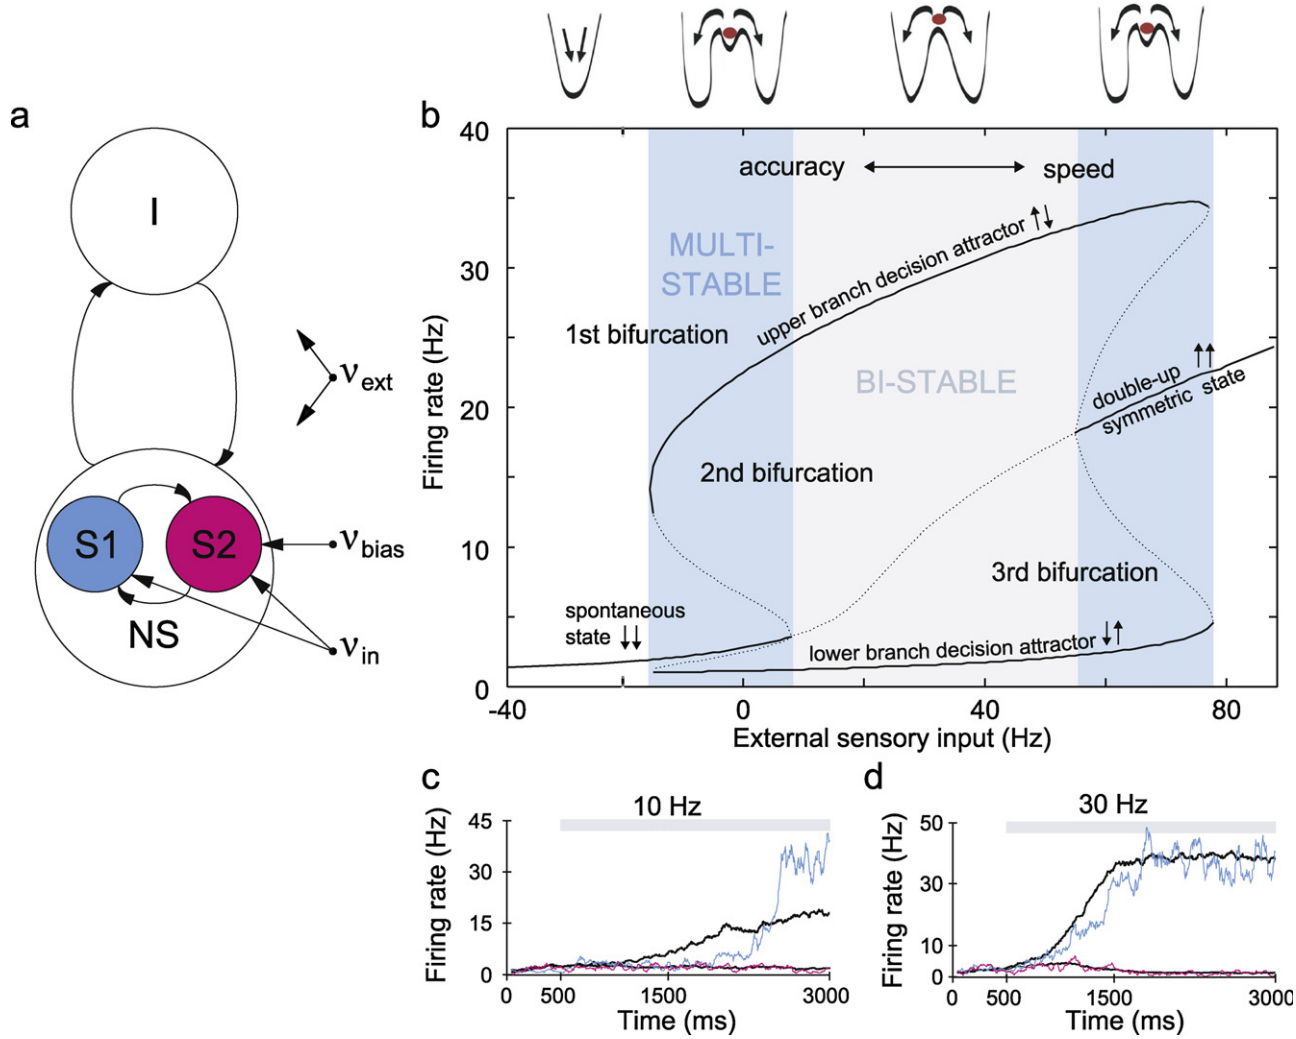

Fig. 5. Biophysically realistic attractor network of decision-making. (a) Schematic representation of the network. The excitatory neurons are organized into three pools: the nonselective neurons (NS) and the two selective pools (S1, S2) that receive the inputs encoding each stimulus (with rate ν_{in}). An additional bias (ν_{bias}) can be applied to one of the two selective pools. All neurons also receive an input (ν_{ext}) that simulates the spontaneous activity in the surrounding cerebral cortex. (b) Stable (solid lines) and unstable (dotted lines) fixed points depend on the external sensory input. They were calculated with the mean-field approximation of the network (Brunel and Wang, 2001). The abscissa shows the firing rates of the equal sensory inputs to S1 and S2, i.e. ν_{in}. The ordinate shows the mean firing rate of the neurons in the selective populations in the different states of the network. (c and d) Single trial (colored traces) and mean (black, averaged over 20 trials) temporal evolution of the firing rate of the selective pools for different inputs, denoted by the gray bars ($\nu_{bias} = 0$). (c) Noise-induced transition from the spontaneous state to the decision state (with low inputs, in the multistable regime). (d) Input-driven transition (in the bistable regime). The spiking simulations as well as the mean-field approximation in (b) were performed with a synaptic strength of $\omega_+ = 1.68$ within selective populations and with equal sensory inputs to both selective populations ($\nu_{bias} = 0$); all other parameters were taken from Wang (2002). Adapted with permission from Masquelier et al. (2011).

stimuli (i.e. coherent motion) in the random dot motion experiment.

As explained above for the general case, decision attractors emerge in the network due to the strong recurrent connectivity ω_+ of neurons *within* one selective pool, while the connections *between* the two selective pools are weaker than average $\omega_- < 1$. This is consistent with an associative, Hebbian, synaptic learning rule, as it is assumed that the activity of neurons that are selective for the same feature has been correlated in the past, while neurons encoding opposite directions instead fired in an anticorrelated manner (Rolls, 2008; Rolls and Deco, 2010). Fig. 5b displays an example of a typical attractor landscape for strong ω_+ as a function of increasing external inputs applied equally to both selective pools.

Without any external sensory inputs to the selective pools (0 Hz), the system will naturally rest in its spontaneous state with similarly low firing rates in all excitatory neural pools.

If a sensory stimulus is applied to the model, which increases the external inputs to the selective pools sufficiently (>10 Hz), the spontaneous state becomes unstable in a 'subcritical pitchfork bifurcation' leading to bistability between the decision attractors (the gray area in Fig. 5b). The network then operates in a region of categorical decision-making, where one selective pool will settle at the upper decision branch with high firing rate ('winner'), while the other will decay to the lower branch ('loser'). In this case, the transition from the spontaneous state to the decision state is 'input-driven' and can be gradual, in the order of several hundred milliseconds, even on single trials (Fig. 5d). These gradual transitions between attractor states, corresponding to the decision process, are a distinguishing feature of the biophysically realistic attractor model, and rely on the slow kinetics of the NMDA receptors ($\tau_{NMDA,decay} = 100$ ms). Consequently, the network's behavior is not just dominated by its steady states, it also exhibits prolonged responses to momentary sensory inputs, with a characteristic time constant of up to a second, during which the model effectively integrates the incoming inputs (Wang, 2008).

As depicted schematically in Fig. 4c, in the multistable regimes decision-making is also possible. There, the transition from the spontaneous or symmetric state to a decision state is induced by

noise fluctuations, and can be rather sharp in a single trial (Fig. 5c). Nevertheless, the trial-averaged activity builds up slowly, as observed experimentally in decision-related neurons. This ramping activity is, however, obtained in a conceptually different way compared to the bistable regime: the gradual build-up is an artifact of averaging across abrupt transitions at random times (Marti et al., 2008).

For sufficiently strong connection weights ω_+, as in the example of Fig. 5, the network can exhibit persistent activity, meaning that the high firing rates of the 'winner' population can be sustained even after all external sensory inputs are switched off. This is because at 0 Hz, with only background activity, the system is already in the multistable regime, where the decision states are stable in addition to the spontaneous state. Under these conditions, the decision states would only destabilize if negative inputs were applied to the selective populations. Persistent activity is a characteristic feature of all biophysically realistic attractor models that we will describe, unless otherwise stated. Therefore, the bifurcation between the multistable regime at low inputs and the bistable regime is the first bifurcation with relevance to the decision-making process ('second bifurcation' in Fig. 5b).

In the vicinity of this 'second bifurcation', slow integration is enhanced above the intrinsically slow kinetics mediated by the NMDA receptors, as the effective time constant of the system increases even beyond that related to the NMDA receptors close to the third bifurcation (Roxin and Ledberg, 2008; Wong and Wang, 2006). Therefore, in this dynamical region performance is high and reaction times are rather long, because of long stimulus-integration times. For this reason, previous analyses of the binary attractor model particularly concentrated on the dynamics in the proximity of the second bifurcation, where the spontaneous state destabilizes (Marti et al., 2008; Wang, 2002; Wong and Wang, 2006).

Nevertheless, exploring the network dynamics at the other end of the bistable regime is also of interest. For sufficiently high external sensory inputs, the network again enters a multistable regime. Crossing the 'third bifurcation', a symmetric 'double-up' state becomes stable, where both selective pools fire with intermediate, elevated rates. This double-up symmetric state has recently been deployed to explain LIP responses to static visual stimuli, as for example the response targets in the random dot motion paradigm (Albantakis and Deco, 2009; Furman and Wang, 2008; Wong et al., 2007). Assuming high selective inputs with R-target onset, the high firing rates of LIP neurons prior to the motion stimulus (Churchland et al., 2008; Huk and Shadlen, 2005; Kiani and Shadlen, 2009; Roitman and Shadlen, 2002) can be reproduced with the attractor model in the double-up state. Moreover, changes of mind, and thus the possibility to reevaluate a first choice, emerge naturally in the attractor network close to the 'third bifurcation' (Albantakis and Deco, 2011).

Across the bistable regime, in between the two bifurcations (2nd and 3rd), higher external inputs to both selective populations lead to faster reaction times and less accuracy, congruent with a speed–accuracy tradeoff (Roxin and Ledberg, 2008; Wong and Wang, 2006). This dependency of decision-behavior on the amount of unbiased external sensory inputs to both selective populations explicitly arises from the nonlinearity of the attractor model. In linear models, such as the drift–diffusion model, changes in the common sensory evidence to both decision alternatives would not affect decision behavior. This is obvious in the case of the DDM, as the DDM characteristically accumulates only the difference between the evidence for the two alternatives. Indeed, input-dependent decision behavior might provide a means to distinguish between linear and nonlinear modeling approaches (Albantakis and Deco, 2011).

We further note that, although the decision-attractor would automatically provide an upper bound for the neural activity of the winning selective population, the model's decision is typically determined by a fixed firing rate threshold independent of the applied amount of sensory inputs, in line with neurophysiological evidence from LIP neurons. How this decision threshold is read out or adjusted by down-stream areas is not explicitly included in the attractor model. Nevertheless, possible extensions have been suggested, which implement the decision threshold involving cortico-collicular and cortico-basal ganglia circuits (Bogacz and Gurney, 2007; Lo and Wang, 2006).

Taken together, the characteristic features of the biophysically realistic spiking-neuron attractor model are:

- strong recurrent connections within the selective neural populations, which generate attractor states,
- global feedback inhibition enabling winner-take-all competition,
- stochasticity because of finite-size effects and random Poisson inputs to the network,
- a long synaptic time-constant (NMDA) facilitating the integration of incoming neural activity.

2.3.2. Model reductions

Simulating populations of individual and realistic neurons as described above is necessary to simulate realistic neuronal dynamics, physiological responses and behavior. Nevertheless, to gain an analytical understanding of the population dynamics, a reduced, mathematically tractable description can yield deeper insights into the essential, collective model behavior (Rolls and Deco, 2010).

2.3.2.1. Mean-field approximation. Taking a mean-field approach, Brunel and Wang (2001) considerably reduced the state variables of the network by replacing the different populations of spiking neurons with an approximation of their average population activity. Because of the recurrent connections in the network, the population firing rates have to be determined self-consistently based on the input currents to the neural pools, which in turn depend on the firing rates. Equalizing the pre- and postsynaptic activity, the possible fixed points or steady states of the population firing rates can be obtained.

Several approximations have to be assumed in order to arrive at a closed system of one nonlinear equation for each population in the network. First, postulating that individual neurons fire spikes independently, according to a stationary Poisson process, the net input is treated as a Gaussian random process. This assumption generally holds in the limit of infinitely large networks, where each neuron receives a large number of presynaptic currents, which each deflect the membrane potential only minimally compared to the voltage distance between the resting and threshold potential. Second, only fast AMPA-mediated external inputs are assumed to contribute to the fluctuations in the synaptic current. Fluctuations in the NMDA and GABA currents are neglected as they are supposed to filter out, due to the longer synaptic time constants.

Finding a self-consistent solution for the population rates is further complicated by the nonlinear properties of the NMDA receptors. NMDA saturation is approximated by calculating the average NMDA gating variable as a nonlinear function of the presynaptic rate. In addition, the voltage dependence of the NMDA conductance is linearized around the average neural potential.

For the detailed mathematical derivation we refer to the original publication (Brunel and Wang, 2001), to Renart et al. (2003), and to Rolls and Deco (2010). Solving the mean-field equations is computationally much less intense than running simulations with the full spiking network. The mean-field analysis thus allows calculation of the steady state firing rates of the attractor model for a wide range of parameters. This makes it feasible to scan the parameter space in order to find a parameter set matching experimental findings.

In sum, by solving the mean-field equations for a set of initial conditions (here the initial firing rates of each neural population) one obtains the approximated average firing rate of each pool when the system has settled into a stationary state after the period of dynamical transients. The mean-field reduction, however, does not provide an accurate description of the temporal dynamics.

2.3.2.2. Two-dimensional reduction. Based on the mean-field approach, Wong and Wang (2006) further reduced the biophysically realistic spiking-neuron model to a two-variable system. In particular, they fitted a simplified input–output function to the complex first-passage time formula used in the mean-field to describe the output firing rates as a function of the mean synaptic inputs. With the assumption that the network dynamics are dominated by the slow NMDA receptors, they further set the firing rate of the nonselective pool to a constant mean firing rate (2 Hz) and linearized the input–output relation of the inhibitory neurons. Thereby, inhibition could be incorporated as mutual negative inputs between the selective pools. The nonselective neurons and inhibitory interneurons could thus be eliminated, leaving two neural units with self-excitation and effective mutual inhibition.

The two-dimensional reduction is particularly useful to perform phase-plane analyses in order to elucidate the different dynamical regimes of the network. To simulate the noisy temporal evolution of the spiking neural network, Wong and Wang (2006) explicitly added a noise term to the external inputs, described by an Ornstein–Uhlenbeck process (white noise filtered by a short, AMPA receptor synaptic time constant). The two-dimensional reduction can thus be viewed as a closely related connectionist version of the full spiking model. In this way, it can also account for decision-related behavior and neural activity, albeit without explicit analogy to real neural parameters (Wong et al., 2007).

2.3.2.3. Nonlinear diffusion. Instead of a two-component system of rate equations as in (Wong and Wang, 2006), Roxin and Ledberg (2008) derived a one-dimensional *nonlinear* diffusion equation to describe the asymptotic behavior of winner-take-all models in the proximity of the bifurcation to bistability, where the spontaneous state destabilizes. Their reduction is universally valid for all winner-take-all models, but also makes it possible to relate the variables of the nonlinear diffusion process to those of the full spiking-neuron model, and thus to neurobiologically meaningful quantities. A particularly relevant prediction based on this one-dimensional reduction is that the speed-accuracy tradeoff can be implemented by changes in the common inputs to both selective neural populations, instead of, or in addition to, an adaptation of the decision-threshold.

2.4. Distinguishing model approaches

As we have seen, models on the accumulation of noisy evidence, as for instance in the random dot motion paradigm, come with a huge variety. Although they differ in fundamental features, such as network structure and connectivity, in practice, it may be very difficult to distinguish between them on the basis of just behavioral data or mean firing rates (Bogacz et al., 2006; Ditterich, 2010).

Ratcliff and Smith (2004) evaluated four types of sequential sampling models and the leaky competing accumulator (LCA) model against three sets of psychophysical two-alternative forced-choice experiments. In particular they compared the three models described above, the drift–diffusion model, the Ornstein–Uhlenbeck model (Eq. (11)), and the race model, and a so-called 'Poisson counter' model (Townsend and Ashby, 1983), all with trial-to-trial variability in drift, starting point, and non-decision time. (The Poisson counter model resembles the race model, with the difference that evidence is counted in discrete units, delivered at random times, with exponentially distributed intervals. Therefore, it can be interpreted as an independent accumulation of two spike trains.)

Of all models considered, only the Poisson counter model failed to match the empirical data, and the faster decision times on error trials still proved problematic for the race model. The Poisson counter model also proved inferior to the DDM when compared to the neural activity of superior colliculus build-up neurons from macaque monkeys performing a two-alternative forced-choice task (Ratcliff et al., 2003). The activity pattern predicted by the DDM, however, resembled the observed neuronal firing rates, suggesting that build-up cells in the superior colliculus might participate in a diffusion-like decision process.

Because of the mutual similarity between the models (Bogacz et al., 2006; Ratcliff and Smith, 2004), finding new analytical methods and well designed experiments to distinguish the different approaches is a major future challenge in the field of perceptual decision-making.

One approach along that line was conducted by Huk and Shadlen (2005). By adding brief motion pulses to a standard RDM stimulus, they first of all provided strong physiological support for a temporal integration in LIP. However, their findings revealed a departure from *perfect* integration, as the effect of the motion pulse decreased with its onset time. Later motion pulses thus influenced behavior and neural activity less than earlier motion pulses. Neither a perfect DDM, nor leaky integrators could reproduce this experimental finding, while the time-varying dynamics of the attractor model explained both the behavioral and the neuronal data (Wong and Huk, 2008; Wong et al., 2007). Still, time-varying effects such as decreasing decision bounds or an 'urgency' signal might produce decreased sensitivity to later perturbations also in the DDM and LCA.

Recently, multiple-choice decision tasks have received increasing attention in the context of distinctions between models (Churchland et al., 2011; Ditterich, 2010; Leite and Ratcliff, 2010; Purcell et al., 2010).

Analyzing higher-order statistical properties of neurophysiological data from their 2- and 4-alternative random dot motion task, Churchland et al. (2008, 2011) were able to distinguish between models categorized by their different sources of variability. Models with just one source of variability, such as the LATER model (2.1.2) and a model by Cisek et al. (2009) with fixed slope, but a random distribution of firing rates at each time-step, failed to account for the higher-order measures, although they agreed with behavior and mean firing rates. On the other hand, all different model implementations of a stochastic accumulation-to-threshold tested in Churchland et al.'s (2011) study could account for variance and within-trial correlations, in addition to behavioral data and first-order firing rates. In particular, the tested models included the drift–diffusion model (Ratcliff and Rouder, 1998), a model based on probabilistic population codes (Beck et al., 2008), and the reduced version of the attractor model by Wong et al. (2007).

Based on human behavioral data from a random dot motion task with three alternatives and three motion components, Ditterich (2010) aimed to distinguish more detailed aspects of conceptual accumulation-to-bound models with regard to their goodness of fit and their neurophysiological predictions. Perfect integrators were compared to leaky, saturating integrators, with either feedback or feedforward inhibition. As we have seen, in the case of two alternatives, most of the discussed models proved equivalent to the DDM for certain parameter ranges (Bogacz et al., 2006). Therefore, it might not be too surprising that none of the models could be excluded based only on the fits to behavioral data of a 3-alternative random dot motion task (Niwa and Ditterich,

2008). Yet, the models differ substantially in their neurophysiological predictions on how the integrator states should evolve over time (see Table 2 in Ditterich, 2010). Neuronal recordings from monkeys performing the same task will hopefully soon provide clarification. Moreover, feedforward and feedback inhibition respectively suggest either negative or positive correlation between the integrator units, which might be tested with multi-electrode recordings. Finally, in the case of an equal amount of motion coherence in all three directions, Niwa and Ditterich (2008) measured faster mean reaction times for higher coherence levels. While models with feedforward inhibition require a scaling of the variance of the sensory signals in order to account for this effect, conceptual models with feedback inhibition could explain the result just with a change of the mean input (Ditterich, 2010).

Considering the evidence presented in this section, so far two types of decision-making models have proven particularly successful. On the one hand, the extended drift–diffusion model and its connectionist implementations account for a vast range of behavioral data. They also conceptually represent neural activity during the decision-making period. On the other hand, the physiologically detailed attractor model (Rolls and Deco, 2010; Wang, 2002, 2008) and its reductions (Wong and Wang, 2006), which mimic real neuronal dynamics, accurately simulate behavioral data and LIP activity during the random dot motion task, and in other decision-making tasks including vibrotactile decision-making in areas such as the ventral premotor cortex (Deco and Rolls, 2006), and value-based decision-making in the medial prefrontal cortex (Grabenhorst and Rolls, 2011; Rolls et al., 2010a,b). Moreover, they account for persistent activity, and the nonlinear, time-dependent, effects of events that interrupt the decision-making process (Wong and Huk, 2008; Wong et al., 2007).

With respect to the biologically plausible attractor decision-making networks, the dynamics relevant to the decision-making depend on the stability of the spontaneous firing activity state (i.e., the low firing rate state in which no decision has yet been made). If, once the second stimulus is presented, the spontaneous state destabilizes, then the dynamics rapidly evolves towards one of the two decision states (Wong and Wang, 2006). This corresponds to a so called 'ballistic scenario' consistent with a linear diffusion model. On the other hand, if the spontaneous state does not lose stability but is instead bistable with the decision states, hence leading to multistability between three possible fixed points (with a multistable region between the first and second bifurcation), then, only a sufficiently strong perturbation would drive the system from the stable spontaneous state to one of the two decision states (Deco and Rolls, 2006; Rolls and Deco, 2010). This differs from the earlier ballistic scenario, in which the system will evolve towards one of the two choices even in the absence of fluctuations (Wong and Wang, 2006). Thus, in the multistable regime, fluctuations, perhaps noise-driven, are essential for decision making. Computational analysis of this scenario showed that behavior consistent with Weber's law for decision making (that the difference between the decision cues $\Delta\lambda$/the intensity of the decision cues λ is a constant) is more consistent with a fluctuation-driven scenario in a multistable regime (Deco and Rolls, 2006; Deco et al., 2007). In particular, experimental behavioral evidence showing a Weber's law like behavior for the decision-making involved in the discrimination of two vibrotactile stimuli in humans suggests that the neurodynamical mechanisms and computational principle underlying this process are consistent with a fluctuation-driven scenario in a multistable regime (Deco et al., 2007). In practice, the system works at the bifurcation between the two scenarios (Roxin and Ledberg, 2008), and there is a nonlinear modification of the linear diffusion model (at the exact point of the bifurcation) towards a concave or convex nonlinear landscape in the direction of one or the other scenarios.

Overall, in this section, we have seen that the biophysically realistic attractor network model provides a theoretical foundation for understanding many experimental observations from neurophysiological single neuron activity to human and animal behavior during decision-making. In the following we will describe developments in this approach that provide an understanding of many further aspects of decision-making, including how decision-making tasks with sequential presentation of the stimuli are implemented, how confidence in a decision before any feedback is received is an emergent property of the decision-making mechanism, how decisions can be predicted much better than chance before the evidence for the decision is even presented, how task difficulty and errors influence the operation of the mechanism, and how the brain implements decisions between stimuli with different reward value.

3. Synaptic facilitation as part of the mechanism for sequential decision-making, and for decision-making with postponed responses

3.1. Synaptic facilitation and sequential decision-making

Most models of decision-making deal with simultaneous presentation of the stimuli to be discriminated, λ_1 and λ_2. But in some tasks, such as the vibrotactile discrimination task, the stimuli are presented sequentially, and the decision is then about whether the second stimulus applied to the finger is higher or lower in vibrotactile frequency than the first stimulus (Romo et al., 2004). One way in which this has been modeled is to assume that there is a short-term memory of the first stimulus, and then to analyze the decision-making between the remembered first stimulus and the second stimulus, both applied simultaneously to the attractor decision-making network as λ_1 and λ_2 (Deco and Rolls, 2006). This approach also assumes that λ_1 and λ_2 are represented by different neurons, that is, that a place code (Rolls and Treves, 2011) is used.

By introducing synaptic facilitation into the attractor model, we have recently been able to develop a model of decision-making that can account for the memory of the first stimulus as well as the temporal evolution of how the stimuli are encoded in the neural firing rate. This means that the model for example captures when and how the vibrotactile flutter frequencies of both the first and the second stimulus are reflected in the neural firing rates (Deco et al., 2010).

The model accounts for the presence and usefulness of partial differential neurons of the type illustrated in Fig. 6 in the ventral premotor cortex, which is implicated in vibrotactile decision-making (Romo et al., 2004). In the ventral premotor cortex (VPC) some neurons reflect the decision-making itself, responding for example if the first vibrotactile stimulus (f_1) applied to the hand is higher in flutter frequency than the second (f_2), while other neurons respond to the decision $f_2 > f_1$ (see Fig. 2GHI of Romo et al., 2004). In addition to these neurons, the so called "partial differential" neurons reflect the memory of f_1. The activity of one of these neurons recorded in the ventral prefrontal cortex (Romo et al., 2004) is illustrated in Fig. 6. Partial differential neurons respond to f_1 during the presentation of f_1, do not respond in the first part of the delay period, then gradually ramp up towards the end of the delay period to a firing frequency that reflects f_1, and then during the decision period when f_2 is presented are influenced by f_2 (Romo et al., 2004). The responses of partial differential neurons may be related to the decision-making (Romo et al., 2004), for as shown in Fig. 6, if $f_1 > f_2$ the firing during f_2 is higher than when $f_1 < f_2$, for a given f_2.

To simulate the neuronal activity during the delay period the model utilizes synaptic facilitation, a process implemented using a

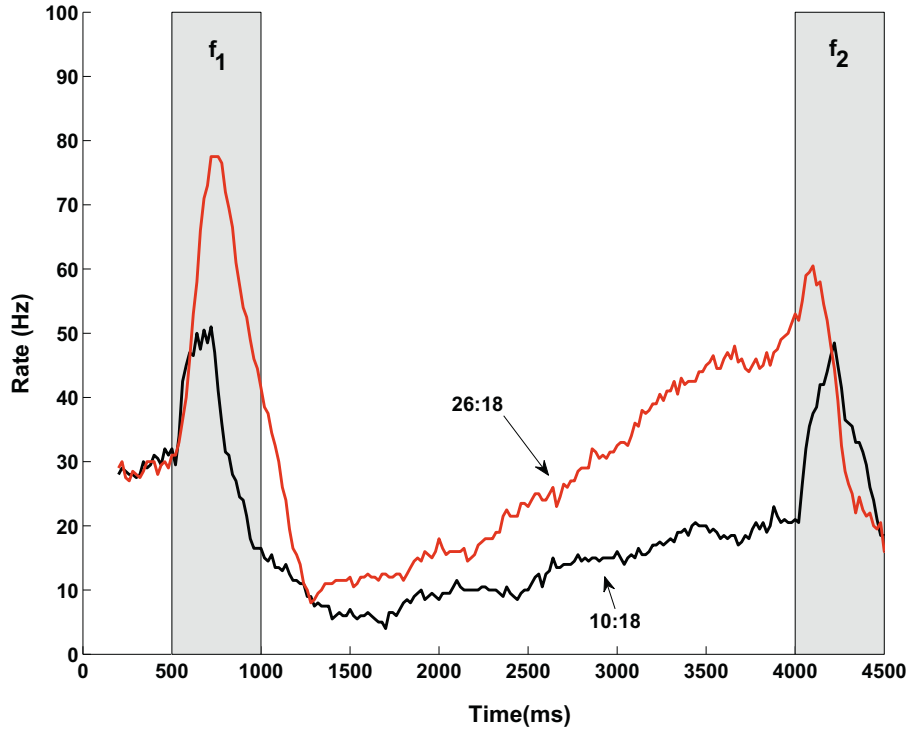

Fig. 6. Activity of a single neuron of the "partially differential" type recorded in the ventral premotor cortex (VPC) during the vibrotactile decision-making task, after Romo et al. (2004), as illustrated in Fig 2JLK of that paper. The f_1 period was from 500 to 1000 ms, there was then a delay period, and the f_2 period when f_2 was applied and the decision was made was from 4000 to 4500 ms. f_2 was in both cases 18 Hz. When f_1 was 26 Hz (red plot), the firing rate during f_1, and at the end of the delay period, and during the comparison period when f_2 was being applied was higher than when f_1 was 10 Hz (black plot). Thus the firing of this type of neuron during f_2 helps in the decision that $f_1 > f_2$ when f_1 is 26 Hz, and that $f_1 < f_2$ when f_1 is 10 Hz. Approximately 30 trials were used to generate these peristimulus time histograms for each pair for this single neuron. After Deco et al. (2010).

phenomenological model of calcium-mediated transmission (Mongillo et al., 2008). Synaptic facilitation is caused by the increased accumulation of residual calcium at the presynaptic terminals, which increases the probability of neurotransmitter release (Zucker and Regehr, 2002). This type of synaptic facilitation is common in higher cortical areas such as the prefrontal cortex (Hempel et al., 2000; Wang et al., 2006; Zucker and Regehr, 2002). The synaptic efficacy of the recurrent connections between excitatory neurons is modulated by the utilization parameter u (the fraction of resources used) reflecting the calcium level. When a spike reaches the presynaptic terminal, calcium influx in the presynaptic terminal causes an increase of u which increases the release probability of transmitter and thus the strength of that synapse. The decay time constant of the synaptic facilitation is regulated by a parameter τ_F which has been determined experimentally to be around 1–2 s (Wang et al., 2006), i.e. large enough to be able to sustain memory and allow comparison for short delays (3 s in our case).

Another difference between the attractor network with synaptic facilitation modeling the partial differential neurons and the basic attractor model described in Section 2.3.1 is that here the recurrent connections are not sufficient to maintain the attractor state without external sensory inputs. Consequently, the model responds to the first stimulus f_1, but then the firing decreases when f_1 is removed. Subsequently the firing gradually increases again due to a non-specific increase in excitation that is externally applied during the delay period after f_1 (Deco et al., 2010). This non-specific increase in excitation, in conjunction with the synaptic facilitation, which remains in the recurrent collateral connections of the network neurons that recently fired due to the application of f_1, is sufficient to reinstate these neurons into firing. This principle usefully models the delay-related firing shown in Fig. 6 (Deco et al., 2010). Then when f_2 is presented, the firing rate of the partial differential neurons reflects the vibrotactile frequencies of both f_1 and f_2 (Deco et al., 2010). Other neurons respond to f_2 and not to f_1. If both these neuron types (the partial differential and the f_2 responding neurons) form the λ_1 and λ_2 inputs to a separate attractor decision-making network, then the correct decision can be computed from λ_1 which reflects a combination of f_1 and f_2, and λ_2 which reflects only f_2. Provided that the inputs from these two types of neuron to the attractor decision-making network are scaled appropriately by synaptic strengths that are trained during a learning period, one population in the attractor decision-making network will respond when $f_1 > f_2$, and the other when $f_1 < f_2$. This is possible because the firing of the partial differential neurons reflects f_2 as well as f_1. The decision-making attractor network then compares the firing of the partial differential neurons with the firing of the f_2 neurons. The decision taken will be in one direction if f_1 is high in frequency relative to f_2, and in the other direction if f_1 is low in frequency relative to f_2. In effect, the partial differential neurons allow the value of f_2 to be taken into account in making the decision, so that the attractor network can reflect just the difference between f_1 and f_2 (Deco et al., 2010).

We argue that it is inevitable given the rate encoding with broad tuning of the neurons in VPC to the vibrotactile stimuli, that f_1 will be contaminated by f_2, and that a solution of the type we propose is needed. Indeed, we argue that the partial differential neurons of the type illustrated in Fig. 6 are fundamental to the solution of the decision-making problem by this brain region (Deco et al., 2010).

3.2. Synaptic facilitation, graded firing rates, and postponed decisions

In the preceding section we showed that a synaptic facilitation mechanism allows the attractor network to maintain the memory

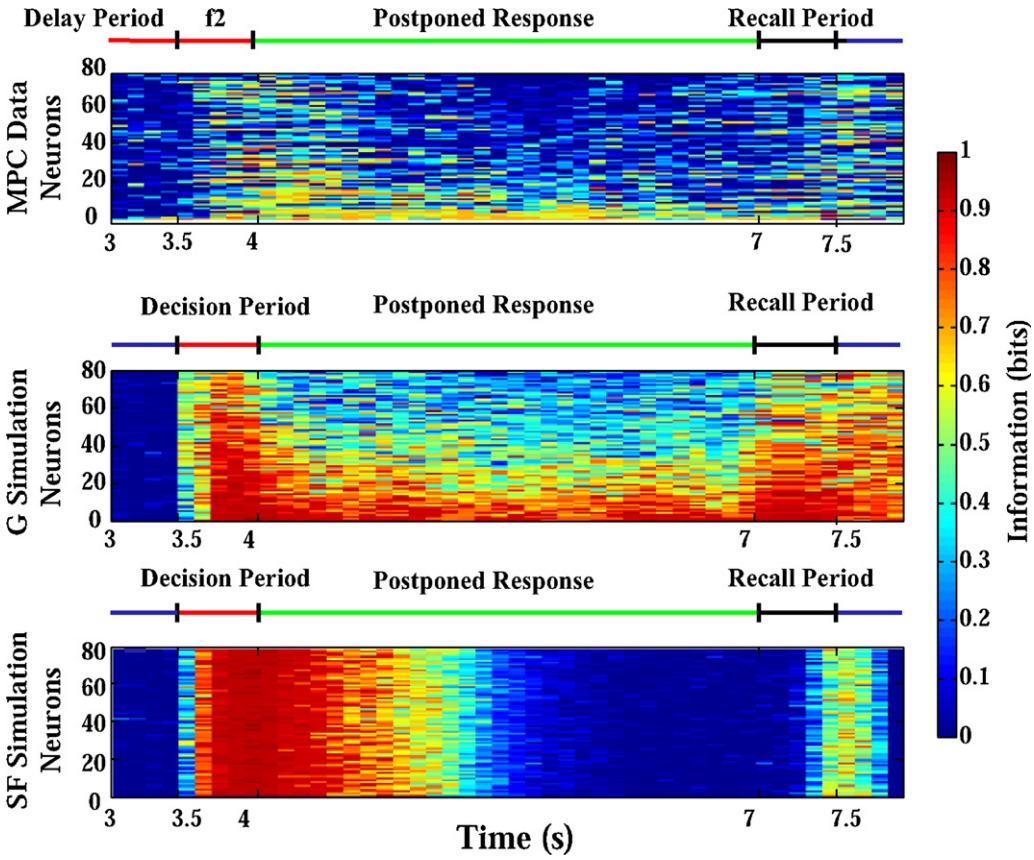

Fig. 7. Mutual information analyses as a function of time. (Top) Eighty neurons from the medial premotor cortex. Each row corresponds to a different neuron. In the experiment the f_2 stimuli have been applied from 3.5 to 4.0 s. The postponed response delay period is 4.0–7.0 s. The behavioral response can be started with a response signal given at 7.0 s. The mutual information shown was calculated between firing rate windows of 200-ms (sliding every 100 ms) and the response made by the monkey. (Middle) Eighty neurons from decision pool 1 of the graded firing rate simulation. Each row is a single neuron. The rows are sorted by the amount of information during the decision period, 3.5–4.0 s, which corresponds to the f_2 period for the MPC neurons. The postponed response period is 4–7 s. The non-selective external input is applied at t = 7.0–7.5 s, labeled recall period. The mutual information of the simulated firing rates was calculated between the firing rate in a 200-ms sliding window and the firing in the decision period. (Bottom) Synaptic Facilitation simulation (conventions as in Middle).
After Martinez-Garcia et al. (2011).

of f_1 during the delay between the stimuli f_1 and f_2 in a sequential decision-making task (Deco et al., 2010).

In addition to being necessary to perform sequential decision-making tasks, memory also plays a role in tasks where the decision is postponed for a period after the evidence has been provided (Lemus et al., 2007). Using an information theoretic approach to single-cell neurophysiological data recorded during a sequential tactile discrimination task with postponed decisions (Lemus et al., 2007), we were able to analyze at what time during the trial information about the forthcoming action becomes available. In particular, we found that information about the forthcoming action becomes available from the activity of neurons recorded in the medial premotor cortex in a sequential decision-making task after the second stimulus is applied that provides the information for a decision about whether the first or second stimulus was higher in vibrotactile frequency (Martinez-Garcia et al., 2011). After this "decision period" the information then decays in a 3 s delay period in which the decision is postponed and the neuronal activity declines before the behavioral response can be made. The information then increases again when the behavioral response is required in the "recall period" (top panel Fig. 7).

We were able to model the observed neuronal activity using an attractor decision-making network in which information reflecting the decision is maintained at a low level during the postponed response delay period, and is then selectively restored by a non-specific input when the response is required (Martinez-Garcia et al., 2011).

One mechanism for the short-term memory is synaptic facilitation which can implement a mechanism for postponed decisions that can be correct even when there is little neuronal firing during the delay period before the postponed decision, as illustrated in Fig. 7 (bottom panel).

A second mechanism to model the experimental data is graded firing rates by different neurons in the delay period (Fig. 7, middle panel). Importantly, also in this scenario, the information about the decision can be restored by the application of the non-specific input during the recall period. This is because the higher firing rate neurons of the "winner" population continue firing during the postponed response period, and help to activate the whole population and thus retrieve the decision information when the non-specific input is applied.

To conclude, either mechanism, synaptic faciliation and graded firing rates, can account for the experimentally observed activity of medial premotor cortex neurons and the evolution of the information about the decision during the decision-making and delay period before the response is made (Martinez-Garcia et al., 2011). The two mechanisms have in common the property that a non-specific input can restore the firing specifically in the correct attractor population of neurons. The relative contribution of the two mechanisms could be tested experimentally by increasing the delay between the decision formation and recall, as the synaptic memory decays away within a few seconds, and would predict that delays of more than a few seconds before the delayed response could occur would not be possible.

4. Confidence in decisions before the outcome is known

4.1. Decision confidence and discriminability

A fundamental issue in understanding how the brain takes decisions is to identify a neural signature for decision-making processes, and to explain how this signature arises as a result of the operation of dynamic cortical neural networks. Evidence from single neuron recordings in monkeys (Kim and Shadlen, 1999) shows that neuronal responses in a motion decision-making task occur earlier on easy vs difficult trials in a decision-related brain region, the dorsolateral prefrontal cortex. In the human dorsolateral prefrontal cortex, higher fMRI BOLD signals can be observed on easy trials vs difficult trials in a similar decision-making task (Heekeren et al., 2004). On the basis of these findings it has been suggested (Heekeren et al., 2004) that the dorsolateral prefrontal cortex implements a general mechanism for decision-making in the brain which is thought to involve the comparison of sensory input signals until a decision criterion is reached and a choice is made. Thus a neural signature of decision-making seems to involve higher levels of neural activity when an easy compared to a difficult decision is taken. However, the fact that some areas of the brain thought to be involved in decision-making show earlier or larger neuronal responses on easy trials (Kim and Shadlen, 1999), and reflect decision confidence (Kiani and Shadlen, 2009), does not by itself provide a theoretical understanding about what is special

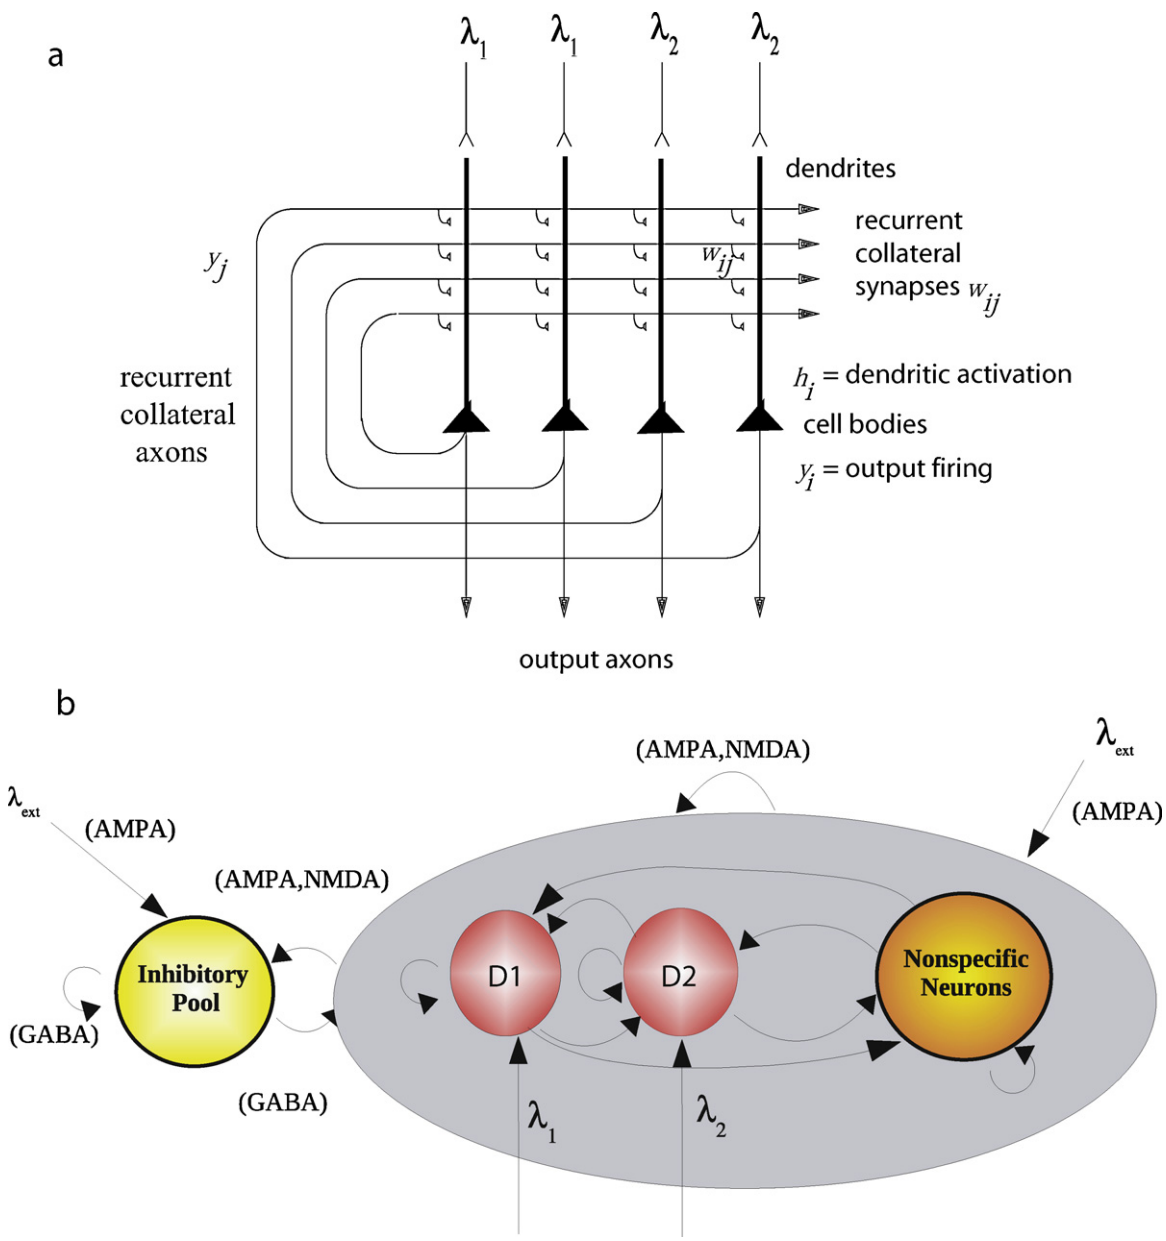

Fig. 8. (a) Attractor or autoassociation single network architecture for decision-making. The evidence for decision 1 is applied via the λ_1 inputs, and for decision 2 via the λ_2 inputs. The synaptic weights w_{ij} have been associatively modified during training in the presence of λ_1 and at a different time of λ_2. When λ_1 and λ_2 are applied, each attractor competes through the inhibitory interneurons (not shown), until one wins the competition, and the network falls into one of the high firing rate attractors that represents the decision. The noise in the network caused by the random spiking of the neurons means that on some trials, for given inputs, the neurons in the decision 1 (D1) attractor are more likely to win, and on other trials the neurons in the decision 2 (D2) attractor are more likely to win. This makes the decision-making probabilistic, for the noise influences when the system will jump out of the spontaneous firing stable (low energy) state, and whether it jumps into the high firing state for decision 1 (D1) or decision 2 (D2). (b) The architecture of the integrate-and-fire network used to model decision-making.
After Rolls et al. (2010a).

about this effect as an indicator of decision-making, or lead to an understanding of how confidence is computed.

We therefore performed a combined theoretical and experimental fMRI study to investigate the neural correlates of discriminability in decision-making (Rolls et al., 2010a). Further, a link from discriminability of the stimuli, measured by $\Delta\lambda$ (the difference between the decision stimuli λ_1 and λ_2), to confidence can be made, for it is well established that subjective decision confidence increases with discriminability (Jonsson et al., 2005; Vickers, 1979; Vickers and Packer, 1982). Consistently, in rats too the probability that a trial will be aborted reflecting low decision confidence also increases with $\Delta\lambda$ on error trials (Kepecs et al., 2008). The same investigation thus provided evidence on how confidence before the outcome is known is represented in the brain (Rolls et al., 2010a).

The architecture of the integrate-and-fire attractor network model of decision making that was used is shown in Fig. 8 (Rolls et al., 2010a). It was found that as $\Delta\lambda$ increased, the firing rates of the neurons in the winning attractor became a little higher. The reason for this was that the decision taken by the attractor, e.g. D1 winning, was supported by the continuing presence of the decision cues $\lambda_1 > \lambda_2$. Because it was an integrate-and-fire network, we were able to integrate over the currents in all the neuronal pools, and to convolve these currents with the hemodynamic response function, to predict that the fMRI BOLD signal would increase as $\Delta\lambda$ increased in brain areas where this type of decision-making process was being implemented (Rolls et al., 2010a). We then performed two fMRI experiments to test the prediction. We used decisions about the value of stimuli, about which of a pair from a set of sequentially presented odors was more pleasant, and about which of a pair from a set of thermal stimuli sequentially presented to the hand was more pleasant. The value, measured by their subjective pleasantness, of these stimuli is represented in the human orbitofrontal cortex (Rolls et al., 2010a). We found that in the medial prefrontal cortex, a region just anterior to the orbitofrontal cortex and previously implicated in reward value-based decision making (Grabenhorst et al., 2008), the BOLD signal increased with $\Delta\lambda$ (Rolls et al., 2010a).

This investigation provided a clear theoretical basis for understanding the relation between the discriminability of stimuli, the subjective confidence in a decision before the outcome is known, the firing rates of and the currents in the neurons in an attractor decision-making network, and the fMRI BOLD signal (Rolls et al., 2010a).

4.2. Decision confidence and correct vs incorrect decisions

If one makes a correct decision, consistent with the evidence, then one's confidence is higher than when one makes an incorrect decision (as shown by confidence ratings) (Vickers, 1979; Vickers and Packer, 1982). Consistent with this, the probability that a rat will abort a trial to try again is higher if a decision just made is incorrect (Kepecs et al., 2008). Why does this occur, and in which brain regions is the underlying processing implemented?

We investigated this in the same decision-making network illustrated in Fig. 8, and found that the integrate-and-fire attractor decision-making network predicts higher and shorter latency neuronal responses, and higher BOLD signals, in brain areas involved in making choices, on correct than on error trials (Fig. 9) (Rolls et al., 2010b). Moreover, the firing rate increases approximately linearly with $\Delta\lambda$ on correct trials, and decreases with $\Delta\lambda$ on error trials. The reason for this behavior of the choice attractor network is that on correct trials, when the network influenced by the spiking noise has reached an attractor state supported by the recurrent connections between the neurons, the external inputs to the network that provide the evidence for the decision will have firing rates that are consistent with the decision, so the firing rate of the winning attractor will be higher when the correct attractor (given the evidence) wins the competition than when the incorrect attractor wins the competition.

We then showed in an experimental fMRI investigation that this BOLD signature of correct vs incorrect decisions, which reflects confidence in whether a decision just taken is correct, is found in the medial prefrontal cortex area 10, the posterior and subgenual cingulate cortex, and the dorsolateral prefrontal cortex, but not in the mid-orbitofrontal cortex, where activations instead reflect the pleasantness or subjective affective value of the stimuli used as inputs to the choice decision-making process (Rolls et al., 2010b).

This approach to decision-making, in contrast to some mathematical models of decision-making as an accumulation or diffusion process reviewed above, thus makes testable predictions about how correct vs error performance is implemented in the brain, and these predictions are supported by experimental results showing that areas involved in choice decision-making (or that receive from them) have activity consistent with these predictions, and that other areas not involved in the choice-making part of the process do not (Rolls and Deco, 2010).

These investigations thus provide evidence that the medial prefrontal cortex area 10 is involved in taking decisions about reward value, using an attractor decision-making network architecture.

4.3. Decisions about decisions: monitoring decisions

We have seen that after a categorical decision there is nevertheless a continuous-valued representation of decision confidence encoded in the firing rates of the neurons in a decision-making integrate-and-fire attractor neuronal network (Fig. 9) (Rolls and Deco, 2010; Rolls et al., 2010a,b). What happens if instead of having to report or assess the continuous-valued representation of confidence in a decision one has taken, one needs to take a decision based on one's (partly) continuous-valued confidence estimate that one has just made a correct or incorrect decision? One might for example wait for a reward if one thinks one's decision was correct, or alternatively stop waiting on that trial and start another trial or action. We propose that in this case, one needs a second decision-making network, which takes decisions based on one's decision confidence (Insabato et al., 2010).

The architecture for such a system is illustrated in Fig. 10. The first network takes decisions conventionally based on its inputs λ_A and λ_B. We now consider the operation of the confidence decision (second) network. The confidence network has two selective pools of neurons, one of which (C) responds to represent confidence in the decision, and the other of which responds when there is little or a lack of confidence in the decision (LC). (In the experiment of Kepecs et al. (2008), C corresponds to a decision to stay and wait for a reward, i.e. what they call the positive outcome population, though it really represents confidence or a prediction that the decision just taken will have a positive outcome. LC corresponds to a decision to abort a trial and not wait for a possible reward, i.e. what they call the negative outcome population, though it really represents lack of confidence that the decision just taken will have a positive outcome, equivalent to confidence that the decision just taken will have a negative outcome.) If the output firing of the winning attractor of the first, decision-making, network (whether DA or DB, reflected in their sum) is high because the decision just taken has sensory inputs consonant with the decision, then the confidence decision network acting as a second level network takes the decision, probabilistically as before, to have confidence in the decision, and the C population probabilistically wins the competition. If the output firing of DA and DB (reflected in their sum) is low because the decision just taken has sensory inputs that are not

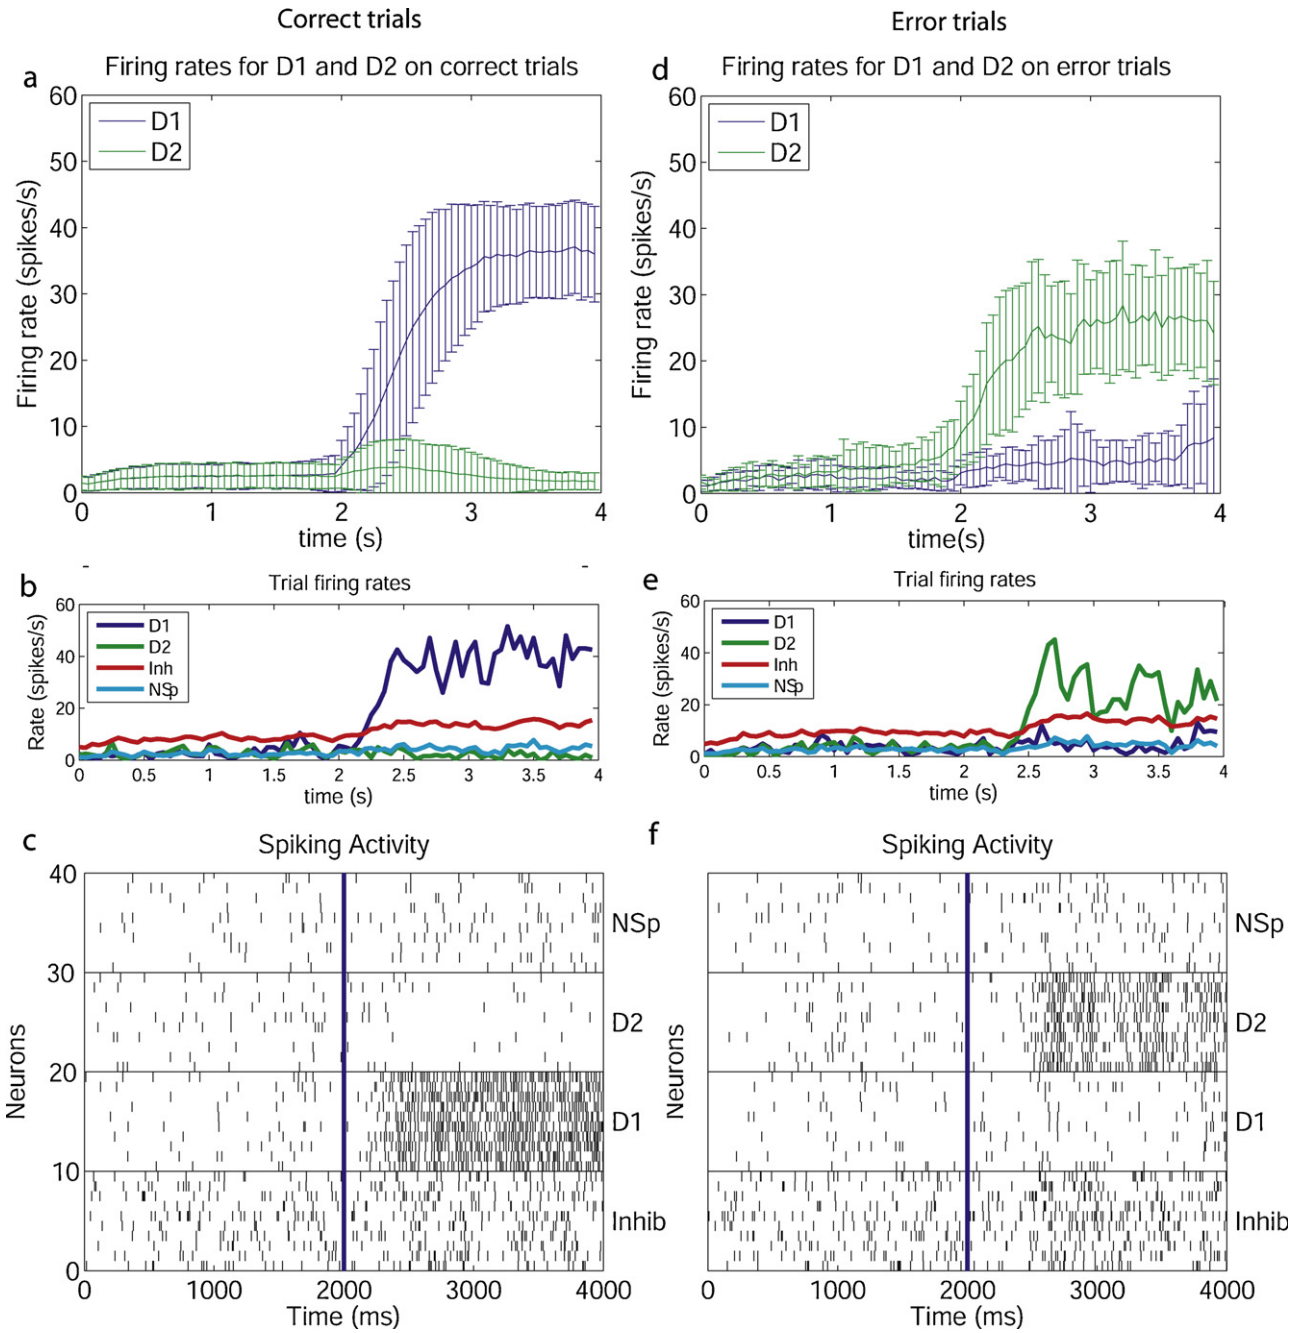

Fig. 9. (a) and (d) Firing rates (mean ± sd) for correct and error trials for an intermediate level of difficulty (ΔI = 32). The period 0–2 s is the spontaneous firing, and the decision cues were turned on at time = 2 s. The means were calculated over 1000 trials. D1: firing rate of the D1 population of neurons, which is the correct population. D2: firing rate of the D2 population of neurons, which is the incorrect population. A correct trial was one in which the mean rate of the D1 attractor averaged > 10 spikes/s than that of the competing population for the last 1000 ms of the simulation runs. (Given the attractor nature of the network and the parameters used, the network reached one of the attractors on >90% of the 1000 trials, and this criterion clearly separated these trials, as indicated by the mean rates and standard deviations for the last s of the simulation as shown.) (b) and (e) The firing rates of the four populations of neurons on a single trial for a correct (b) and error (e) decision. Inh is the inhibitory population that uses GABA as a transmitter. NSp is the non-specific population of neurons (see Fig. 8). (c) and (f) Rastergrams for the trials shown in (b) and (e). 10 neurons from each of the four pools of neurons are shown.
After Rolls et al. (2010b).

consonant with the decision, then with weaker driving inputs to the C network, it loses the competition with LC. The confidence network in this case takes the decision, probabilistically as before, to have a lack of confidence in the decision, in that the LC population wins the competition. The confidence decision network thus acts as a decision-making network to make confident decisions if the firing rates from the first, decision-making, network are high, and to make lack of confidence decisions if the firing rates from the first, decision-making, network are low (Insabato et al., 2010).

The finding that the changes of firing rates found in the rat by Kepecs et al. (2008), as a function of $\Delta\lambda$ are comparable with those in the model provides good support for the model (Insabato et al., 2010).

This approach indicates that an integrate-and-fire attractor neuronal decision-making network encodes confidence in its firing rates, and that adding a second attractor network allows decisions to be made about whether to change the decision made by the first network, and for example abort the trial or strategy. The second network, the confidence decision network, is in effect monitoring

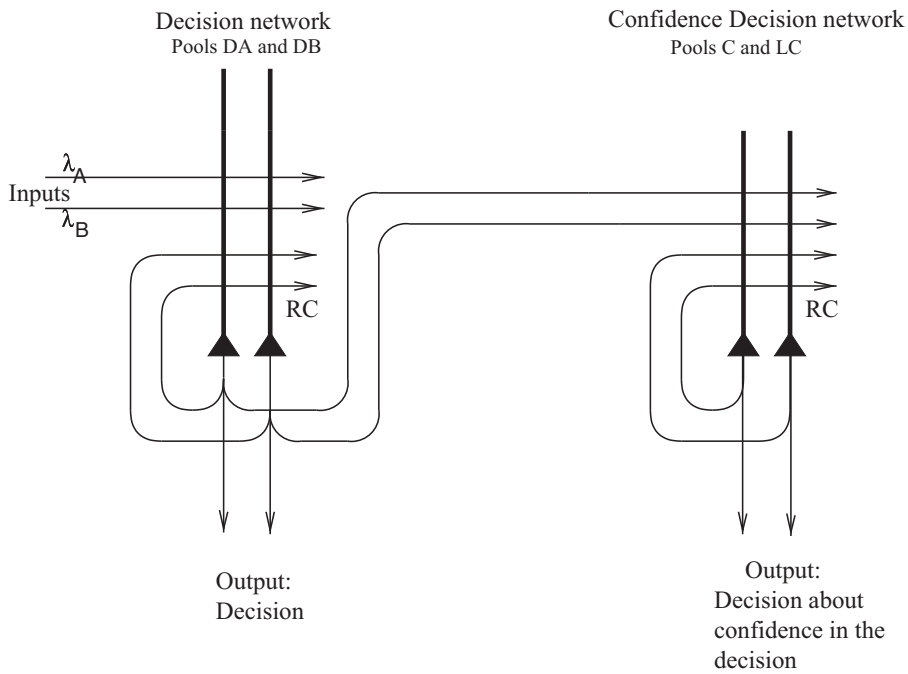

Fig. 10. Network architecture for decisions about confidence estimates. The first network is a decision-making network, and its outputs are sent to a second network that makes decisions based on the firing rates from the first network, which reflect the decision confidence. In the first network, high firing of neuronal population (or pool) DA represents decision A, and high firing of population DB represents decision B. Pools DA and DB receive stimulus-related inputs λ_1 and λ_2, the evidence for each of the decisions, and these bias the attractor networks, which have internal positive feedback produced by the recurrent excitatory connections (RC). Pools DA and DB compete through inhibitory interneurons. The neurons are integrate-and-fire spiking neurons with random spiking times (for a given mean firing rate) which introduce noise into the network and influence the decision-making, making it probabilistic. The second network is a confidence decision attractor network, and receives inputs from the first network. The confidence decision network has two selective pools of neurons, one of which (C) responds to represent confidence in the decision, and the other of which responds when there is little or a lack of confidence in the decision (LC). The C neurons receive the outputs from the selective pools of the (first) decision-making network, and the LC neurons receive $\lambda_{Reference}$ which is from the same source but saturates at 40 spikes/s, a rate that is close to the rates averaged across correct and error trials of the sum of the firing in the selective pools in the (first) decision-making network. (After Rolls and Deco, 2010; and Insabato et al., 2010.)

the decisions taken by the first network, and can cause a change of strategy or behavior if the assessment of the decision taken by the first network does not seem a confident decision. The implication is that some types of 'self-monitoring' can be accounted for by simple, two-attractor network, computational processes. Other types of self-monitoring may be more complex (Rolls, 2011, 2012b).

5. Predictions of decisions from noise in the brain even before the evidence for the decision has been provided

One of the crucial concepts of the biologically plausible integrate-and-fire attractor decision-making networks that we have been investigating is that there is internal noise or randomness in the system due to the close to Poisson spike time distribution of the neuronal action potentials, which when combined with the finite size of the network, leads to statistical fluctuations and thus to probabilistic decision-making (Rolls and Deco, 2010). These statistical fluctuations are evident in the firing rates of the neurons in the spontaneous firing rate period before the decision cues are applied. If these fluctuations increase the firing rates of one of the decision pools before the cues are applied more than the other decision pool, then this can make the pool with the higher rates when the decision cues are applied more likely to win the competition (Rolls and Deco, 2011b). Indeed, we have shown that it is possible to predict which decision will be taken, with 0 or low values of $\Delta\lambda$, to more than 70% correct (Rolls and Deco, 2011b). This has interesting implications for understanding determinism and free will (Rolls, 2011, 2012b).

6. What influences noise in the brain, and does noise still apply in large biologically plausible networks?

6.1. Graded firing rate representations increase noise in decision-making networks

Representations in the cortex are often distributed with graded firing rates in the neuronal populations (Rolls, 2008; Rolls and Treves, 2011). The firing rate probability distribution of each neuron to a set of stimuli is often exponential or gamma. In processes in the brain such as decision-making that are influenced by the noise produced by the close to random spike timings of each neuron for a given mean rate, the noise with this graded type of representation may be larger than with the binary firing rate distribution that is usually investigated.

In integrate-and-fire simulations of an attractor decision-making network, we showed that the noise is indeed greater for a given sparseness of the representation for graded, exponential, than for binary firing rate distributions (Webb et al., 2011). The greater noise was measured by faster escaping times from the spontaneous firing rate state when the decision cues are applied, and this corresponds to faster decision or reaction times. The greater noise was also evident as less stability of the spontaneous firing state before the decision cues are applied.

The implication is that spiking-related noise will continue to be a factor that influences processes such as decision-making, signal detection, short-term memory, and memory recall even with the quite large networks found in the cerebral cortex. In these networks there are several thousand recurrent collateral synapses onto each neuron. The greater noise with graded firing rate

distributions has the advantage that it can increase the speed of operation of cortical circuitry (Webb et al., 2011).

6.2. Diluted network connectivity decreases the noise in decision-making networks

The connectivity of the cerebral cortex is diluted, with the probability of excitatory connections between even nearby pyramidal cells rarely more than 0.1, and in the hippocampus 0.04 (Rolls, 2008).

To investigate the extent to which this diluted connectivity affects the dynamics of attractor networks in the cerebral cortex, we simulated an integrate-and-fire attractor network taking decisions between competing inputs with diluted connectivity of 0.25 or 0.1, and with the same number of synaptic connections per neuron for the recurrent collateral synapses within an attractor population as for full connectivity (Rolls and Webb, 2012).

The results indicated that there was less spiking-related noise with the diluted connectivity in that the stability of the network when in the spontaneous state of firing increased, and the accuracy of the correct decisions increased. The decision times were a little slower with diluted than with complete connectivity (Rolls and Webb, 2012).

Given that the capacity of the network is set by the number of recurrent collateral synaptic connections per neuron (Rolls, 2008), on which there is a biological limit, the findings indicate that the stability of cortical networks, and the accuracy of their correct decisions or memory recall operations, can be increased by utilizing diluted connectivity and correspondingly increasing the number of neurons in the network, with little impact on the speed of processing of the cortex. Thus diluted connectivity can decrease cortical spiking-related noise.

In addition, we showed that in this decision-making network the Fano factor for the trial-to-trial variability of the neuronal firing decreases from the spontaneous firing state value when the attractor network makes a decision (Rolls and Webb, 2012), providing a neuronal network account of this change in variability found in many cortical regions when they are activated (Churchland et al., 2010).

7. Dynamical neuropsychiatry

There are two types of stability that it is important to recognize in attractor networks in the brain, including those involved in decision-making.

The first is that the high firing rate attractor states must be maintained sufficiently high in firing rate that they are stable, due to the positive feedback in the recurrent collaterals dominating the spiking noise-driven fluctuations that may tend to knock the system out of its high firing rate attractor. If the rates are insufficiently high, due to for example reduced glutamatergic excitatory transmitter efficacy, then the poor stability may result in poor working memory and attention (Rolls, 2005, 2008; Rolls and Deco, 2010). This has been related to the cognitive symptoms of schizophrenia, which may include difficulty in working memory and attention (Loh et al., 2007a; Rolls, 2005, 2008, 2012a; Rolls and Deco, 2010, 2011a; Rolls et al., 2008b), and normal aging (Rolls, 2008; Rolls and Deco, 2010). Too much stability of the high firing rate attractor states, produced for example by hyperglutamatergia, may contribute to the symptoms of obsessive–compulsive disorder, by making attractors (in different brain systems in different individuals) too persistent (Rolls, 2012a; Rolls et al., 2008a).

The second type of stability is that of the low firing rate spontaneous state, when decision or memory recall cues are not being applied. Too little stability here, caused for example by GABA downregulation, could cause attractor networks to jump without provocation by a stimulus, into a high firing rate state produced by the spiking-related noise (Loh et al., 2007a,b; Rolls, 2008). It has been suggested that the positive symptoms of schizophrenia may be related to reduced inhibition in attractor networks (Loh et al., 2007a,b; Rolls, 2008, 2012a; Rolls and Deco, 2011a; Rolls et al., 2008b).

There are implications for treatment suggested by these approaches for neuropsychiatry (Rolls, 2012a; Rolls and Deco, 2011a; Rolls et al., 2008b).

8. Decision-making, oscillations, and communication through coherence

Oscillations in the brain (Buzsaki, 2006; Wang, 2010) may influence decision-making. For example, theta oscillations increased the speed of operation of a decision-making attractor network (Smerieri et al., 2010). The mechanisms involved may have included concentrating the spikes in the network close together to facilitate the speed of response of neurons, and increasing the firing rate above a threshold for part of the theta cycle in a mechanism like stochastic resonance (Rolls and Treves, 2011; Smerieri et al., 2010). Gamma band synchronization also has effects that could influence decision-making (Buehlmann and Deco, 2010), though information transmission from one attractor decision-making network to another occurs at far lower values of the synaptic connection weights between the two networks than are required for gamma synchronization (Rolls et al., 2012). This raises important questions (Rolls and Treves, 2011; Rolls et al., 2012) about the hypothesis of communication through coherence (Fries, 2005, 2009) at least in the context of connected decision-making attractor networks.

A mechanism by which oscillations may influence neuronal firing and decision-making is by resetting neuronal activity. For example, a process such as memory recall may occur within a single theta cycle, and then be quenched so that a new attempt at recall can be made in the next theta cycle. This has the potential advantage that in a changing, ambiguous, or uncertain situation, several attempts can be made at the memory recall, without previous attempts dominating the memory state for a period due to attractor dynamics in autoassociation networks (Rolls and Treves (1998) page 118). Effects consistent with this prediction have recently been observed in the rat hippocampus (Jezek et al., 2011): in response to an instantaneous transition between two familiar and similar spatial contexts, hippocampal neurons in one theta cycle indicated one place, and in another theta cycle another place. These findings indicate that, in hippocampal CA3, pattern-completion dynamics can occur within each individual theta cycle. Reset, with potentially different recall in different theta cycles, may facilitate rapid updating and correction of recall. Memory recall may be seen as the same neural network process as decision-making in an attractor network (Rolls, 2008).

9. Conclusions

The operation of attractor decision-making networks described here applies to the use of attractor networks in the brain for many other processes, depending on the cortical region, including short-term memory, long-term memory, attention, associative thought processes, creativity, and disordered thought processes (Rolls, 2008, 2012b; Rolls and Deco, 2010). In all these cases, there is a transition to a categorical state determined by the input, the noise, and the possible attractor states in the network.

The approach taken here and elsewhere (Rolls, 2008, 2012b; Rolls and Deco, 2010) is towards how the dynamics of attractor networks used for decision-making and related functions is

implemented neurobiologically, and in this respect is what we term a 'mechanistic' model. This may be contrasted with what we describe as 'phenomenological' models, which may model phenomena at some abstract level. Examples in the area of decision-making are the accumulator, counter, or race models of decision-making in which the noisy evidence for different choices accumulates by integration until some threshold is reached (see Section 2). In contrast to the 'mechanistic' approach taken here, such 'phenomenological' models do not describe a neurobiological mechanism, do not make specific predictions about how neuronal firing rates, synaptic currents, or BOLD signals will alter as decisions become easier, and do not make specific predictions about confidence in a decision, but do account for behavioral performance. In this context, it is of interest that the integrate-and-fire attractor network approach as well as some of the other approaches can be analyzed as diffusion processes (Roxin and Ledberg, 2008), and that the attractor approach is not inconsistent with these accumulator or with Bayesian (Beck et al., 2008; Ma et al., 2006, 2008) approaches, but instead defines a mechanism, from which many important properties emerge, such as a representation of decision confidence on correct and error trials, as shown here, and specific predictions that can be tested with the experimental methods of neuroscience. We emphasize that the empirical data alone are not sufficient to provide an understanding of how a complex process such as decision-making takes place in the brain, and that to understand the mechanism of decision-making (and related categorization processes), to make predictions, and to suggest treatments for disorders of the dynamical processes involved, the empirical data must be complemented by a mechanistic neurobiologically based theory and model.

Acknowledgements

G.D. and L.A. were supported by the "CONSOLIDER-INGENIO 2010 Programme in the "Brainglot" project (Bilingualism and Cognitive Neuroscience) CSD2007-00012. G.D. was further supported by the European Community's Seventh Framework Programme under the project "BrainScales" (project number 269921) and the Spanish Ministry of Science in the project "Communication and Information Processing Between Cortical Circuits: Oscillations and Plasticity" SAF2010-16085. Parts of the research described were supported by the Oxford Centre for Computational Neuroscience, http://www.oxcns.org."

References

Albantakis, L., Deco, G., 2009. The encoding of alternatives in multiple-choice decision making. Proc. Natl. Acad. Sci. U.S.A. 106, 10308–10313.
Albantakis, L., Deco, G., 2011. Changes of mind in an attractor network of decision-making. PLoS Comput. Biol. 7, e1002086.
Amit, D.J., Brunel, N., 1997. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex 7, 237–252.
Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., Pouget, A., 2008. Probabilistic population codes for Bayesian decision making. Neuron 60, 1142–1152.
Bogacz, R., Brown, E., Moehlis, J., Holmes, P., Cohen, J.D., 2006. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. 113, 700–765.
Bogacz, R., Gurney, K., 2007. The basal ganglia and cortex implement optimal decision making between alternative actions. Neural Comput. 19, 442–477.
Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A., 1992. The analysis of visual motion: a comparison of neuronal and psychophysical performance. J. Neurosci. 12, 4745–4765.
Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A., 1993. Responses of neurons in macaque MT to stochastic motion signals. Vis. Neurosci. 10, 1157–1169.
Brown, S.D., Heathcote, A., 2008. The simplest complete model of choice response time: linear ballistic accumulation. Cogn. Psychol. 57, 153–178.
Brunel, N., Wang, X.J., 2001. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. J. Comput. Neurosci. 11, 63–85.
Buehlmann, A., Deco, G., 2010. Optimal information transfer in the cortex through synchronization. PLoS Comput. Biol. 6, e1000934.
Busemeyer, J.R., Townsend, J.T., 1993. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. Psychol. Rev. 100, 432–459.
Buzsaki, G., 2006. Rhythms of the Brain. Oxford University Press, New York.
Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.J., Pouget, A., Shadlen, M.N., 2011. Variance as a signature of neural computations during decision making. Neuron 69, 818–831.
Churchland, A.K., Kiani, R., Shadlen, M.N., 2008. Decision-making with multiple alternatives. Nat. Neurosci. 11, 693–702.
Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., Bradley, D.C., Smith, M.A., Kohn, A., Movshon, J.A., Armstrong, K.M., Moore, T., Chang, S.W., Snyder, L.H., Lisberger, S.G., Priebe, N.J., Finn, I.M., Ferster, D., Ryu, S.I., Santhanam, G., Sahani, M., Shenoy, K.V., 2010. Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat. Neurosci. 13, 369–378.
Cisek, P., Puskas, G.A., El-Murr, S., 2009. Decisions in changing conditions: the urgency-gating model. J. Neurosci. 29, 11560–11571.
Deco, G., Rolls, E.T., 2006. Decision-making and Weber's Law: a neurophysiological model. Eur. J. Neurosci. 24, 901–916.
Deco, G., Rolls, E.T., Romo, R., 2010. Synaptic dynamics and decision-making. Proc. Natl. Acad. Sci. U.S.A. 107, 7545–7549.
Deco, G., Scarano, L., Soto-Faraco, S., 2007. Weber's law in decision making: integrating behavioral data in humans with a neurophysiological model. J. Neurosci. 27, 11192–11200.
Ditterich, J., 2006a. Evidence for time-variant decision making. Eur. J. Neurosci. 24, 3628–3641.
Ditterich, J., 2006b. Stochastic models of decisions about motion direction: behavior and physiology. Neural Netw. 19, 981–1012.
Ditterich, J., 2010. A comparison between mechanisms of multi-alternative perceptual decision making: ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. Front. Neurosci. 4, 184.
Dragalin, V.P., Tartakovsky, A.G., Veeravalli, V.V., 1999. Multihypothesis sequential probability ratio tests – Part I: Asymptotic optimality. IEEE Trans. Inf. Theory 45, 2448–2461.
Fries, P., 2005. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends Cogn. Sci. 9, 474–480.
Fries, P., 2009. Neuronal gamma-band synchronization as a fundamental process in cortical computation. Annu. Rev. Neurosci. 32, 209–224.
Furman, M., Wang, X.J., 2008. Similarity effect and optimal control of multiple-choice decision making. Neuron 60, 1153–1168.
Gardiner, C.W., 1985. Handbook of Stochastic Methods. New York, Springer.
Gnadt, J.W., Andersen, R.A., 1988. Memory related motor planning activity in posterior parietal cortex of macaque. Experimentelle Hirnforschung[[n1]]Exp. brain Res. 70, 216–220.
Gold, J.I., Shadlen, M.N., 2001. Neural computations that underlie decisions about sensory stimuli. Trends Cogn. Sci. 5, 10–16.
Gold, J.I., Shadlen, M.N., 2002. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron 36, 299–308.
Gold, J.I., Shadlen, M.N., 2007. The neural basis of decision making. Annu. Rev. Neurosci. 30, 535–574.
Grabenhorst, F., Rolls, E.T., 2011. Value, pleasure, and choice in the ventral prefrontal cortex. Trends Cogn. Sci. 15, 56–67.
Grabenhorst, F., Rolls, E.T., Parris, B.A., 2008. From affective value to decision-making in the prefrontal cortex. Eur. J. Neurosci. 28, 1930–1939.
Green, D., Swets, J., 1966. Signal Detection Theory and Psychophysics. New York, Wiley.
Heekeren, H.R., Marrett, S., Bandettini, P.A., Ungerleider, L.G., 2004. A general mechanism for perceptual decision-making in the human brain. Nature 431, 859–862.
Hempel, C.M., Hartman, K.H., Wang, X.J., Turrigiano, G.G., Nelson, S.B., 2000. Multiple forms of short-term plasticity at excitatory synapses in rat medial prefrontal cortex. J. Neurophysiol. 83, 3031–3041.
Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vazquez, Y., Camarillo, L., Romo, R., 2010. Decoding a perceptual decision process across cortex. Neuron 66, 300–314.
Hick, W.E., 1952. On the rate of gain of information. Q. J. Exp. Psychol. 4, 11–26.
Huk, A.C., Shadlen, M.N., 2005. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J. Neurosci. 25, 10420–10436.
Insabato, A., Pannunzi, M., Rolls, E.T., Deco, G., 2010. Confidence-related decision-making. J. Neurophysiol. 104, 539–547.
Jezek, K., Henriksen, E.J., Treves, A., Moser, E.I., Moser, M.-B., 2011. Theta-paced flickering between place-cell maps in the hippocampus. Nature 278, 246–249.
Jonsson, F.U., Olsson, H., Olsson, M.J., 2005. Odor emotionality affects the confidence in odor naming. Chem. Senses 30, 29–35.
Kepecs, A., Uchida, N., Zariwala, H.A., Mainen, Z.F., 2008. Neural correlates, computation and behavioural impact of decision confidence. Nature 455, 227–231.
Kiani, R., Shadlen, M.N., 2009. Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324, 759–764.
Kim, J.N., Shadlen, M.N., 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. 2, 176–185.
Laming, D., 1968. Information Theory of Choice Reaction Time. Chichester, Wiley.
Leite, F.P., Ratcliff, R., 2010. Modeling reaction time and accuracy of multiple-alternative decisions. Atten. Percept. Psychophys. 72, 246–273.

Lemus, L., Hernandez, A., Luna, R., Zainos, A., Nacher, V., Romo, R., 2007. Neural correlates of a postponed decision report. Proc. Natl. Acad. Sci. U.S.A. 104, 17174–17179.
Lo, C.C., Wang, X.J., 2006. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nat. Neurosci. 9, 956–963.
Loh, M., Rolls, E.T., Deco, G., 2007a. A dynamical systems hypothesis of schizophrenia. PLoS Comput. Biol. 3, e228 doi:10.1371/journal.pcbi.0030228.
Loh, M., Rolls, E.T., Deco, G., 2007b. Statistical fluctuations in attractor networks related to schizophrenia. Pharmacopsychiatry 40, S78–S84.
Luce, R.D., 1991. Response Times: Their Role in Inferring Elementary Mental Organization. Oxford University Press, New York.
Lund, J.S., Angelucci, A., Bressloff, P.C., 2003. Anatomical substrates for functional columns in macaque monkey primary visual cortex. Cereb. Cortex 13, 15–24.
Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A., 2006. Bayesian inference with probabilistic population codes. Nat. Neurosci. 9, 1432–1438.
Ma, W.J., Beck, J.M., Pouget, A., 2008. Spiking networks for Bayesian inference and choice. Curr. Opin. Neurobiol. 18, 217–222.
Marti, D., Deco, G., Mattia, M., Gigante, G., Del Giudice, P., 2008. A fluctuation-driven mechanism for slow decision processes in reverberant networks. PLoS ONE 3, e2534.
Martinez-Garcia, M., Rolls, E.T., Deco, G., Romo, R., 2011. Neural and computational mechanisms of postponed decisions. Proc. Natl. Acad. Sci. U.S.A. 108, 11626–11631.
Masquelier, T., Albantakis, L., Deco, G., 2011. The timing of vision—how neural processing links to different temporal dynamics. Front. Psychol. 2, 151.
Mazurek, M.E., Roitman, J.D., Ditterich, J., Shadlen, M.N., 2003. A role for neural integrators in perceptual decision making. Cereb. Cortex 13, 1257–1269.
McMillen, T., Holmes, P., 2006. The dynamics of choice among multiple alternatives. J. Math. Psychol. 50, 30–57.
Mongillo, G., Barak, O., Tsodyks, M., 2008. Synaptic theory of working memory. Science 319, 1543–1546.
Neyman, J., Pearson, E.S., 1933. On the problem of the most efficient tests of statistical hypotheses. Phil. Trans. R. Soc. 231, 289–337.
Niwa, M., Ditterich, J., 2008. Perceptual decisions between multiple directions of visual motion. J. Neurosci. 28, 4435–4445.
Opris, I., Bruce, C.J., 2005. Neural circuitry of judgment and decision mechanisms. Brain Res. 48, 509–526.
Palmer, J., Huk, A.C., Shadlen, M.N., 2005. The effect of stimulus strength on the speed and accuracy of a perceptual decision. J. Vis. 5, 376–404.
Platt, M.L., Glimcher, P.W., 1999. Neural correlates of decision variables in parietal cortex. Nature 400, 233–238.
Pleskac, T.J., Busemeyer, J.R., 2010. Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. Psychol. Rev. 117, 864–901.
Purcell, B.A., Heitz, R.P., Cohen, J.Y., Schall, J.D., Logan, G.D., Palmeri, T.J., 2010. Neurally constrained modeling of perceptual decision making. Psychol. Rev. 117, 1113–1143.
Ratcliff, R., 1978. Theory of memory retrieval. Psychol. Rev. 85, 59–108.
Ratcliff, R., Cherian, A., Segraves, M., 2003. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. J. Neurophysiol. 90, 1392–1407.
Ratcliff, R., McKoon, G., 2008. The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. 20, 873–922.
Ratcliff, R., Rouder, J.N., 1998. Modeling response times for two-choice decisions. Psychol. Sci. 9, 347–356.
Ratcliff, R., Smith, P.L., 2004. A comparison of sequential sampling models for two-choice reaction time. Psychol. Rev. 111, 333–367.
Reddi, B.A., Carpenter, R.H., 2000. The influence of urgency on decision time. Nat. Neurosci. 3, 827–830.
Renart, A., Brunel, N., Wang, X.-J., 2003. Mean field theory of irregularly spiking neuronal populations and working memory in recurrent cortical networks. In: Feng, J. (Ed.), Computational Neuroscience: A Comprehensive Approach. Chapman and Hall, Boca-Raton, pp. 431–490.
Roitman, J.D., Shadlen, M.N., 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. 22, 9475–9489.
Rolls, E.T., 2005. Emotion Explained. Oxford University Press, Oxford.
Rolls, E.T., 2008. Memory, Attention, and Decision-Making: A Unifying Computational Neuroscience Approach. Oxford University Press, Oxford.
Rolls, E.T., 2011. Consciousness, decision-making, and neural computation. In: Cutsuridis, V., Hussain, A., Taylor, J.G. (Eds.), Perception-Action Cycle: Models, Algorithms and Systems. Springer, Berlin, pp. 287–333.
Rolls, E.T., 2012a. Glutamate, obsessive-compulsive disorder, schizophrenia, and the stability of cortical attractor neuronal networks. Pharmacol. Biochem. Behav. 100, 736–751.
Rolls, E.T., 2012b. Neuroculture. On the Implications of Brain Science. Oxford University Press, Oxford.
Rolls, E.T., Deco, G., 2010. The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function. Oxford University Press, Oxford.
Rolls, E.T., Deco, G., 2011a. A computational neuroscience approach to schizophrenia and its onset. Neurosci. Biobehav. Rev. 35, 1644–1653.
Rolls, E.T., Deco, G., 2011b. Prediction of decisions from noise in the brain before the evidence is provided. Front. Neurosci. 5, 33.
Rolls, E.T., Grabenhorst, F., Deco, G., 2010a. Choice, difficulty, and confidence in the brain. Neuroimage 53, 694–706.
Rolls, E.T., Grabenhorst, F., Deco, G., 2010b. Decision-making, errors, and confidence in the brain. J. Neurophysiol. 104, 2359–2374.
Rolls, E.T., Loh, M., Deco, G., 2008a. An attractor hypothesis of obsessive-compulsive disorder. Eur. J. Neurosci. 28, 782–793.
Rolls, E.T., Loh, M., Deco, G., Winterer, G., 2008b. Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. Nat. Rev. Neurosci. 9, 696–709.
Rolls, E.T., Treves, A., 1998. Neural Networks and Brain Function. Oxford University Press, Oxford.
Rolls, E.T., Treves, A., 2011. The neuronal encoding of information in the brain. Prog. Neurobiol. 95, 448–490.
Rolls, E.T., Webb, T.J., 2012. Cortical attractor network dynamics with diluted connectivity. Brain Res. 1434, 212–225.
Rolls, E.T., Webb, T.J., Deco, G., 2012. Communication before coherence. Eur. J. Neurosci. in revision.
Romo, R., Hernandez, A., Zainos, A., 2004. Neuronal correlates of a perceptual decision in ventral premotor cortex. Neuron 41, 165–173.
Romo, R., Salinas, E., 2003. Flutter discrimination: neural codes, perception, memory and decision making. Nat. Rev. Neurosci. 4, 203–218.
Roxin, A., Ledberg, A., 2008. Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. PLoS Comput. Biol. 4, e1000046.
Shadlen, M.N., Newsome, W.T., 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. 86, 1916–1936.
Smerieri, A., Rolls, E.T., Feng, J., 2010. Decision reaction time, slow inhibition, and theta rhythm. J. Neurosci. 30, 14173–14181.
Stone, M., 1960. Models for choice-reaction time. Psychometrika 25, 251–260.
Tanner Jr., W.P., Swets, J.A., 1954. A decision-making theory of visual detection. Psychol. Rev. 61, 401–409.
Townsend, J.T., Ashby, F.G., 1983. The Stochastic Modeling of Elementary Psychological Processes. Cambridge University Press, Cambridge.
Tuckwell, H., 1988. Introduction to Theoretical Neurobiology. Cambridge University Press, Cambridge.
Usher, M., McClelland, J.L., 2001. The time course of perceptual choice: the leaky, competing accumulator model. Psychol. Rev. 108, 550–592.
Vandekerckhove, J., Tuerlinckx, F., 2007. Fitting the Ratcliff diffusion model to experimental data. Psychon. Bull. Rev. 14, 1011–1026.
Vickers, D., 1970. Evidence for an accumulator model of psychophysical discrimination. Ergonomics 13, 37–58.
Vickers, D., 1979. Decision Processes in Visual Perception. Academic Press, New York.
Vickers, D., Packer, J., 1982. Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. Acta Psychol. (Amst.) 50, 179–197.
Wald, A., 1947. Sequential Analysis. New York, Wiley.
Wang, X.J., 2002. Probabilistic decision making by slow reverberation in cortical circuits. Neuron 36, 955–968.
Wang, X.J., 2008. Decision making in recurrent neuronal circuits. Neuron 60, 215–234.
Wang, X.J., 2010. Neurophysiological and computational principles of cortical rhythms in cognition. Physiol. Rev. 90, 1195–1268.
Wang, Y., Markram, H., Goodman, P.H., Berger, T.K., Ma, J., Goldman-Rakic, P.S., 2006. Heterogeneity in the pyramidal network of the medial prefrontal cortex. Nat. Neurosci. 9, 534–542.
Webb, T., Rolls, E.T., Deco, G., Feng, J., 2011. Noise in attractor networks in the brain produced by graded firing rate representations. PLoS ONE 6, e23630.
Wong, K.F., Huk, A.C., 2008. Temporal dynamics underlying perceptual decision making: Insights from the Interplay between an attractor model and parietal neurophysiology. Front. Neurosci. 2, 245–254.
Wong, K.F., Huk, A.C., Shadlen, M.N., Wang, X.J., 2007. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. Front. Comput. Neurosci. 1, 6.
Wong, K.F., Wang, X.J., 2006. A recurrent network mechanism of time integration in perceptual decisions. J. Neurosci. 26, 1314–1328.
Zucker, R.S., Regehr, W.G., 2002. Short-term synaptic plasticity. Annu. Rev. Physiol. 64, 355–405.