



Research report

Non-reward neural mechanisms in the orbitofrontal cortex

Edmund T. Rolls ^{a,b,*} and Gustavo Deco ^{c,d}

^a Oxford Centre for Computational Neuroscience, Oxford, UK

^b University of Warwick, Department of Computer Science, Coventry, UK

^c Universitat Pompeu Fabra, Theoretical and Computational Neuroscience, Barcelona, Spain

^d Institutio Catalana de Recerca i Estudis Avancats (ICREA), Spain

ARTICLE INFO

Article history:

Received 28 January 2016

Reviewed 18 April 2016

Revised 3 May 2016

Accepted 24 June 2016

Action editor Angela Sirigu

Published online 15 July 2016

Keywords:

Emotion

Reward

Non-reward

Reward prediction error

Orbitofrontal cortex

Attractor network

Depression

Impulsiveness

ABSTRACT

Single neurons in the primate orbitofrontal cortex respond when an expected reward is not obtained, and behaviour must change. The human lateral orbitofrontal cortex is activated when non-reward, or loss occurs. The neuronal computation of this negative reward prediction error is fundamental for the emotional changes associated with non-reward, and with changing behaviour. Little is known about the neuronal mechanism. Here we propose a mechanism, which we formalize into a neuronal network model, which is simulated to enable the operation of the mechanism to be investigated. A single attractor network has a reward population (or pool) of neurons that is activated by expected reward, and maintain their firing until, after a time, synaptic depression reduces the firing rate in this neuronal population. If a reward outcome is not received, the decreasing firing in the reward neurons releases the inhibition implemented by inhibitory neurons, and this results in a second population of non-reward neurons to start and continue firing encouraged by the spiking-related noise in the network. If a reward outcome is received, this keeps the reward attractor active, and this through the inhibitory neurons prevents the non-reward attractor neurons from being activated. If an expected reward has been signalled, and the reward attractor neurons are active, their firing can be directly inhibited by a non-reward outcome, and the non-reward neurons become activated because the inhibition on them is released. The neuronal mechanisms in the orbitofrontal cortex for computing negative reward prediction error are important, for this system may be over-reactive in depression, under-reactive in impulsive behaviour, and may influence the dopaminergic 'prediction error' neurons.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction: Non-reward neurons in the orbitofrontal cortex, and the mechanism by which they may be activated

The aim of the research described here is to investigate the neuronal mechanisms that may underlie the computation of

non-reward in the primate including human orbitofrontal cortex (Rolls, 2014; Rolls & Grabenhorst, 2008; Thorpe, Rolls, & Maddison, 1983). This is a fundamentally important process involved in changing behaviour when an expected reward is not received. The process is important in many emotions, which can be produced if an expected reward is not received

* Corresponding author. Oxford Centre for Computational Neuroscience, Oxford, UK.

E-mail address: Edmund.Rolls@oxcns.org (E.T. Rolls).

URL: <http://www.oxcns.org>

<http://dx.doi.org/10.1016/j.cortex.2016.06.023>

0010-9452/© 2016 Elsevier Ltd. All rights reserved.

or lost, with examples including sadness and anger (Rolls, 2014). Consistent with this the psychiatric disorder of depression may arise when the non-reward system leading to sadness is too sensitive, or maintains its activity for too long (Rolls, 2016b), or has increased functional connectivity (Cheng et al., 2016). Conversely, if the non-reward system is underactive or is damaged by lesions of the orbitofrontal cortex, the decreased sensitivity to non-reward may contribute to increased impulsivity (Berlin, Rolls, & Iversen, 2005; Berlin, Rolls, & Kischka, 2004), and even antisocial and psychopathic behaviour (Rolls, 2014). For these reasons, understanding the mechanisms that underlie non-reward is important not only for understanding normal human behaviour and how it changes when rewards are not received, but also for understanding and potentially treating better some emotional and psychiatric disorders.

In the investigation described here, we develop hypotheses about the mechanisms involved in the computation of non-reward in the orbitofrontal cortex (Section 3) based on neurophysiological, functional neuroimaging, and neuropsychological evidence on the orbitofrontal cortex (Section 2). [These hypotheses are based on the functions of the orbitofrontal cortex in primates including humans, given the evidence that most of the orbitofrontal part that is present in rodents is agranular, and does not correspond to most of the primate including human orbitofrontal cortex (Passingham & Wise, 2012; Rolls, 2014; Wise, 2008).] We then develop an integrate-and-fire neuronal network model of how the non-reward signal is computed, and simulate the model to elucidate its properties, and to examine the parameters for the correct operation of the model (Sections 3 and 4). The particular neuronal responses that are modelled are how the non-reward neurons in the orbitofrontal cortex (Rosenkilde, Bauer, & Fuster, 1981; Thorpe et al., 1983) are triggered into a high firing rate with continuing activity, if no reward outcome is received after an expected reward has been indicated. The non-reward neurons should also be activated if after an expected reward signal, a punishment is received. The non-reward neurons should not be triggered if after an expected reward signal, a reward outcome is received (Thorpe et al., 1983).

2. Background to the hypotheses: Neurophysiological, neuroimaging, and neuropsychological evidence on what needs to be modelled to account for non-reward neurons

Some of the background evidence on the role of the primate including human orbitofrontal cortex in non-reward, including crucially what needs to be modelled at the neuronal, mechanism, level, is as follows.

Single neurons in the primate orbitofrontal cortex respond when an expected reward is not obtained, and behaviour must change. This was discovered by Thorpe et al. (1983), who found that 3.5% of neurons in the macaque orbitofrontal cortex detect different types of non-reward. These neurons thus signal negative reward prediction error, the reward outcome value minus the expected value. For example, some neurons responded in extinction, immediately after a lick had been made to a visual stimulus that had previously been

associated with fruit juice reward. Other non-reward neurons responded in a reversal task, immediately after the monkey had responded to the previously rewarded visual stimulus, but had obtained the punisher of salt taste rather than reward, indicating that the choice of stimulus should change in this visual discrimination reversal task. Importantly, at least some of these non-reward neurons continue firing for several seconds when an expected reward is not obtained, as illustrated in Fig. 1. These neurons do not respond when an expected punishment is received, for example the taste of salt from a correctly labelled dispenser. Different non-reward neurons may respond in different tasks, providing the potential for behaviour to reverse in one task, but not in all tasks (Rolls, 2014; Rolls & Grabenhorst, 2008; Thorpe et al., 1983). The existence of neurons in the middle part of the macaque orbitofrontal cortex that respond to non-reward (Thorpe et al., 1983) [originally described by Thorpe, Maddison, and Rolls (1979) and Rolls (1981, chap. 16)] is confirmed by recordings that revealed 10 such non-reward neurons (of 140 recorded, or approximately 7%) found in delayed match to sample and delayed response tasks by Joaquin Fuster and colleagues (Rosenkilde et al., 1981).

The orbitofrontal cortex is likely to be where non-reward is computed, for not only does it contain non-reward neurons, but it also contains neurons that encode the signals from which non-reward is computed. These signals include a representation of expected reward, and reward outcome (Rolls, 2014). Expected reward is represented in the primate orbitofrontal cortex in that many neurons respond to the sight of a food reward (Thorpe et al., 1983), learn to respond to any visual stimulus associated with a food reward (Rolls, Critchley, Mason, & Wakeman, 1996), very rapidly reverse the stimulus to which they respond when the reward outcome associated with each stimulus changes (Rolls et al., 1996), and represent reward value in that they stop responding to a visual stimulus when the reward is devalued by satiation (Critchley & Rolls, 1996a). Other orbitofrontal cortex neurons represent the expected reward value of olfactory stimuli, based on similar evidence (Critchley & Rolls, 1996a, 1996b; Rolls & Baylis, 1994; Rolls et al., 1996). Reward outcome is represented in the orbitofrontal cortex by its neurons that respond to taste and fat texture (Rolls & Baylis, 1994; Rolls, Critchley, Browning, Hernadi, & Lenard, 1999; Rolls, Verhagen, & Kadohisa, 2003; Rolls, Yaxley, & Sienkiewicz, 1990; Verhagen, Rolls, & Kadohisa, 2003), and do this based on their reward value as shown by the fact that their responses decrease to zero when the reward is devalued by feeding to satiety (Rolls, 2015b, 2016c; Rolls et al., 1999; Rolls, Sienkiewicz, & Yaxley, 1989). Moreover, the primate orbitofrontal cortex is key in these reward value representations, for it receives the necessary inputs from the inferior temporal visual cortex, insular taste cortex, and pyriform olfactory cortex, yet in these preceding areas reward value is not represented (Rolls, 2014, 2015a, 2015b). There is consistent evidence that expected value and reward outcome value are represented in the human orbitofrontal cortex (Grabenhorst, Rolls, & Bilderbeck, 2008; Kringelbach, O'Doherty, Rolls, & Andrews, 2003; O'Doherty, Kringelbach, Rolls, Hornak, & Andrews, 2001; Rolls, 2014, 2015b; Rolls O'Doherty, et al., 2003; Rolls, McCabe, & Redoute, 2008).

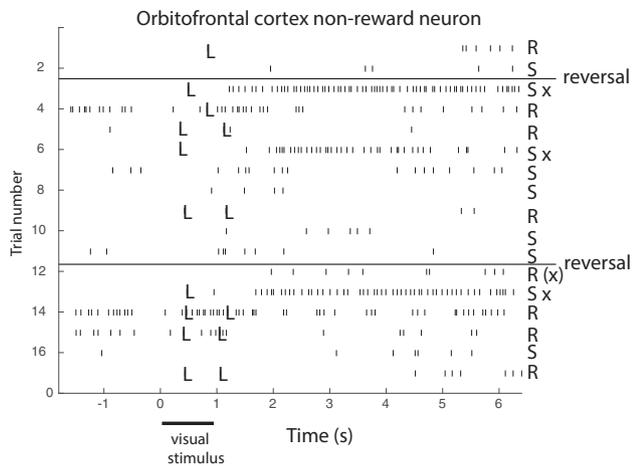


Fig. 1 – Error neuron: Responses of an orbitofrontal cortex neuron that responded only when the monkey licked to a visual stimulus during reversal, expecting to obtain fruit juice reward, but actually obtained the taste of aversive saline because it was the first trial of reversal (trials 3, 6, and 13). Each vertical line represents an action potential; each L indicates a lick response in the Go-NoGo visual discrimination task. The visual stimulus was shown at time 0 for 1 sec. The neuron did not respond on most reward (R) or saline (S) trials, but did respond on the trials marked S x, which were the first trials after a reversal of the visual discrimination on which the monkey licked to obtain reward, but actually obtained saline because the task had been reversed. The two times at which the reward contingencies were reversed are indicated. After responding to non-reward, when the expected reward was not obtained, the neuron fired for many seconds, and was sometimes still firing at the start of the next trial. It is notable that after an expected reward was not obtained due to a reversal contingency being applied, on the very next trial the macaque selected the previously non-rewarded stimulus. This shows that rapid reversal can be performed by a non-associative process, and must be rule-based. (After Thorpe et al., 1983).

In that most neurons in the macaque orbitofrontal cortex respond to reinforcers and punishers, or to stimuli associated with rewards and punishers, and do not respond in relation to responses, the orbitofrontal cortex is closely related to stimulus processing, including the stimuli that give rise to affective states (Rolls, 2014). When the orbitofrontal cortex computes errors, it computes mismatches between stimuli that are expected, and stimuli that are obtained, and in this sense the errors are closely related to those required to produce affective states, and to change the reward value representation of a stimulus. This type of error representation provided by the orbitofrontal cortex may thus be different from that represented in the cingulate cortex, in which behavioural responses are represented, where the errors may be more closely related to errors that arise when action–outcome expectations are not met, and where action–outcome rather than stimulus–reinforcer representations need to be corrected

(Grabenhorst & Rolls, 2011; Rolls, 2014; Rushworth, Kolling, Sallet, & Mars, 2012).

We have also been able to obtain evidence that non-reward used as a signal to reverse behavioural choice is represented in the human orbitofrontal cortex. Kringelbach and Rolls (2003) used the faces of two different people, and if one face was selected then that face smiled, and if the other was selected, the face showed an angry expression. After good performance was acquired, there were repeated reversals of the visual discrimination task. Kringelbach and Rolls (2003) found that activation of a lateral part of the orbitofrontal cortex in the fMRI study was produced on the error trials, that is when the human chose a face, and did not obtain the expected reward. Control tasks showed that the response was related to the error, and the mismatch between what was expected and what was obtained as the reward outcome, in that just showing an angry face expression did not selectively activate this part of the lateral orbitofrontal cortex. An interesting aspect of this study that makes it relevant to human social behaviour is that the conditioned stimuli were faces of particular individuals, and the unconditioned stimuli were face expressions. Moreover, the study reveals that the human orbitofrontal cortex is very sensitive to social feedback when it must be used to change behaviour (Kringelbach & Rolls, 2003, 2004). Correspondingly, it has now been shown in macaques using fMRI that the lateral orbitofrontal cortex is activated by non-reward during reversal (Chau et al., 2015).

The non-reward neurons in the orbitofrontal cortex are implicated in changing behaviour when non-reward occurs. Monkeys with orbitofrontal cortex damage are impaired on Go/NoGo task performance, in that they Go on the NoGo trials (Iversen & Mishkin, 1970), and in an object-reversal task in that they respond to the object that was formerly rewarded with food, and in extinction in that they continue to respond to an object that is no longer rewarded (Butter, 1969; Jones & Mishkin, 1972; Meunier, Bachevalier, & Mishkin, 1997). The visual discrimination reversal learning deficit shown by monkeys with orbitofrontal cortex damage (Baylis & Gaffan, 1991; Jones & Mishkin, 1972; Murray & Izquierdo, 2007) may be due at least in part to the tendency of these monkeys not to withhold responses to non-rewarded stimuli (Jones & Mishkin, 1972) including objects that were previously rewarded during reversal (Rudebeck & Murray, 2011), and including foods that are not normally accepted (Baylis & Gaffan, 1991; Butter, McDonald, & Snyder, 1969). Consistently, orbitofrontal cortex (but not amygdala) lesions impaired instrumental extinction (Murray & Izquierdo, 2007). Similarly, in humans, patients with ventral frontal lesions made more errors in the reversal (or in a similar extinction) task, and completed fewer reversals, than control patients with damage elsewhere in the frontal lobes or in other brain regions (Rolls, Hornak, Wade, & McGrath, 1994), and continued to respond to the previously rewarded stimulus when it was no longer rewarded. A reversal deficit in a similar task in patients with ventromedial frontal cortex damage was also reported by Fellows and Farah (2003). Further, in patients with well localised surgical lesions of the orbitofrontal cortex (made to treat epilepsy, tumours, etc.), it was found that they were severely impaired at the reversal task, in that they accumulated less money (Hornak et al., 2004). These patients often

failed to switch their choice of stimulus after a large loss; and often did switch their choice even though they had just received a reward, and this has been quantified in a more recent study (Berlin et al., 2004). The importance of the failure to rapidly learn about the value of stimuli from negative feedback has also been described as a critical difficulty for patients with orbitofrontal cortex lesions (Fellows, 2007, 2011; Wheeler & Fellows, 2008), and has been contrasted with the effects of lesions to the anterior cingulate cortex which impair the use of feedback to learn about actions (Camille, Tsuchida, & Fellows, 2011; Fellows, 2011; Rolls, 2014).

Dopamine neurons, some of which appear to respond to positive reward prediction error (when a reward outcome is larger than expected) (Schultz, 2013), do not provide the answer to how non-reward is computed for a number of reasons (Rolls, 2014). First, they may reflect only weakly any non-reward, by having somewhat reduced firing rates below their already low firing rates. Second, there are no known computations that could be performed by dopamine neurons based on expected reward and reward outcome signals, which are not known to be represented in the midbrain region where the dopamine neurons are represented. Instead, it is suggested that the dopamine neurons reflect at least in part processing of expected reward value, outcome reward value, and negative reward prediction error performed in the orbitofrontal cortex (Rolls, 2014). Consistent with this, the orbitofrontal cortex does project to the ventral striatum in which neurons of the type present in the orbitofrontal cortex are found (Rolls, Thorpe, & Maddison, 1983; Rolls & Williams, 1987; Williams, Rolls, Leonard, & Stern, 1993), and which then projects to the midbrain area where the dopamine neuron cell bodies are located (Rolls, 2014).

Given these findings, the aim of the current investigation was to produce hypotheses and a model that could account for the non-reward firing of orbitofrontal cortex neurons, as analyzed by Thorpe et al. (1983). The particular neuronal responses to model were how the non-reward neurons are triggered into a high firing rate with continuing activity, if no reward outcome is received after an expected reward has been indicated. The non-reward neurons should also be activated if after an expected reward signal, a punishment is received. The non-reward neurons should not be triggered if after an expected reward signal, a reward outcome is received (Thorpe et al., 1983). The hypothesis and model utilize the properties of expected reward and reward outcome neurons in the orbitofrontal cortex, as well as the responses of non-reward neurons in the orbitofrontal cortex, as just described.

3. Methods

3.1. Outline of the hypothesis and model

A network with two competing attractor subpopulations of neurons, one for reward, and one for non-reward, is proposed. The single attractor network has one population (or pool) of neurons that is activated by expected reward, and maintains its firing until, after a time, synaptic depression reduces the firing rate in this reward neuronal population. A second population of neurons encodes non-reward, and is triggered into

high firing rate continuing activity because of reduced inhibition from the common inhibitory neurons if no reward outcome is received to keep the reward attractor population active. The emergence of activity when it is not being inhibited in the non-reward neurons is facilitated by the stochastic spiking-time-related noise in the system (Rolls & Deco, 2010). This accounts for non-reward signalling in extinction, if an expected reward is not followed by a reward outcome. If a reward outcome is received, this keeps the reward attractor active, and this through the inhibitory neurons prevents the non-reward attractor neurons from being activated. If an expected reward has been signalled, and the reward attractor is active, the reward attractor firing can be directly inhibited by a non-reward outcome, and the non-reward neurons, which signal negative reward prediction error, are activated.

3.2. Attractor framework

Our aim is to investigate these mechanisms in a biophysically realistic attractor framework, so that the properties of receptors, synaptic currents and the statistical effects related to the probabilistic spiking of the neurons can be part of the model. We use a minimal architecture, a single attractor or autoassociation network (Amit, 1989; Hertz, Krogh, & Palmer, 1991; Hopfield, 1982; Rolls, 2008; Rolls & Deco, 2002; Rolls & Treves, 1998). We chose a recurrent (attractor) integrate-and-fire network model which includes synaptic channels for AMPA, NMDA and GABA_A receptors (Brunel and Wang 2001).

Both excitatory and inhibitory neurons are represented by a leaky integrate-and-fire model (Tuckwell, 1988). The basic state variable of a single model neuron is the membrane potential. It decays in time when the neurons receive no synaptic input down to a resting potential. When synaptic input causes the membrane potential to reach a threshold, a spike is emitted and the neuron is set to the reset potential at which it is kept for the refractory period. The emitted action potential is propagated to the other neurons in the network. The excitatory neurons transmit their action potentials via the glutamatergic receptors AMPA and NMDA which are both modelled by their effect in producing exponentially decaying currents in the postsynaptic neuron. The rise time of the AMPA current is neglected, because it is typically very short. The NMDA channel is modelled with an alpha function including both a rise and a decay term. In addition, the synaptic function of the NMDA current includes a voltage dependence controlled by the extracellular magnesium concentration (Jahr & Stevens, 1990). The inhibitory postsynaptic potential is mediated by a GABA_A receptor model and is described by a decay term.

The single attractor network contains 800 excitatory and 200 inhibitory neurons, which is consistent with the observed proportions of pyramidal cells and interneurons in the cerebral cortex (Abeles, 1991; Braitenberg & Schütz, 1991). The connection strengths are adjusted using mean-field analysis (Brunel and Wang 2001; Rolls & Deco, 2010), so that the excitatory and inhibitory neurons exhibit a spontaneous activity of 3 Hz and 9 Hz, respectively (Koch & Fuster, 1989; Wilson, O'Scalaidhe, & Goldman-Rakic, 1994). The recurrent excitation mediated by the AMPA and NMDA receptors is

dominated by the NMDA current to avoid instabilities during the delay periods (Wang, 2002).

Our cortical network model features a minimal architecture to investigate stability of recalled memories in a short-term memory period, and consists of two selective pools, a reward population (S1) and a non-reward population (S2) (Fig. 2). We use just two selective pools to eliminate possible disturbing factors. The non-selective pool NS models the spiking of cortical neurons and serves to generate an approximately Poisson spiking dynamics in the model (Brunel & Wang, 2001), which is what is observed in the cortex. The inhibitory pool IH contains the 200 inhibitory neurons. The connection weights between the neurons within each selective pool or population are called the intra-pool connection strengths w_+ . The increased strength of the intra-pool connections is counterbalanced by the other excitatory connections (w_-) to keep the average input constant.

The network receives Poisson input spikes via AMPA receptors which are envisioned to originate from 800 external neurons at an average spontaneous firing rate of 3 Hz from each external neuron, consistent with the spontaneous activity observed in the cerebral cortex (Rolls, 2008; Rolls & Treves, 1998; Wilson et al. 1994). Given that there are 800 synapses on each neuron in the network for these external

inputs, the number of spikes being received by every neurons in the network is 2400 spikes/s. This external input can be altered for specific neuronal populations to introduce the effects of external stimuli on the network. In addition to these external excitatory inputs, each excitatory neuron receives 80 excitatory inputs from other neurons in the same population in which the firing rate is modulated by w_+ , and 720 excitatory inputs from other excitatory neurons (given that there are 800 excitatory neurons) in which the firing rate is modulated by w_- (see Fig. 2). A detailed mathematical description is provided in the [Supplementary Material](#).

3.3. Short term synaptic depression

A synaptic depression mechanism was used following (Dayan & Abbott, 2001, p. 185). In particular, the probability of transmitter release P_{rel} was decreased after each presynaptic spike by a factor $P_{\text{rel}} = P_{\text{rel}} \cdot f_D$ with $f_D = .988$. Between presynaptic action potentials the release probability P_{rel} is updated by

$$\tau_p \frac{dP_{\text{rel}}}{dt} = P_0 - P_{\text{rel}} \quad (1)$$

with $P_0 = 1$ and $\tau_p = 1000$ msec.

3.4. Analysis

Simulations were performed with the following protocols. Reference to Figs. 3–5, especially the event part at the bottom of each Figure, may be helpful.

3.4.1. Non-reward: Extinction

First there was a period of spontaneous firing from 0 to 500 msec in which the input to the reward and non-reward pools was 2.90 spikes/s per synapse. (There were 800 external synapses onto each neuron, making the mean Poisson firing rate input to each neuron in these pools correspond to 2320 spikes/s.) This 500 msec period corresponded to the inter-trial period.

At 500 msec the external input to the reward pool was increased to 3.1 spikes/s per synapse, to reflect an expected reward input. (The expected reward input might correspond to the sight of food or of a visual stimulus currently associated with the sight of food.) This was maintained at this level until the end of the trial at 5000 msec. Because no reward outcome (such as a taste of food in the mouth) was received for the rest of the trial, this was an Extinction trial, in which no reward outcome was delivered. Systematic parameter exploration including mean field analyses showed that a suitable value for w_+ was 2.1 for the reward population, for this enabled a high firing rate attractor state to be maintained by that population in the absence of synaptic depression even when the activating stimulus was removed and the external input to the reward attractor was reset to the default value of 3.0 spikes/s per synapse (Deco & Rolls, 2006; Rolls & Deco, 2010). That choice defined w_- as .88, as described above. Systematic parameter exploration showed that a value for w_+ of 2.22 for the non-reward population was in the region where this population of neurons would emerge from its spontaneous firing rate under the influence of the spiking-related noise in the system into a high firing rate attractor

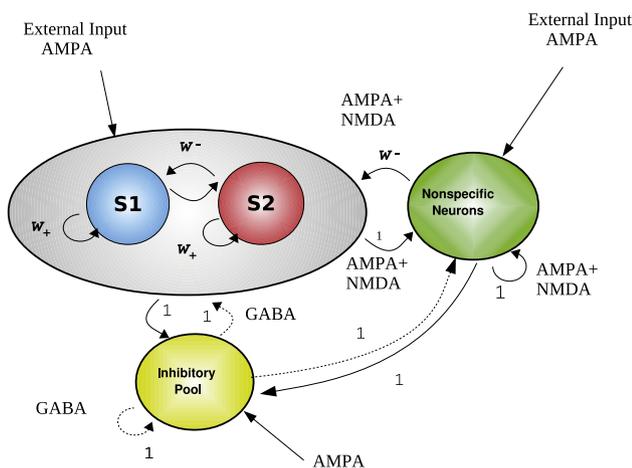


Fig. 2 – The attractor network model. The excitatory neurons are divided into two selective pools S1 (termed the reward attractor population of neurons) and S2 (termed the non-reward attractor population of neurons) (with 80 neurons each) with strong intra-pool connection strengths w_+ and one non-selective pool (NS) (with 640 neurons). The value of w_+ for S2 is a little higher than that for S1, so that S2 tends to emerge into a high firing rate if activity is not maintained in S1 after an expected reward input has been received (see text and [Supplementary Material](#)). The other connection strengths are 1 or weak w_- . The network contains 1000 neurons, of which 800 are in the excitatory pools and 200 are in the inhibitory pool IH. Every neuron in the network also receives inputs from 800 external neurons, and these neurons increase their firing rates to apply a stimulus to one of the pools S1 or S2. The [Supplementary Material](#) contains the synaptic connection matrices.

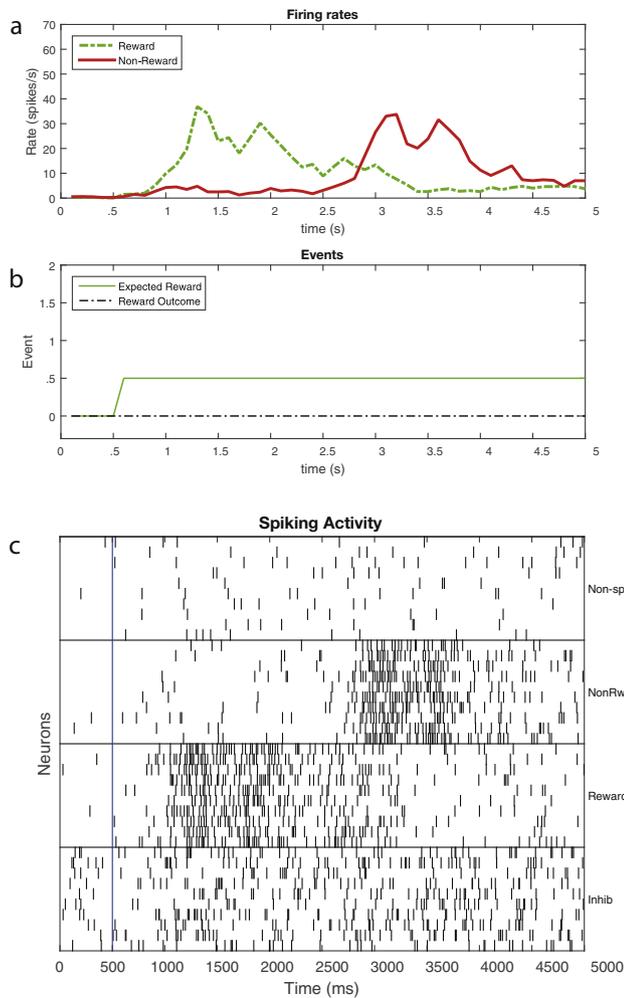


Fig. 3 – The operation of the network when an expected reward is not obtained (extinction). a. The firing rates of the reward population of neurons and the non-reward population of neurons during a 5 sec trial. b. After a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to the reward attractor population of neurons (3.1 Hz per synapse), and maintained at that level for the remainder of the trial. No reward outcome input was received. The input to the non-reward population of neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial. (Note: there were 800 external synapses onto each neuron through which the external inputs were applied with the firing rates specified for each synapse.) c. Rastergrams showing for each of the four populations of neurons, non-specific, non-reward, reward, and inhibitory, the spiking of ten neurons chosen at random from each population. Each small vertical line represents a spike from a neuron. Each horizontal row shows the spikes of one neuron. The different neurons are from the same trial.

state if the reward neurons were not firing. The two values of w_+ became fixed parameters that were not altered in any of the simulations. For these extinction simulations, at 500 msec the input to the Non-reward population of neurons

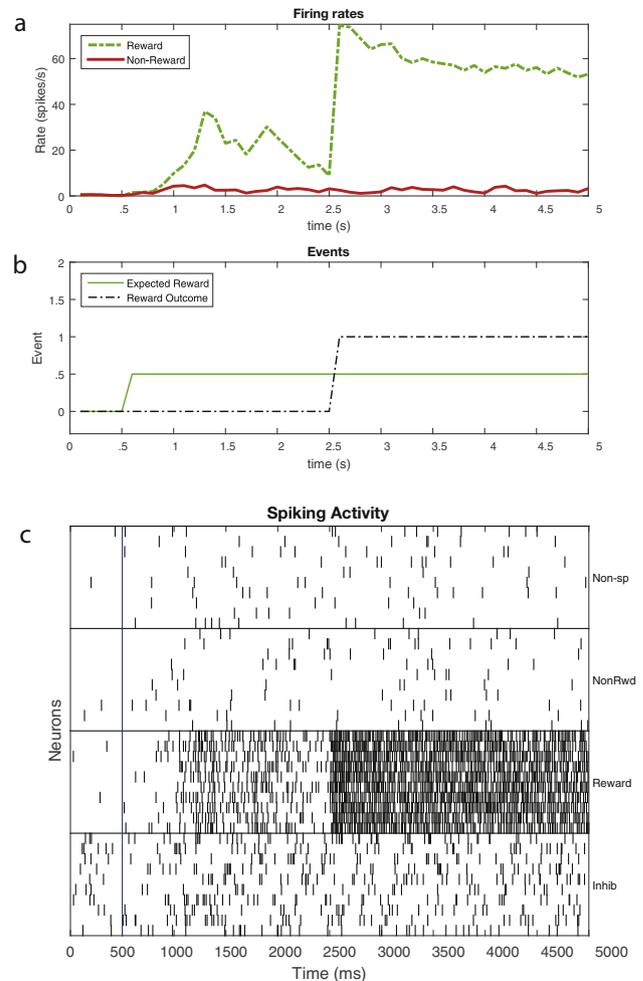


Fig. 4 – The operation of the network when an expected reward is followed by a reward outcome. a. The firing rates of the reward population of neurons and the non-reward population of neurons during a 5 sec trial. b. After a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to the reward attractor population of neurons (3.1 Hz per external synapse). At 2500 msec a reward outcome input was received (3.7 Hz per synapse), and this was maintained until the end of the trial. The input to the non-reward population of neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial. c. Rastergrams showing for each of the four populations of neurons, non-specific, non-reward, reward, and inhibitory, the spiking of ten neurons chosen at random from each population. Each small vertical line represents a spike from a neuron. Each horizontal row shows the spikes of one neuron.

was increased to 3.05 spikes/s per synapse, and was maintained constant for the rest of the trial, that is, until 5500 msec.

3.4.2. Expected reward followed by a reward

In this scenario (Fig. 4), after a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to

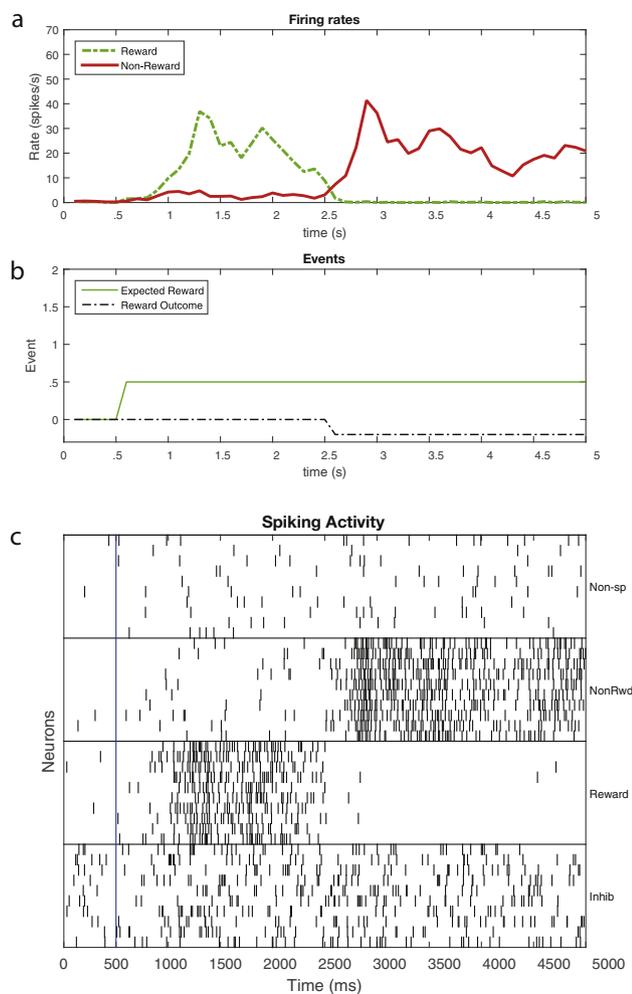


Fig. 5 – The operation of the network when an expected reward is followed by a punishment. **a.** The firing rates of the reward population of neurons and the non-reward population of neurons during a 5 sec trial. **b.** After a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to the reward attractor population of neurons (3.1 Hz per synapse). At 2500 msec a punishment was received, and this decreased the expected reward input to the reward population to a low value (2.8 Hz per synapse), and this was maintained until the end of the trial. The input to the non-reward population of neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial. **c.** Rastergrams showing for each of the four populations of neurons, non-specific, non-reward, reward, and inhibitory, the spiking of ten neurons chosen at random from each population. Each small vertical line represents a spike from a neuron. Each horizontal row shows the spikes of one neuron.

the reward attractor population of neurons at 3.1 Hz per synapse. At 2500 msec a reward outcome input (corresponding for example to the delivery of a taste reward) was received (3.7 Hz per external synapse, of which there were 800), and this was maintained until the end of the trial. As in the other simulations, the input to the non-reward population of

neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial.

3.4.3. Expected reward followed by punishment

In this scenario (Fig. 5), after a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to the reward attractor population of neurons (at 3.1 Hz per synapse). At 2500 msec a punishment was received (corresponding for example to the delivery of an aversive saline taste instead of the reward of a good taste), and this decreased the expected reward input to the reward population to a low value (2.8 Hz per synapse), and this was maintained until the end of the trial. The input to the non-reward population of neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial.

4. Results

4.1. Expected reward followed by non-reward: Extinction

Fig. 3 shows the firing rates of the populations of neurons when an expected reward input is not followed by a reward outcome input, that is, in extinction. w_+ was 2.1 for the reward population, and 2.22 for the non-reward population to make it more excitable, and w_- was .88. The expected reward input started at .5 sec and was maintained at a value of 3.1 Hz per synapse until the end of the trial. The external input to the non-reward population also increased at the end of the spontaneous period to reflect the fact that a trial was in progress, but its value was 3.05 Hz per synapse as this was an expected reward trial, so that the expected reward had to be greater than any expected non-reward. Fig. 3 shows that the reward population increased its firing rate peaking at approximately 1300 msec (800 msec after the expected reward input started), and then decreased to a low firing rate of approximately 15 spikes/s by 2500 msec because of the synaptic depression that is implemented in the recurrent collateral inputs (and not in the external inputs to the network). At about this time, 2500 msec, the non-reward population started to increase its firing rate, because it no longer was receiving inhibition through the inhibitory neurons from the reward population. The non-reward population started firing when the inhibition on it was released because it had a slightly higher than the typical value for w_+ (i.e., it is 2.22), and because of the spiking-timing-related noise in the system. The Non-reward population maintained its activity for approximately 2 sec before synaptic depression in its recurrent collateral connections reduced its firing rate back to baseline.

This simulation thus shows how a non-reward population of neurons can be triggered into firing when an expected reward is not followed by a reward outcome.

4.2. Expected reward followed by a reward outcome

The operation of the network when an expected reward is followed by a reward outcome is illustrated in Fig. 4. After a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to the reward attractor population

of neurons (3.1 Hz per synapse). At 2500 msec a reward outcome input was received (3.7 Hz per external synapse), and this was maintained until the end of the trial. The input to the non-reward population of neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial. Fig. 4 shows that the effect of the reward outcome was to increase the firing of the reward population (even though synaptic depression was present in the recurrent collateral synapses). The high firing rate of the reward population of neurons produced sufficient inhibition via the inhibitory neurons on the non-reward population of neurons that they fired only at a low rate of approximately 3 spikes/s.

This simulation thus shows how the non-reward population of neurons is not triggered into a high firing when an expected reward is followed by a reward outcome.

4.3. Expected reward followed by punishment outcome activates the non-reward population

The operation of the network when an expected reward is followed by a punishment outcome is illustrated in Fig. 5. After a period of spontaneous activity from 0 until .5 sec, the expected reward input was applied to the reward attractor population of neurons (at 3.1 Hz per synapse). At 2500 msec a punishment was received, and this decreased the expected reward input to the reward population to a low value (2.8 Hz per synapse), and this was maintained until the end of the trial. The input to the non-reward population of neurons was maintained constant at 3.05 Hz per synapse from time = .5 sec until the end of the trial. Fig. 5 shows that when the expected reward input to the reward population is decreased at 2500 msec by the receipt of punishment, the release of inhibition through the inhibitory neurons caused the non-reward population to emerge into a high firing rate state reflecting the non-reward contingency.

This simulation thus shows how the non-reward population of neurons is triggered into a high firing when an expected reward is followed by a punishment outcome.

5. Discussion

The simulations show how non-reward neurons could be activated when an expected reward is followed by non-reward in the form of no reward outcome, or a punishment outcome. The mechanism implemented in the simulations is the release of inhibition on a non-reward population of neurons when synaptic depression in the recurrent collateral connections decreases the firing of the reward population of neurons, and the firing of the non-reward population is not maintained by the receipt of a reward outcome input to maintain the reward neuronal firing. The simulations are valuable in showing the parametric conditions under which this process can be implemented. In the simulations, the speed of the synaptic depression, and the time before the non-reward population become active, can be influenced by for example the time constant of the synaptic depression. It is useful to note that the same synaptic depression operated in all the excitatory connections of the neurons in the model, and is not a special property of the reward neurons.

Other mechanisms than the synaptic depression used to illustrate the concept of how non-reward neuronal firing is produced can be envisaged. For example, if the trials typically involved a delay of several seconds before the reward outcome was received, then a separate short-term memory set to estimate the typical delay before reward outcome might be used to decrease the expected reward input to the network after the appropriate delay, allowing the non-reward neurons to then fire. What is a key concept in the mechanism proposed is however the balance between the reward and the non-reward attractor populations, which are effectively in competition with each other via the common population of inhibitory neurons. This helps to account for the neurophysiological evidence that if a food reward is moved towards a macaque, then the expected reward neurons in the primate orbitofrontal cortex fire fast, and as soon as the reward stops moving (within sight) towards the macaque, or is slowly removed (by being moved backwards), the non-reward neurons start to fire, and the reward neurons stop firing (Rolls et al., 1983). The competing attractor hypothesis for reward and non-reward representations in the orbitofrontal cortex described here provides an account for this important property of the populations of neurons in the primate orbitofrontal cortex. Consistent with this competing attractor hypothesis, there is evidence for reciprocal activity in the human medial orbitofrontal cortex reward area and the lateral orbitofrontal cortex non-reward/loss area (O'Doherty et al., 2001). In particular, activations in the human medial orbitofrontal increase in proportion to the amount of (monetary) reward received, and decrease in proportion to the amount of (monetary) loss; and in the lateral orbitofrontal cortex, the activations are reciprocally related to these changes (O'Doherty et al. 2001).

The hypothesis that there are timing mechanisms in the prefrontal cortex that do time how long to wait for an expected reward is supported by our finding that patients with orbitofrontal cortex damage and patients with borderline personality disorder not only are impulsive with reduced sensitivity to non-reward, but have a faster perception of time (they under-produce a time interval) (Berlin et al., 2005; Rolls, 2014). We have hypothesized that this faster perception of time may indeed be part of the mechanism by which some patients behave with what is described as impulsivity, because what they expected to happen has not happened yet.

Non-reward neurons of the type described here may be very important in changing behaviour (Rolls, 2014). For example, in a visual discrimination task in which one stimulus is rewarded (for example a triangle, which if selected leads to the delivery of a sweet taste), and the other is punished (for example a square, which if selected leads to the delivery of aversive salt taste), then the selection switches if the rule reverses, and selection of the triangle leads to a salt taste. This reversal can be very fast, in one trial, and can be non-associative in that after the selection of the triangle led to the delivery of saline, the next choice of the square was to select it, even though its previous presentation had been associated with salt taste (Rolls, 2014; Thorpe et al., 1983). A model for this rule-based reversal has been described (Deco & Rolls, 2005c), and the non-reward signal needed for that model could be produced by the neuronal network mechanism described here.

The present model, although based on our previous research on attractor based biologically plausible dynamical modelling of computations performed by the cortex (Battaglia & Treves, 1998; Deco, Ledberg, Almeida, & Fuster, 2005; Deco, Rolls, & Romo, 2010; Deco & Rolls, 2005a, 2005b, 2005c, 2006; Panzeri, Rolls, Battaglia, & Lavis, 2001; Rolls, 2016a; Rolls & Deco, 2010, 2011, 2015a, 2015b; Rolls, Dempere-Marco, & Deco, 2013; Rolls, Grabenhorst, & Deco, 2010; Rolls, Webb, & Deco, 2012; Treves, 1993), is different from earlier models by what is being modelled (the generation of a non-reward signal), by the brain region modelled (the orbitofrontal cortex, and in particular the lateral orbitofrontal cortex), and by the architecture in which the reward attractor network is activated by an expected reward and inhibits a non-reward attractor network, which latter only becomes active if the habituating reward attractor is not maintained in its activity by the reward outcome.

The implementation of the model described here used non-overlapping populations of neurons for the reward and non-reward attractors in the single network, for clarity and simplicity. However separate attractors with partially overlapping neurons in the different populations using a sparse distributed representation more like that found in the cerebral cortex (Rolls, 2016a; Rolls & Treves, 2011) do operate with similar properties (Battaglia & Treves, 1998; Rolls, 2012; Rolls, Treves, Foster, & Perez-Vicente, 1997).

Although the system described here was modelled with networks that have realistic dynamics as they move into and out of stable attractor states (Battaglia & Treves, 1998; Deco & Rolls, 2006; Deco et al. 2005; Panzeri et al. 2001; Rolls, 2016a; Rolls & Deco, 2010; Treves, 1993) another approach to such modelling is a network that when the system does not have sufficient time to fall into a stable attractor state, nevertheless may visit a succession of states in a stable trajectory (Rabinovich, Huerta, & Laurent, 2008; Rabinovich, Varona, Tristan, & Afraimovich, 2014). A possible way to understand the operation of such a system is that of a stable heteroclinic channel. A stable heteroclinic channel is defined by a sequence of successive metastable ('saddle') states. Under the proper conditions, all the trajectories in the neighbourhood of these saddle points remain in the channel, ensuring robustness and reproducibility of the dynamical trajectory through the state space (Rabinovich et al. 2008, 2014). This approach might be considered in future as a somewhat different implementation of the system described here.

A model of a different process, action–outcome learning, for a different part of the brain, the anterior cingulate cortex, has in common with the present model that an error is signalled if a prediction is not inhibited by the expected outcome (Alexander & Brown, 2011). However, the implementation as well as the purpose of that model is quite different, for no biologically plausible networks with attractor dynamics are modelled as here, but instead that is a mathematical model based on reinforcement learning algorithms.

In conclusion, an approach has been described to how the orbitofrontal cortex may compute non-reward from expected value and reward outcome. A key attribute of the new approach described here is that there are competing attractors in the orbitofrontal cortex for reward and non-reward, and that these have reciprocal activity because of the competition

between them. This can account for the neurophysiological findings that the firing rates of reward and non-reward neurons in the primate orbitofrontal cortex are dynamically reciprocally related to each other (Rolls, 2014; Thorpe et al., 1983). This is seen not only in the activity of different single neurons, some of which respond as a reward approaches, and others of which respond when the reward stops approaching (Thorpe et al., 1983), but also at the fMRI level, where activations in the medial orbitofrontal cortex increase linearly with the amount of money gained, and activations in the lateral orbitofrontal cortex increase linearly with the amount of money lost (O'Doherty et al., 2001). Separate populations of reward and non-reward neurons may be needed because in the cerebral cortex the high firing rates of neurons are usually the drivers of outputs (Rolls, 2016a; Rolls & Treves, 2011), and different output systems must be activated to lead to the different behavioural effects of receiving an expected reward, and of not receiving an expected reward. We know of no other hypothesis and model that accounts for the neurophysiological findings described here. The process that we analyze here is fundamentally important in changing behaviour when an expected reward is not received. The process is important in many emotions, which can be produced if an expected reward is not received or lost, with examples including sadness and anger (Rolls, 2014). Consistent with this the psychiatric disorder of depression may arise when the non-reward system leading to sadness is too sensitive, or maintains its activity for too long (Rolls, 2016b), or has increased functional connectivity (Cheng et al., 2016). Conversely, if the non-reward system is underactive or is damaged by lesions of the orbitofrontal cortex, the decreased sensitivity to non-reward may contribute to increased impulsivity (Berlin et al., 2005; Berlin et al., 2004), and even antisocial and psychopathic behaviour (Rolls, 2014). For these reasons, understanding the mechanisms that underlie non-reward is important not only for understanding normal human behaviour and how it changes when rewards are not received, but also for understanding and potentially treating better some emotional and psychiatric disorders. Indeed, a prediction of the current model is that if there are non-reward attractors in the orbitofrontal cortex, and these are over-sensitive in depression, then treatments such as ketamine which by blocking NMDA receptors knock the network out of its attractor state, may in this way reduce depression, and there is already evidence consistent with this (Carlson et al., 2013; Lally et al., 2015; Rolls, 2016b).

Acknowledgements

This research was supported by the Oxford Centre for Computational Neuroscience, Oxford, UK. GD was supported by the ERC Advanced Grant: DYSTRUCTURE (no. 295129), and by the Spanish Research Project SAF2010-16085.

Supplementary data

Supplementary data related to this article can be found at <http://dx.doi.org/10.1016/j.cortex.2016.06.023>.

REFERENCES

- Abeles, A. (1991). *Corticonics*. New York: Cambridge University Press.
- Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action–outcome predictor. *Nature Neuroscience*, 14, 1338–1344.
- Amit, D. J. (1989). *Modeling brain function. The world of attractor neural networks*. Cambridge: Cambridge University Press.
- Battaglia, F., & Treves, A. (1998). Stable and rapid recurrent processing in realistic autoassociative memories. *Neural Computation*, 10, 431–450.
- Baylis, L. L., & Gaffan, D. (1991). Amygdalectomy and ventromedial prefrontal ablation produce similar deficits in food choice and in simple object discrimination learning for an unseen reward. *Experimental Brain Research*, 86, 617–622.
- Berlin, H., Rolls, E. T., & Iversen, S. D. (2005). Borderline personality disorder, impulsivity, and the orbitofrontal cortex. *American Journal of Psychiatry*, 162, 234–245.
- Berlin, H., Rolls, E. T., & Kischka, U. (2004). Impulsivity, time perception, emotion, and reinforcement sensitivity in patients with orbitofrontal cortex lesions. *Brain*, 127, 1108–1126.
- Braitenberg, V., & Schütz, A. (1991). *Anatomy of the cortex*. Berlin: Springer-Verlag.
- Brunel, N., & Wang, X. J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11, 63–85.
- Butter, C. M. (1969). Perseveration in extinction and in discrimination reversal tasks following selective prefrontal ablations in *Macaca mulatta*. *Physiology & Behavior*, 4, 163–171.
- Butter, C. M., McDonald, J. A., & Snyder, D. R. (1969). Orality, preference behavior, and reinforcement value of non-food objects in monkeys with orbital frontal lesions. *Science*, 164, 1306–1307.
- Camille, N., Tsuchida, A., & Fellows, L. K. (2011). Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *Journal of Neuroscience*, 31, 15048–15052.
- Carlson, P. J., Diazgranados, N., Nugent, A. C., Ibrahim, L., Luckenbaugh, D. A., Brutsche, N., et al. (2013). Neural correlates of rapid antidepressant response to ketamine in treatment-resistant unipolar depression: A preliminary positron emission tomography study. *Biological Psychiatry*, 73, 1213–1221.
- Chau, B. K., Sallet, J., Papageorgiou, G. K., Noonan, M. P., Bell, A. H., Walton, M. E., et al. (2015). Contrasting roles for orbitofrontal cortex and amygdala in credit assignment and learning in macaques. *Neuron*, 87, 1106–1118.
- Cheng, W., Rolls, E. T., Qiu, J., Liu, W., Tang, Y., Huang, C.-C., et al. (2016). Medial reward and lateral non-reward orbitofrontal cortex circuits change in opposite directions in depression. *Brain*, in revision.
- Critchley, H. D., & Rolls, E. T. (1996a). Hunger and satiety modify the responses of olfactory and visual neurons in the primate orbitofrontal cortex. *Journal of Neurophysiology*, 75, 1673–1686.
- Critchley, H. D., & Rolls, E. T. (1996b). Olfactory neuronal responses in the primate orbitofrontal cortex: Analysis in an olfactory discrimination task. *Journal of Neurophysiology*, 75, 1659–1672.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press.
- Deco, G., Ledberg, A., Almeida, R., & Fuster, J. (2005). Neural dynamics of cross-modal and cross-temporal associations. *Experimental Brain Research*, 166, 325–336.
- Deco, G., & Rolls, E. T. (2005a). Neurodynamics of biased competition and cooperation for attention: A model with spiking neurons. *Journal of Neurophysiology*, 94, 295–313.
- Deco, G., & Rolls, E. T. (2005b). Sequential memory: A putative neural and synaptic dynamical mechanism. *Journal of Cognitive Neuroscience*, 17, 294–307.
- Deco, G., & Rolls, E. T. (2005c). Synaptic and spiking dynamics underlying reward reversal in the orbitofrontal cortex. *Cerebral Cortex*, 15, 15–30.
- Deco, G., & Rolls, E. T. (2006). A neurophysiological model of decision-making and Weber's law. *European Journal of Neuroscience*, 24, 901–916.
- Deco, G., Rolls, E. T., & Romo, R. (2010). Synaptic dynamics and decision-making. *Proceedings of the National Academy of Sciences of the United States of America*, 107, 7545–7549.
- Fellows, L. K. (2007). The role of orbitofrontal cortex in decision making: A component process account. *Annals of the New York Academy of Sciences*, 1121, 421–430.
- Fellows, L. K. (2011). Orbitofrontal contributions to value-based decision making: Evidence from humans with frontal lobe damage. *Annals of the New York Academy of Sciences*, 1239, 51–58.
- Fellows, L. K., & Farah, M. J. (2003). Ventromedial frontal cortex mediates affective shifting in humans: Evidence from a reversal learning paradigm. *Brain*, 126, 1830–1837.
- Grabenhorst, F., & Rolls, E. T. (2011). Value, pleasure, and choice systems in the ventral prefrontal cortex. *Trends in Cognitive Sciences*, 15, 56–67.
- Grabenhorst, F., Rolls, E. T., & Bilderbeck, A. (2008). How cognition modulates affective responses to taste and flavor: Top-down influences on the orbitofrontal and pregenual cingulate cortices. *Cerebral Cortex*, 18, 1549–1559.
- Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation*. Wokingham, UK: Addison Wesley.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79, 2554–2558.
- Hornak, J., O'Doherty, J., Bramham, J., Rolls, E. T., Morris, R. G., Bullock, P. R., et al. (2004). Reward-related reversal learning after surgical excisions in orbitofrontal and dorsolateral prefrontal cortex in humans. *Journal of Cognitive Neuroscience*, 16, 463–478.
- Iversen, S. D., & Mishkin, M. (1970). Perseverative interference in monkey following selective lesions of the inferior prefrontal convexity. *Experimental Brain Research*, 11, 376–386.
- Jahr, C., & Stevens, C. (1990). Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *Journal of Neuroscience*, 10, 3178–3182.
- Jones, B., & Mishkin, M. (1972). Limbic lesions and the problem of stimulus–reinforcement associations. *Experimental Neurology*, 36, 362–377.
- Koch, K. W., & Fuster, J. M. (1989). Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Experimental Brain Research*, 76, 292–306.
- Kringelbach, M. L., O'Doherty, J., Rolls, E. T., & Andrews, C. (2003). Activation of the human orbitofrontal cortex to a liquid food stimulus is correlated with its subjective pleasantness. *Cerebral Cortex*, 13, 1064–1071.
- Kringelbach, M. L., & Rolls, E. T. (2003). Neural correlates of rapid reversal learning in a simple model of human social interaction. *NeuroImage*, 20, 1371–1383.
- Kringelbach, M. L., & Rolls, E. T. (2004). The functional neuroanatomy of the human orbitofrontal cortex: Evidence from neuroimaging and neuropsychology. *Progress in Neurobiology*, 72, 341–372.
- Lally, N., Nugent, A. C., Luckenbaugh, D. A., Niciu, M. J., Roiser, J. P., & Zarate, C. A., Jr. (2015). Neural correlates of

- change in major depressive disorder anhedonia following open-label ketamine. *Journal of Psychopharmacology*, 29, 596–607.
- Meunier, M., Bachevalier, J., & Mishkin, M. (1997). Effects of orbital frontal and anterior cingulate lesions on object and spatial memory in rhesus monkeys. *Neuropsychologia*, 35, 999–1015.
- Murray, E. A., & Izquierdo, A. (2007). Orbitofrontal cortex and amygdala contributions to affect and action in primates. *Annals of the New York Academy of Sciences*, 1121, 273–296.
- O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, 4, 95–102.
- Panzeri, S., Rolls, E. T., Battaglia, F., & Lavis, R. (2001). Speed of feedforward and recurrent processing in multilayer networks of integrate-and-fire neurons. *Network: Computation in Neural Systems*, 12, 423–440.
- Passingham, R. E. P., & Wise, S. P. (2012). *The neurobiology of the prefrontal cortex*. Oxford: Oxford University Press.
- Rabinovich, M., Huerta, R., & Laurent, G. (2008). Transient dynamics for neural processing. *Science*, 321, 48–50.
- Rabinovich, M. I., Varona, P., Tristan, I., & Afraimovich, V. S. (2014). Chunking dynamics: Heteroclinics in mind. *Frontiers in Computational Neuroscience*, 8, 22.
- Rolls, E. T. (1981). Processing beyond the inferior temporal visual cortex related to feeding, learning, and striatal function. In Y. Katsuki, R. Norgren, & M. Sato (Eds.), *Brain mechanisms of sensation* (pp. 241–269). New York: Wiley.
- Rolls, E. T. (2008). *Memory, attention, and decision-making. A unifying computational neuroscience approach*. Oxford: Oxford University Press.
- Rolls, E. T. (2012). Advantages of dilution in the connectivity of attractor networks in the brain. *Biologically Inspired Cognitive Architectures*, 1, 44–54.
- Rolls, E. T. (2014). *Emotion and decision-making explained*. Oxford: Oxford University Press.
- Rolls, E. T. (2015a). Functions of the anterior insula in taste, autonomic, and related functions. *Brain and Cognition*. <http://dx.doi.org/10.1016/j.bandc.2015.07.002>.
- Rolls, E. T. (2015b). Taste, olfactory, and food reward value processing in the brain. *Progress in Neurobiology*, 127–128, 64–90.
- Rolls, E. T. (2016a). *Cerebral cortex: Principles of operation*. Oxford: Oxford University Press.
- Rolls, E. T. (2016b). A non-reward attractor theory of depression. *Neuroscience and Biobehavioral Reviews*, 68, 47–58.
- Rolls, E. T. (2016c). Reward systems in the brain and nutrition. *Annual Review of Nutrition*, 36, 14.
- Rolls, E. T., & Baylis, L. L. (1994). Gustatory, olfactory and visual convergence within the primate orbitofrontal cortex. *Journal of Neuroscience*, 14, 5437–5452.
- Rolls, E. T., Critchley, H. D., Browning, A. S., Hernadi, A., & Lenard, L. (1999). Responses to the sensory properties of fat of neurons in the primate orbitofrontal cortex. *Journal of Neuroscience*, 19, 1532–1540.
- Rolls, E. T., Critchley, H. D., Mason, R., & Wakeman, E. A. (1996). Orbitofrontal cortex neurons: Role in olfactory and visual association learning. *Journal of Neurophysiology*, 75, 1970–1981.
- Rolls, E. T., & Deco, G. (2002). *Computational neuroscience of vision*. Oxford: Oxford University Press.
- Rolls, E. T., & Deco, G. (2010). *The noisy Brain: Stochastic dynamics as a principle of brain function*. Oxford: Oxford University Press.
- Rolls, E. T., & Deco, G. (2011). Prediction of decisions from noise in the brain before the evidence is provided. *Frontiers in Neuroscience*, 5, 33.
- Rolls, E. T., & Deco, G. (2015a). Networks for memory, perception, and decision-making, and beyond to how the syntax for language might be implemented in the brain. *Brain Research*, 1621, 316–334.
- Rolls, E. T., & Deco, G. (2015b). A stochastic neurodynamics approach to the changes in cognition and memory in aging. *Neurobiology of Learning and Memory*, 118, 150–161.
- Rolls, E. T., Dempere-Marco, L., & Deco, G. (2013). Holding multiple items in short term memory: A neural mechanism. *Plos One*, 8, e61078.
- Rolls, E. T., & Grabenhorst, F. (2008). The orbitofrontal cortex and beyond: From affect to decision-making. *Progress in Neurobiology*, 86, 216–244.
- Rolls, E. T., Grabenhorst, F., & Deco, G. (2010). Decision-making, errors, and confidence in the brain. *Journal of Neurophysiology*, 104, 2359–2374.
- Rolls, E. T., Hornak, J., Wade, D., & McGrath, J. (1994). Emotion-related learning in patients with social and emotional changes associated with frontal lobe damage. *Journal of Neurology, Neurosurgery and Psychiatry*, 57, 1518–1524.
- Rolls, E. T., McCabe, C., & Redoute, J. (2008). Expected value, reward outcome, and temporal difference error representations in a probabilistic decision task. *Cerebral Cortex*, 18, 652–663.
- Rolls, E. T., O'Doherty, J., Kringelbach, M. L., Francis, S., Bowtell, R., & McGlone, F. (2003). Representations of pleasant and painful touch in the human orbitofrontal and cingulate cortices. *Cerebral Cortex*, 13, 308–317.
- Rolls, E. T., Sienkiewicz, Z. J., & Yaxley, S. (1989). Hunger modulates the responses to gustatory stimuli of single neurons in the caudolateral orbitofrontal cortex of the macaque monkey. *European Journal of Neuroscience*, 1, 53–60.
- Rolls, E. T., Thorpe, S. J., & Maddison, S. P. (1983). Responses of striatal neurons in the behaving monkey. 1. Head of the caudate nucleus. *Behavioural Brain Research*, 7, 179–210.
- Rolls, E. T., & Treves, A. (1998). *Neural networks and brain function*. Oxford: Oxford University Press.
- Rolls, E. T., & Treves, A. (2011). The neuronal encoding of information in the brain. *Progress in Neurobiology*, 95, 448–490.
- Rolls, E. T., Treves, A., Foster, D., & Perez-Vicente, C. (1997). Simulation studies of the CA3 hippocampal subfield modelled as an attractor neural network. *Neural Networks*, 10, 1559–1569.
- Rolls, E. T., Verhagen, J. V., & Kadohisa, M. (2003). Representations of the texture of food in the primate orbitofrontal cortex: Neurons responding to viscosity, grittiness, and capsaicin. *Journal of Neurophysiology*, 90, 3711–3724.
- Rolls, E. T., Webb, T. J., & Deco, G. (2012). Communication before coherence. *European Journal of Neuroscience*, 36, 2689–2709.
- Rolls, E. T., & Williams, G. V. (1987). Neuronal activity in the ventral striatum of the primate. In M. B. Carpenter, & A. Jayaraman (Eds.), *The Basal Ganglia II – Structure and function – Current concepts* (pp. 349–356). New York: Plenum.
- Rolls, E. T., Yaxley, S., & Sienkiewicz, Z. J. (1990). Gustatory responses of single neurons in the orbitofrontal cortex of the macaque monkey. *Journal of Neurophysiology*, 64, 1055–1066.
- Rosenkilde, C. E., Bauer, R. H., & Fuster, J. M. (1981). Single unit activity in ventral prefrontal cortex in behaving monkeys. *Brain Research*, 209, 375–394.
- Rudebeck, P. H., & Murray, E. A. (2011). Dissociable effects of subtotal lesions within the macaque orbital prefrontal cortex on reward-guided behavior. *Journal of Neuroscience*, 31, 10569–10578.
- Rushworth, M. F., Kolling, N., Sallet, J., & Mars, R. B. (2012). Valuation and decision-making in frontal cortex: One or many serial or parallel systems? *Current Opinion in Neurobiology*, 22, 946–955.
- Schultz, W. (2013). Updating dopamine reward signals. *Current Opinion in Neurobiology*, 23, 229–238.

- Thorpe, S. J., Maddison, S., & Rolls, E. T. (1979). Single unit activity in the orbitofrontal cortex of the behaving monkey. *Neuroscience Letters*, *S3*, S77.
- Thorpe, S. J., Rolls, E. T., & Maddison, S. (1983). Neuronal activity in the orbitofrontal cortex of the behaving monkey. *Experimental Brain Research*, *49*, 93–115.
- Treves, A. (1993). Mean-field analysis of neuronal spike dynamics. *Network*, *4*, 259–284.
- Tuckwell, H. (1988). *Introduction to theoretical neurobiology*. Cambridge: Cambridge University Press.
- Verhagen, J. V., Rolls, E. T., & Kadohisa, M. (2003). Neurons in the primate orbitofrontal cortex respond to fat texture independently of viscosity. *Journal of Neurophysiology*, *90*, 1514–1525.
- Wang, X. J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, *36*, 955–968.
- Wheeler, E. Z., & Fellows, L. K. (2008). The human ventromedial frontal lobe is critical for learning from negative feedback. *Brain*, *131*, 1323–1331.
- Williams, G. V., Rolls, E. T., Leonard, C. M., & Stern, C. (1993). Neuronal responses in the ventral striatum of the behaving macaque. *Behavioural Brain Research*, *55*, 243–252.
- Wilson, F. A. W., O'Scalaidhe, S. P., & Goldman-Rakic, P. (1994). Functional synergism between putative gamma-aminobutyrate-containing neurons and pyramidal neurons in prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *91*, 4009–4013.
- Wise, S. P. (2008). Forward frontal fields: Phylogeny and fundamental function. *Trends in Neuroscience*, *31*, 599–608.