

Visual Functions of the Orbitofrontal Cortex[☆]

ET Rolls, Oxford Centre for Computational Neuroscience, Oxford, UK

© 2015 Elsevier Inc. All rights reserved.

Visual Inputs to the Orbitofrontal Cortex, and Visual Stimulus-Reinforcer Association Learning and Reversal	1
Error Neurons in the Orbitofrontal Cortex, and Visual Stimulus – Reinforcer Association Learning and Reversal	4
A Visual Representation of Faces in the Orbitofrontal Cortex	5
Effects of Damage to the Orbitofrontal Cortex on Visual Processing	6
References	7

Visual Inputs to the Orbitofrontal Cortex, and Visual Stimulus-Reinforcer Association Learning and Reversal

Rolls and colleagues have shown that there is a major visual input to many neurons in the orbitofrontal cortex, and that what is represented by these neurons is in many cases the reinforcer (reward or punisher) association of visual stimuli (Rolls and Grabenhorst, 2008; Grabenhorst and Rolls, 2011; Rolls, 2014). Many of these neurons reflect the relative preference or reward value of different visual stimuli, in that their responses decrease to zero to the sight of one food on which the monkey is being fed to satiety, but remain unchanged to the sight of other food stimuli. In this sense the visual reinforcement-related neurons predict the reward value that is available from the primary reinforcer, the taste. The visual input is from the ventral, temporal lobe, visual stream concerned with 'what' object is being seen, in that orbitofrontal visual neurons frequently respond differentially to objects or images (but depending on their reward association). The primary reinforcer that has been used is taste.

The fact that these neurons represent the reinforcer associations of visual stimuli has been shown to be the case in formal investigations of the activity of orbitofrontal cortex visual neurons, which in many cases reverse their responses to visual stimuli when the taste with which the visual stimulus is associated is reversed by the experimenter. An example of the responses of an orbitofrontal cortex neuron that reversed the visual stimulus to which it responded during reward-reversal is shown in [Figure 1](#). This reversal by orbitofrontal visual neurons can occur in as little as one trial over a few seconds.

These neurons thus reflect the information about which stimulus is currently associated with reward during reversals of visual discrimination tasks – they are reward predicting neurons. If a reversal occurs, then taste cells in the orbitofrontal cortex provide the information that an unexpected taste reinforcer has been obtained, another group of cells shows a vigorous discharge that reflects the error between the expected reward and the reward actually obtained (see below), and the visual cells with reinforcer association-related responses reverse the stimulus to which they are responsive. These neurophysiological changes take place rapidly, in as little as 5 s, and are presumed to be part of the neuronal learning mechanism that enables primates to alter their knowledge of the reinforcer association of visual stimuli so rapidly. This capacity is important whenever behaviour must be corrected when expected reinforcers are not obtained, in, for example, feeding, emotional, and social situations. In that these neurons reflect whether a visual stimulus is associated with reward or a punisher, they reflect the relative preference for different stimuli. Consistent with this evidence that the responses of some orbitofrontal cortex neurons reflect the learned predictive reward value of visual stimuli, [Thorpe et al. \(1983\)](#) found that orbitofrontal cortex neurons learned to respond differently to new stimuli that did or did not predict reward. Different neurons in the orbitofrontal cortex are tuned to different learned or conditioned reinforcers, with for example approximately 5% responding to visual stimuli associated with taste reward, and 3% to visual stimuli associated with taste punishment. These may be described as reward value neurons, for they represent the economic value of available offers, influenced for example by the quality of the reward, and the amount offered ([Grabenhorst and Rolls, 2011](#); [Padoa-Schioppa, 2011](#); [Rolls, 2014](#)).

In the visual discrimination reversal task, a second class of neuron was found that codes for particular stimuli only if they are associated with reward, and not if they are associated with punishment. Such a neuron might respond to a green stimulus associated with reward; after reversal not respond to the green stimulus when it was associated with punishment; and not respond to a blue stimulus irrespective of whether it was associated with reward or punishment ([Thorpe et al., 1983](#)). An example is shown in [Figure 2](#). They may be described as conditional visual stimulus-to-taste reward neurons, and may be important in a process by which the rule for which stimulus is treated as being currently rewarded can be reversed in one trial, as described by [Deco and Rolls \(2005\)](#).

This reversal learning found in orbitofrontal cortex neurons probably is implemented in the orbitofrontal cortex, for it does not occur one synapse earlier in the visual inferior temporal cortex, and it is in the orbitofrontal cortex that there is convergence of visual and taste pathways on to the same neurons. The way in which the ventral visual stream builds position, size and even view invariant representations of objects is described by [Rolls \(2012\)](#), and is very important, for it enables brain regions such as the

[☆]*Change History:* September 2014. E Rolls updated the text, the references and further readings to this entire article. N Narayanan and G Richerson have made minor editorial changes.

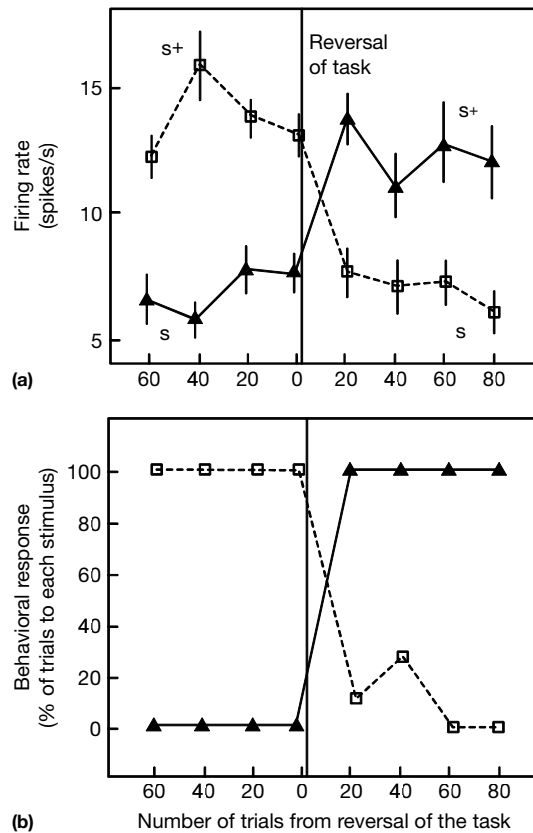


Figure 1 Orbitofrontal cortex: visual discrimination reversal. The activity of an orbitofrontal visual neuron during performance of a visual discrimination task and its reversal. The stimuli were a triangle and a square presented on a video monitor. (a) Each point represents the mean poststimulus activity in a 500 ms period of the neuron based on approximately 10 trials of the different visual stimuli. The standard errors of these responses are shown. After 60 trials of the task the reward associations of the visual stimuli were reversed. (+ indicates that a lick response to that visual stimulus produces fruit juice reward; - indicates that a lick response to that visual stimulus results in a small drop of aversive tasting saline. This neuron reversed its responses to the visual stimuli following the task reversal. (b) The behavioural response of the monkey to the task. It is shown that the monkey performs well, in that he rapidly learns to lick only to the visual stimulus associated with fruit juice reward. After Rolls, Critchley, Mason and Wakeman 1996.




Figure 2 A conditional reward neuron recorded in the orbitofrontal cortex by Thorpe et al. (1983) which responded only to the Green stimulus when it was associated with reward (G+), and not to the Blue stimulus when it was associated with Reward (B+), or to either stimuli when they were associated with a punisher, the taste of salt (G- and B-). The mean firing rate +/- the s.e.m. is shown.

orbitofrontal cortex which learn about the associations of objects to generalize to different positions, sizes, and views of the same object when they are seen later.

A possible mechanism for visual-reinforcer learning in the orbitofrontal cortex is associative (Hebbian) modification of synapses conveying visual input on to taste-responsive neurons, implementing a pattern-association network. In this model the unconditioned stimulus forcing the output neurons to respond is the (taste) primary reinforcer, and the visual conditioned stimulus becomes associated with this by associatively modifiable synapses. Such a pattern association network could in principle

unlearn the association by using associative synapses that incorporate long term depression. Although reversal might be implemented by having long-term synaptic depression (LTD) for synapses that represented the reward-associated stimulus before the reversal, and long-term potentiation (LTP) of the synapses activated by the new stimulus that after reversal is associated with reward, this would require one-trial LTP and one-trial heterosynaptic LTD to account for one-trial stimulus – reward reversal (Rolls, 2014). To implement the reversal learning very rapidly, in as little as one trial after a number of reversals when reversal learning set has been acquired, a special rule switching network in the orbitofrontal cortex may be required, and a model of this has been described (Deco and Rolls, 2005).

These studies in macaques provide evidence on the details of the neuronal representations in the primate orbitofrontal cortex which are essential for building a computational understanding for exactly what information is represented in it, how it is represented, how these differ from preceding and succeeding stages, and thus how the orbitofrontal cortex operates computationally. However, the types of visual-reward association that have been studied in primates (and confirmed as applying in humans) include objects associated with taste rewards or punishers. It has therefore been useful not only to confirm that these concepts do indeed apply to humans, but also to extend the types of visual conditioned reinforcers to quite abstract reinforcers such as monetary reward. In an fMRI study, we (O'Doherty et al., 2001) used a visual discrimination task in which one stimulus was associated with monetary reward, and a different visual stimulus with monetary loss (punishment). The actual amounts of money won on reward trials and lost on punishment trials were probabilistic. This part of the design, and the fact that unexpected visual discrimination reversals occurred so that there were trials on which money was lost, enabled us to show that the magnitude of the activation of the medial orbitofrontal cortex was correlated with the amount of money won on each trial, and the magnitude of the activation of the lateral orbitofrontal cortex was correlated with the amount of money lost on each trial, as shown in Figure 3.

Another way in which it has been shown that the visual neurons in the orbitofrontal cortex reflect the expected reward value predicted by visual stimuli is by reducing the reward value by feeding to satiety. With this sensory-specific satiety (or reward

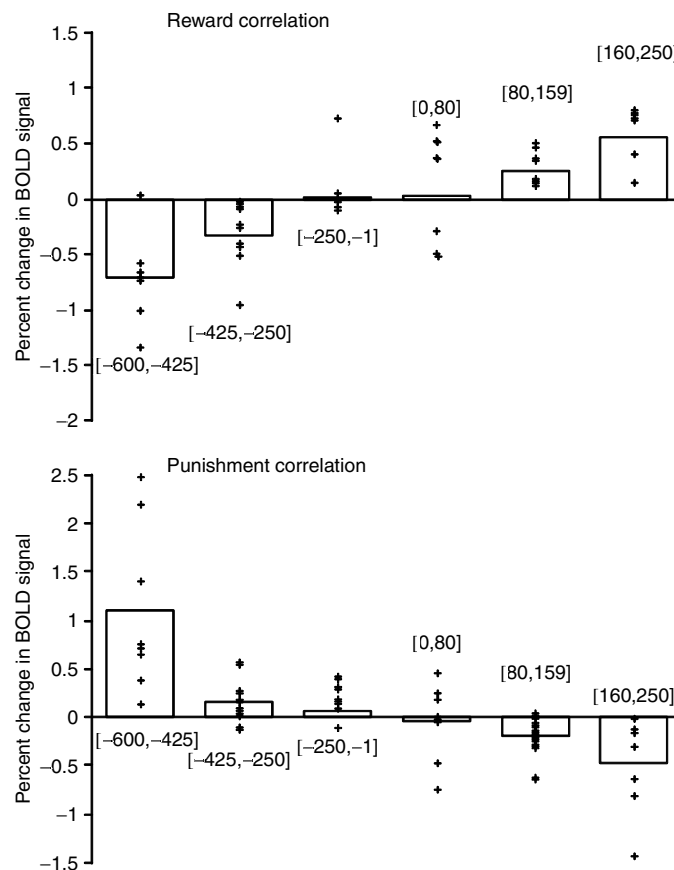


Figure 3 Correlation of brain activations with the amount of money won or lost in a visual discrimination reversal task with probabilistic monetary reward and loss. The mean percent change in BOLD signal from baseline across subjects for 6 different category ranges of reward and punishment. The signal was averaged across a category range within each subject and then the average signal change from each category was averaged across subjects. This is plotted on the right for voxels in the medial OFC that significantly correlated with reward and on the left for voxels in the lateral OFC that significantly correlated with punishment. The ranges of monetary reward and punishment in each category are shown on the chart and were determined by their relative frequencies, which follow from the experimental design. (Large positive values are large wins; large negative values are large losses.) After O'Doherty, Kringelbach, Rolls, Hornak and Andrews, 2001a.

devaluation) paradigm, it has been shown that the visual (as well as the olfactory and taste) responses of orbitofrontal cortex neurons in the macaque decrease to zero as the monkey is fed to satiety with one food, but remain unchanged to another food not eaten in the meal. In that these neurons parallel the changing preference of the monkey for the food being eaten to satiety vs the food not being eaten to satiety, they reflect the relative preference for different visual stimuli.

Error Neurons in the Orbitofrontal Cortex, and Visual Stimulus – Reinforcer Association Learning and Reversal

In addition to the neurons that encode the reward association of visual stimuli, other neurons (3.5%) in the orbitofrontal cortex detect different types of non-reward (Thorpe et al., 1983). For example, some neurons respond in extinction, immediately after a lick has been made to a visual stimulus that has previously been associated with fruit juice reward, and other neurons respond in a reversal task, immediately after the monkey has responded to the previously rewarded visual stimulus, but has obtained the punisher of salt taste rather than reward (see example in Figure 4).

Different populations of such neurons respond to other types of non-reward, including the removal of a formerly approaching taste reward, and the termination of a taste reward. The fact that different non-reward neurons respond to different types of non-reward (e.g. some to the noise of a switch that indicated that extinction of free licking for fruit juice had occurred, and others to the first presentation of a visual stimulus that was not followed by reward in a visual discrimination task) potentially enables context-specific extinction or reversal to occur. For example, the fact that fruit juice is no longer available from a lick tube does not necessarily mean that a visual discrimination task will be operating without reward. Thus the error neurons can be specific to different tasks, and this could provide a mechanism for reversal in one task to be implemented, while at the same time not reversing behaviour in another task.

The presence of these neurons is fully consistent with the hypothesis that they are part of the mechanism by which the orbitofrontal cortex enables very rapid reversal of behaviour by stimulus – reinforcer association relearning when the association of stimuli with reinforcers is altered or reversed. This information appears to be necessary for primates to rapidly alter behavioural responses when reinforcement contingencies are changed, as shown by the effects of damage to the orbitofrontal cortex (Rolls, 2014). To the extent that the firing of some dopamine neurons may reflect error signals (Schultz, 2013), one might ask where the error information comes from, given that the dopamine neurons themselves may not receive information about expected rewards (e.g. a visual stimulus associated with the sight of food), obtained rewards (e.g. taste), and would have to compute an error from these signals. On the other hand, the orbitofrontal cortex

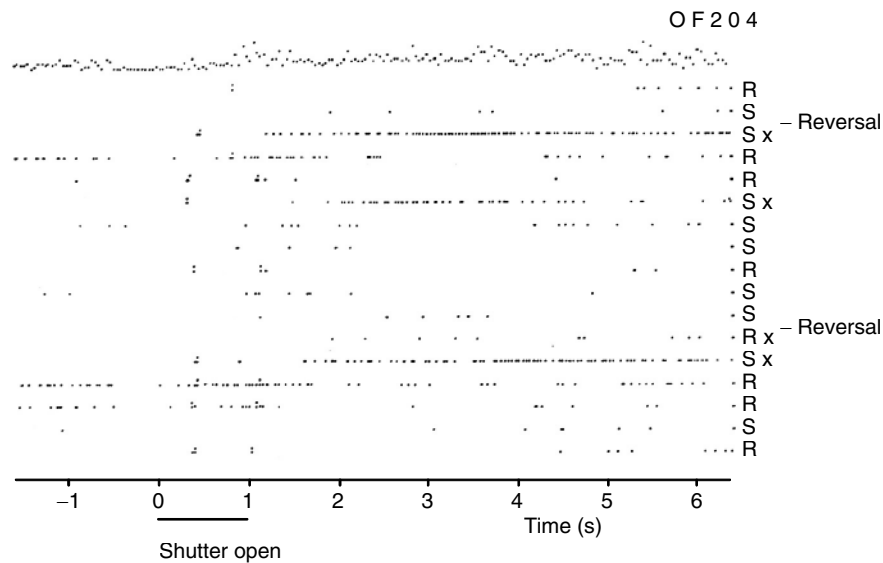


Figure 4 Error neuron: Responses of an orbitofrontal cortex neuron that responded only when the monkey licked to a visual stimulus during reversal, expecting to obtain fruit juice reward, but actually obtaining the taste of aversive saline because it was the first trial of reversal. Each single dot represents an action potential; each vertically arranged double dot represents a lick response. The visual stimulus was shown at time 0 for 1 s. The neuron did not respond on most reward (R) or saline (S) trials, but did respond on the trials marked 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. 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. A model for this has been described by Deco and Rolls (2005). After Thorpe, Rolls and Maddison 1983.

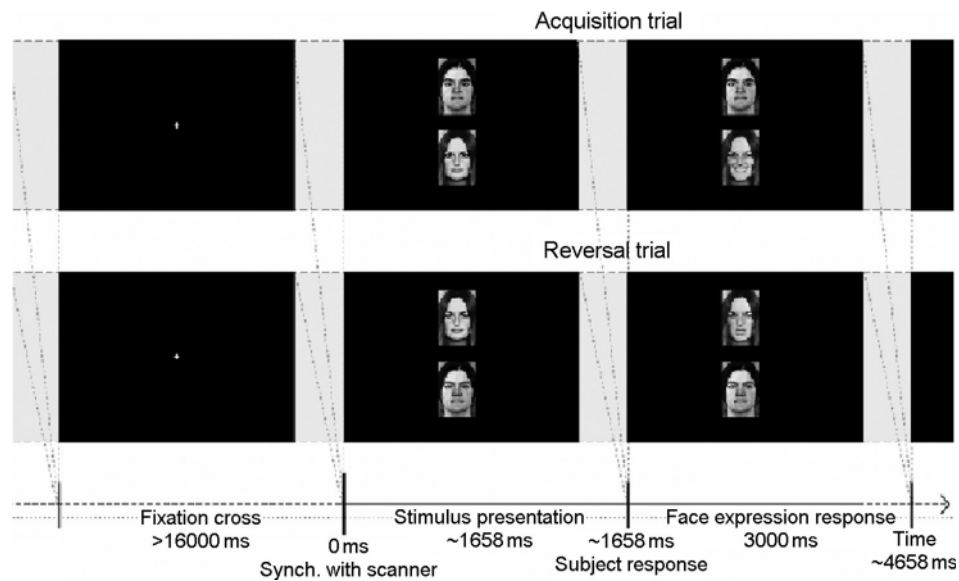


Figure 5 Social reversal task: The trial starts synchronised with the scanner and two people with neutral face expressions are presented to the subject. The subject has to select one of the people by pressing the corresponding button, and the person will then either smile or show an angry face expression for 3000 ms depending on the current mood of the person. The task for the subject is to keep track of the mood of each person and choose the 'happy' person as much as possible (upper row). Over time (after between 4 and 8 correct trials) this will change so that the 'happy' person becomes 'angry' and vice versa, and the subject has to learn to adapt her choices accordingly (bottom row). Randomly intermixed trials with either two men, or two women, were used to control for possible gender and identification effects, and a fixation cross was presented between trials for at least 16000 ms. After Kringelbach and Rolls, 2003.

does have all three types of neuron and the required neuroanatomically defined inputs, and this is an important site in the brain for computing error signals (Rolls, 2014).

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. When it 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 correct affective states. This type of error representation 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.

We have also been able to obtain evidence that non-reward used as a signal to reverse behavioural choice to visual stimuli 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 (see Figure 5). It was found that activation of a lateral part of the orbitofrontal cortex in the fMRI study was produced on the error trials when the human chose a face and did not obtain the expected reward (Figure 6). Control tasks showed that the response was related to the error, and the mismatch between what was expected and what was obtained, 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.

A Visual Representation of Faces in the Orbitofrontal Cortex

Another type of information represented in the orbitofrontal cortex is information about faces. There is a population of orbitofrontal cortex face-selective neurons that respond in many ways similarly to those in the temporal cortical visual areas (see Computational Neuroscience of Vision Rolls and Deco 2002 for a description of their properties). The orbitofrontal face-responsive neurons, first observed by Thorpe et al. (1979), then by Rolls et al. (2006), tend to respond with longer latencies than temporal




Figure 6 Social reversal: Composite figure showing that changing behaviour based on face expression is correlated with increased brain activity in the human orbitofrontal cortex. (a) The figure is based on two different group statistical contrasts from the neuroimaging data which are superimposed on a ventral view of the human brain with the cerebellum removed, and with indication of the location of the two coronal slices (b,c) and the transverse slice (d). The red activations in the orbitofrontal cortex (denoted OFC, maximal activation: $z = 4.94$: 42,42,-8; and $z = 5.51$; $X, Y, Z = -46, 30, -8$) shown on the rendered brain arise from a comparison of reversal events with stable acquisition events, while the blue activations in the fusiform gyrus (denoted Fusiform, maximal activation: $z > 8$; 36,-60,-20 and $z = 7.80$; -30,-56,-16) arise from the main effects of face expression. (b) The coronal slice through the frontal part of the brain shows the cluster in the right orbitofrontal cortex across all nine subjects when comparing reversal events with stable acquisition events. Significant activity was also seen in an extended area of the anterior cingulate/paracingulate cortex (denoted Cingulate, maximal activation: $z = 6.88$; -8,22,52; green circle). (c) The coronal slice through the posterior part of the brain shows the brain response to the main effects of face expression with significant activation in the fusiform gyrus and the cortex in the intraparietal sulcus (maximal activation: $z > 8$; 32,-60, 46; and $z > 8$; -32,-60,44). (d) The transverse slice shows the extent of the activation in the anterior cingulate/paracingulate cortex when comparing reversal events with stable acquisition events. Group statistical results are superimposed on a ventral view of the human brain with the cerebellum removed, and on coronal and transverse slices of the same template brain (activations are thresholded at $p = 0.0001$ for purposes of illustration to show their extent). After Kringsbach and Rolls 2003.

lobe neurons (140–200 ms typically, compared with 80–100 ms); they also convey information about which face is being seen, by having different responses to different faces (see Figure 7); and are typically rather harder to activate strongly than temporal cortical face-selective neurons, in that many of them respond much better to real faces than to two-dimensional images of faces on a video monitor. Some of the orbitofrontal cortex face-selective neurons are responsive to face gesture or movement. The findings are consistent with the likelihood that these neurons are activated via the inputs from the temporal cortical visual areas in which face-selective neurons are found. The significance of the neurons is likely to be related to the fact that faces convey information that is important in social reinforcement, both by conveying face expression, which can indicate reinforcement, and by encoding information about which individual is present, also important in evaluating and utilizing reinforcing inputs in social situations.

Consistent with these findings in macaques, and as described above, in humans, activation of the lateral orbitofrontal cortex occurs when a rewarding smile expression is expected, but an angry face expression is obtained, in a visual discrimination reversal task. This is an example of the operation of a social reinforcer, and, consistent with these results, activation of the orbitofrontal cortex is found when humans are making social judgements. In addition, activation of the medial orbitofrontal cortex is correlated with face attractiveness.

In relation to possible functions in memory of the orbitofrontal cortex, it is of interest that a population of orbitofrontal cortex neurons responds to novel visual stimuli (Rolls et al., 2005).

Effects of Damage to the Orbitofrontal Cortex on Visual Processing

Consistent with the neurophysiological and neuroimaging results described in this contribution, lesions in humans and macaques of the orbitofrontal cortex impair visual discrimination reversal. In humans, the ability to identify face expressions is also impaired, though face recognition is not. These and related findings about changes in emotion, and their implications, are described by Rolls (2014).

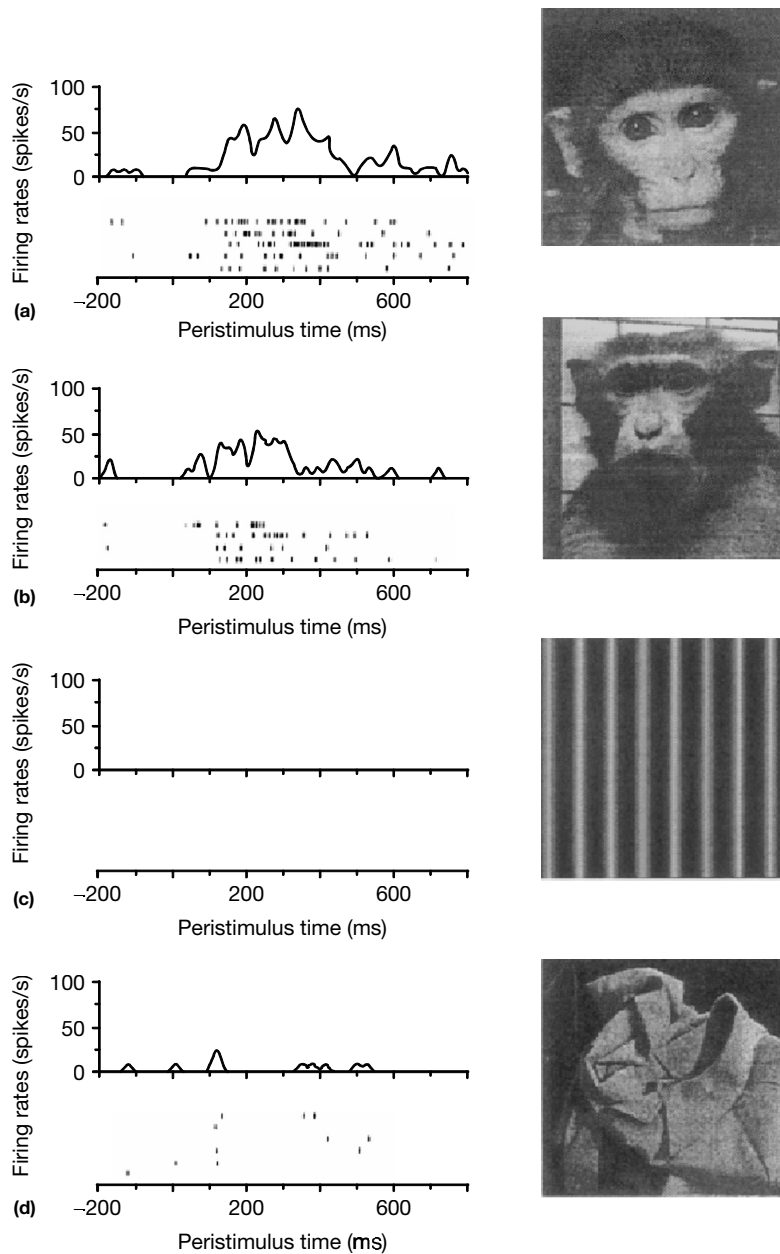


Figure 7 Orbitofrontal cortex face-selective neuron as found in macaques. Peristimulus rastergrams and time histograms are shown. Each trial is a row in the rastergram. Several trials for each stimulus are shown. The ordinate is in spikes/s. The neuron responded best to face (a), also responded, though less to face (b), had different responses to other faces (not shown), and did not respond to non-face stimuli (e.g. (c) and (d)). The stimulus appeared at time 0 on a video monitor. From Rolls, Critchley, Browning, and Inoue, 2006.

References

- Deco G and Rolls ET (2005) Synaptic and spiking dynamics underlying reward reversal in orbitofrontal cortex. *Cerebral Cortex* 15: 15–30.
- Grabenhorst F and Rolls ET (2011) Value, pleasure, and choice in the ventral prefrontal cortex. *Trends in Cognitive Sciences* 15: 56–67.
- Kringelbach ML and Rolls ET (2003) Neural correlates of rapid reversal learning in a simple model of human social interaction. *NeuroImage* 20: 1371–1383.
- O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, and Andrews C (2001) Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience* 4: 95–102.
- Padoa-Schioppa C (2011) Neurobiology of economic choice: A good-based model. *Annual Review of Neuroscience* 34: 333–359.
- Rolls ET, Critchley H, Mason R, and Wakeman EA (1996) Orbitofrontal cortex neurons: role in olfactory and visual association learning. *Journal of Neurophysiology* 75: 1970–1981.
- Rolls ET and Deco G (2002) *Computational Neuroscience of Vision*. Oxford: Oxford University Press.
- Rolls ET (2012) Invariant visual object and face recognition: Neural and computational bases, and a model, VisNet. *Frontiers in Computational Neuroscience* 6: 1–70 35.
- Rolls ET (2014) *Emotion and decision-making explained*. Oxford: Oxford University Press.
- Rolls ET and Grabenhorst F (2008) The orbitofrontal cortex and beyond: From affect to decision-making. *Progress in Neurobiology* 86: 216–244.

- Rolls ET, Browning AS, Inoue K, and Hernadi S (2005) Novel visual stimuli activate a population of neurons in the primate orbitofrontal cortex. *Neurobiology of Learning and Memory* 84: 111–123.
- Rolls ET, Critchley HD, Browning AS, and Inoue K (2006) Face-selective and auditory neurons in the primate orbitofrontal cortex. *Experimental Brain Research* 170: 74–87.
- Schultz W (2013) Updating dopamine reward signals. *Current Opinion in Neurobiology* 23: 229–238.
- Thorpe SJ, Maddison S, and Rolls ET (1979) Single unit activity in the orbitofrontal cortex of the behaving animal. *Neuroscience Letters* S3:.
- Thorpe SJ, Rolls ET, and Maddison S (1983) Neuronal activity in the orbitofrontal cortex of the behaving monkey. *Experimental Brain Research* 49: 93–115.