
Top – down control of visual perception: Attention in natural vision

Edmund T Rolls

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, UK; e-mail: Edmund.Rolls@psy.ox.ac.uk

Presented at Interdisciplinary Conference on Pre-Emptive Perception, Hanse Institute for Advanced Studies, Delmenhorst, Germany, 15–18 October 2005

Abstract. Top–down perceptual influences can bias (or pre-empt) perception. In natural scenes, the receptive fields of neurons in the inferior temporal visual cortex (IT) shrink to become close to the size of objects. This facilitates the read-out of information from the ventral visual system, because the information is primarily about the object at the fovea. Top–down attentional influences are much less evident in natural scenes than when objects are shown against blank backgrounds, though are still present. It is suggested that the reduced receptive-field size in natural scenes, and the effects of top–down attention contribute to change blindness. The receptive fields of IT neurons in complex scenes, though including the fovea, are frequently asymmetric around the fovea, and it is proposed that this is the solution the IT uses to represent multiple objects and their relative spatial positions in a scene. Networks that implement probabilistic decision-making are described, and it is suggested that, when in perceptual systems they take decisions (or ‘test hypotheses’), they influence lower-level networks to bias visual perception. Finally, it is shown that similar processes extend to systems involved in the processing of emotion-provoking sensory stimuli, in that word-level cognitive states provide top–down biasing that reaches as far down as the orbitofrontal cortex, where, at the first stage of affective representations, olfactory, taste, flavour, and touch processing is biased (or pre-empted) in humans.

1 Introduction

Pre-emptive perception can be used to describe a situation in which a motor command influences perception. An example might be a command for a saccadic eye movement, in which a corollary discharge cancels the changed sensory input produced by the movement of the eyes, contributing to stability of visual perception (Wurtz and Sommer 2004). Here, I consider a somewhat different situation in which prior expectations influence visual perception. These prior expectations might consist of top–down attentional influences on object or spatial processing, or they might consist of prior knowledge of the world which can influence what is seen, as in, for example, the perception of a Necker cube as three-dimensional. A particular focus of this paper is on how attention operates in natural, cluttered, visual scenes, where its influence is much more restricted than when two objects are presented on a blank background. The neurophysiological approach to how attention operates in natural vision leads to hypotheses about the binding problem in perception; about how object representations are interfaced to action systems; and about change blindness. It is shown that even cognitive influences produced by a word descriptor can by top–down processing ‘pre-empt’ affective perception in the olfactory, taste, flavour, somatosensory, and emotional systems, by a top–down modulation of representations in the orbitofrontal cortex. These top–down effects do not pre-empt perception in the sense that they override it. Indeed, it is important that bottom–up inputs dominate the processing to avoid hallucinatory activity. Instead, top–down processes can bias visual perception, in ways for which we now have clear computational-neuroscience models at the level of integrate-and-fire neurons, which show that the top–down effects are most evident when the bottom–up input is weak or ambiguous. The processes described here provide a way for understanding not only top–down effects in attention, but also top–down effects that are

captured by the ideas of inference and hypothesis-testing in perception (Gregory 1970, 1998; Helmholtz 1867/1962; Rolls 2008).

Some of the evidence described comes from neuronal recordings in the macaque inferior temporal visual cortex (IT) which is much more developed than in non-primates and provides a model useful for understanding this system and systems connected to it, and their disorders, in humans (Rolls 2005b; Rolls and Deco 2002). Evidence on the representations of objects that are invariant with respect to the position on the retina, the size, and even the view, that are provided by populations of neurons in the IT is described elsewhere (Rolls 2000, 2007, 2008; Rolls and Deco 2002; Rolls and Stringer 2006b).

2 The representation of objects in complex natural scenes

2.1 Object-based attention and object selection in complex natural scenes

Object-based attention refers to attention to an object. For example, in a visual-search task the object might be specified as what should be searched for, and its location must be found. In spatial attention, a particular location in a scene is pre-cued, and the object at that location may need to be identified.

Much of the neurophysiology, psychophysics, and modeling of attention has been with a small number—typically two—of objects in an otherwise blank scene. In this section, I consider how attention operates in complex natural scenes, and in particular describe how the IT operates to enable the selection of an object in a complex natural scene. In relation to pre-emptive perception, the model of the processing is that a short-term-memory system in, for example, the prefrontal cortex provides a top-down signal that acts via biased competition to influence the representations in visual cortical processing areas such as the IT, in ways described in sections 3 and 4.

To investigate how attention operates in complex natural scenes, and how information is passed from the IT to other brain regions to enable stimuli to be selected from natural scenes for action, Rolls et al (2003) analysed the responses of IT neurons to stimuli presented in complex natural backgrounds while performing a top-down object-based attentional-search task. The monkey had to search for two objects on a screen: a touch of one object was rewarded with juice, and a touch of another object was punished with saline (see figure 1). Neuronal responses to the effective stimuli for the neurons were compared when the objects were presented in the natural scene or on a plain background. The response of the neuron to objects at the fovea was hardly reduced when they were presented in natural scenes, and the selectivity of the neurons remained (see figures 2 and 3). However, the main finding was that the magnitudes of the responses of the neurons typically became much less in the real scene the further the monkey fixated in the scene away from the object (see figure 2). It is proposed that this reduced translation invariance (ie invariance with respect to position on the retina) in natural scenes helps an unambiguous representation of an object which may be the target for action to be passed to the brain regions which receive from the primate IT. It helps with the binding problem, by reducing in natural scenes the effective receptive field of at least some IT neurons to approximately the size of an object in the scene. In a very similar task, in which one of two objects had to be selected against a complex background, almost all the information (> 95% of the total information) was found to be encoded by the firing rates of simultaneously recorded IT neurons, and very little information was encoded by stimulus-dependent synchronisation which may therefore not be important for implementing feature binding (Aggelopoulos et al 2005; Rolls 2008).

In natural scenes, the effect of object-based attention on the response properties of IT neurons was found to be relatively small, as illustrated in figure 3 (Rolls et al 2003). The results summarised in figure 3 for 5 deg stimuli show that the receptive fields were large (77.6 deg) with a single stimulus in a blank background, and were

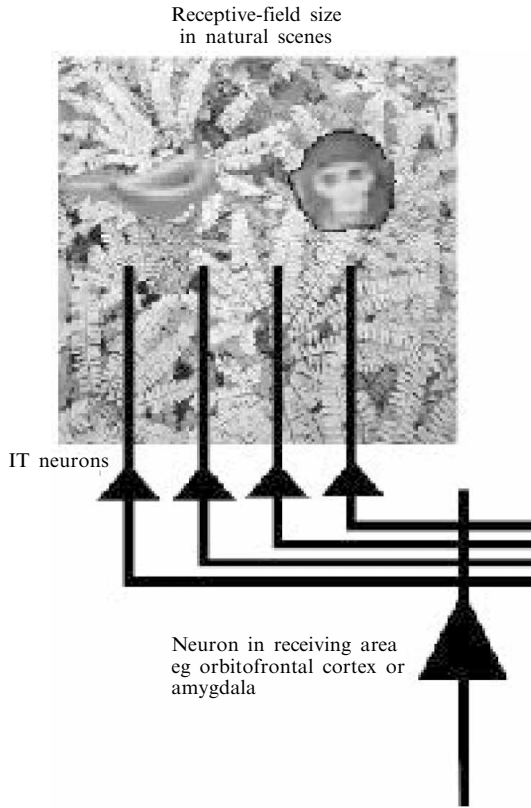


Figure 1. Objects shown in a natural scene, in which the task was to search for and touch one of the stimuli. The objects in the task as run were smaller. The diagram shows that if the receptive fields of IT neurons are large in natural scenes with multiple objects, then any receiving neuron in structures such as the orbitofrontal and amygdala would receive information from many stimuli in the field of view, and would not be able to provide evidence about each of the stimuli separately.




Figure 2. Firing of a temporal cortex cell to an effective stimulus, presented either in a blank background or in a natural scene, as a function of the angle at which the monkey was fixating away from the effective stimulus. The task was to search for and touch the stimulus. (After Rolls et al 2003.)

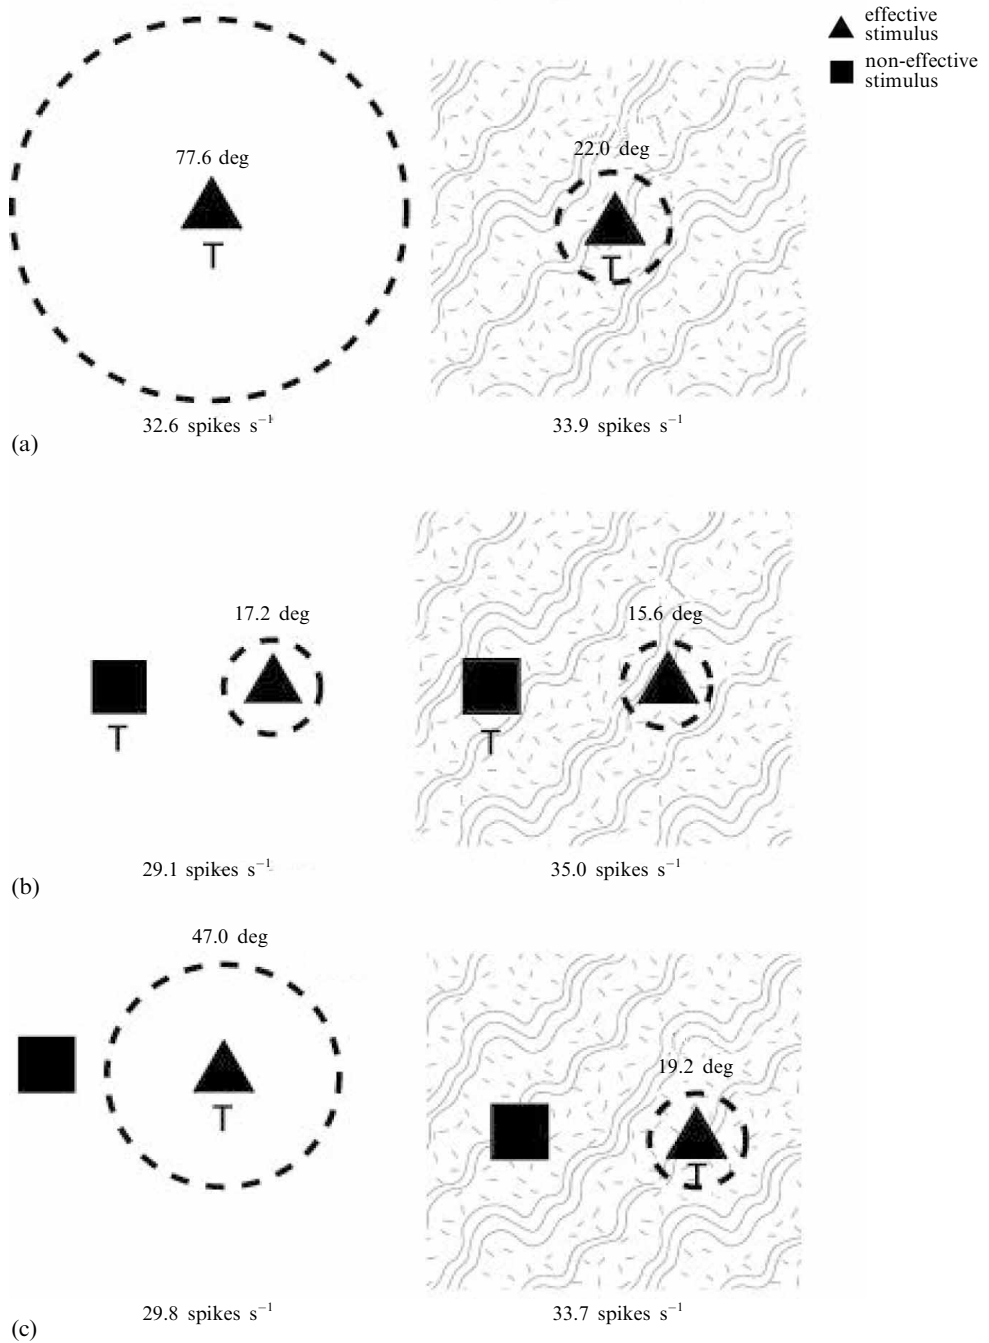


Figure 3. Summary of the receptive-field sizes of IT neurons to a 5 deg effective stimulus presented in either a blank background (blank screen) or in a natural scene (complex background). The stimulus that was a target for action in the different experimental conditions is marked by T. When the target stimulus was touched, a reward was obtained. The mean receptive-field diameter of the population of neurons analysed and the mean firing rate are shown. The stimuli subtended 5 deg × 3.5 deg at the retina, and occurred on each trial in a random position on the 70 deg × 55 deg screen. The dashed circle is proportional to the receptive-field size. (a) Responses with one visual stimulus in a blank (left) or complex (right) background. (b) Responses with two stimuli, when the effective stimulus was not the target of the visual search. (c) Responses with two stimuli, when the effective stimulus was the target of the visual search. (After Rolls et al 2003.)

greatly reduced in size (to 22.0 deg) when presented in a complex natural scene. The results also show that there was little difference in receptive-field size or firing rate in the complex background when the effective stimulus was selected for action (19.2 deg), and when it was not (15.6 deg) (Rolls et al 2003). (For comparison, the effects of attention against a blank background were much larger, with the receptive field increasing from 17.2 deg to 47.0 deg as a result of object-based attention, as shown in figure 3.) The computational basis for these relatively minor effects of object-based attention when objects are viewed in natural scenes is considered in section 3.

2.2 The interface from object representations to action

These findings on how objects are represented in natural scenes make the interface to memory and to action systems simpler, in that what is at the fovea can be interpreted (eg by an associative memory in the orbitofrontal cortex or amygdala) partly independently of the surroundings, and choices and actions can be directed if appropriate to what is at the fovea (Ballard 1993; Rolls and Deco 2002). There thus may be no need to have the precise coordinates of objects in space represented in the IT and passed to the motor system for action to be directed at the target. Instead, given that the output of the IT in complex visual scenes is primarily about an object at the fovea, the dorsal visual system may be able to initiate action to whatever is at the fovea (Rolls 2008; Rolls and Deco 2006). The condition for an action to be performed is that the ventral visual system must have provided a representation of the object, and this must have been identified as a goal (ie as a rewarded object) by, for example, associative reward-based lookup in brain structures such as the orbitofrontal cortex and amygdala, where visual stimuli are interfaced to reward systems (Rolls 2005b).

2.3 Change blindness and inattentional blindness

Change blindness refers to our inability to detect a change to objects in a scene that can occur if the scene is changed during for example a blink or saccade (Simons and Rensink 2005). A dramatic example of the related phenomenon, inattentional blindness, is that when watching a basketball-passing event and counting, and therefore paying attention to ball passes, one may be unaware that a gorilla is walking across the scene (Simons and Chabris 1999). It is proposed here that an important factor that contributes to change blindness is the reduced diameter of the receptive fields of IT neurons that occur in natural scenes, as described in section 2.1. This would result in a failure to activate representations of objects (such as the gorilla that appears in the middle of the basketball-passing event), if they are not close to the fovea, in complex cluttered scenes. The instructions to count the number of ball passes in the basketball-passing event would tend to result in fixation on the ball, and thus for objects that are not close to the ball not to be seen, owing to the reduced receptive-field size of IT neurons in complex scenes. The computational mechanisms that produce this reduction in receptive-field size in complex natural scenes, and that thus it is proposed make an important contribution to change blindness, are described in section 3. One is a bottom-up mechanism that contributes to change blindness, involving interactions between sensory inputs being received by the visual system. Another factor that is likely to contribute to change blindness [or, more particularly, inattentional blindness (Simons and Chabris 1999)] is the top-down attentional biasing of the object representations that have been cued as relevant in the scene, such as the ball, as described in sections 2.1 and 2.4 (though this top-down attentional biasing is less effective in complex natural scenes than when two objects are shown against a blank background, as also described in section 2.1). The effect of the top-down attention in natural scenes is to provide a small increase in the receptive-field size, relative to the relatively small size in a complex scene when attention is not being paid to the object (Rolls et al 2003). This top-down biasing effect would thus make a contribution to

inattentive blindness, by increasing a little the reduced receptive-field sizes of IT neurons, making it more likely that an object would be seen with characteristics similar to those to which attention was being paid. The mechanisms can be thought to implement a way in which visual perception is pre-empted, in this case by the bottom-up process of the reduced receptive fields of IT neurons in complex scenes, and by top-down attentional influences. Both the bottom-up and top-down mechanisms are likely to contribute to a small level of activity in neurons representing the object not being fixated and/or not being attended to, so that the level of activity falls below that required to reach conscious perception, which has been analysed in backward-masking experiments (Rolls 2003, 2005a; Rolls and Tovee 1994; Rolls et al 1994).

2.4 The representation of information about the relative positions of multiple objects in a scene
These experiments have been extended to address the issue of how several objects are represented in a complex scene. The issue arises because the relative spatial locations of objects in a scene must be encoded [and is possible even in short presentation times without eye movements (Biederman 1972)] (and this has been held to involve some spotlight of attention); and because, as shown above, what is represented in complex natural scenes is primarily about what is at the fovea, yet we can locate more than one object in a scene even without eye movements. Aggelopoulos and Rolls (2005) showed that with five objects simultaneously present in the receptive field of IT neurons, although all the neurons responded to their effective stimulus when it was at the fovea, some could also respond to their effective stimulus when it was in a parafoveal position 10° from the fovea. An example of such a neuron is shown in figure 4. The asymmetry is much more evident in a scene with five images present (figure 4a) than when only one image is shown on an otherwise blank screen (figure 4b). Competition between different stimuli in the receptive field thus reveals the asymmetry in the receptive field of IT neurons.

The asymmetry provides a way of encoding the position of multiple objects in a scene. Depending on which asymmetric neurons are firing, the population of neurons provides information [using a distributed representation of the type that a population of receiving neurons with mutual, lateral, inhibition can decode—see Rolls (2008)] to the next processing stage, not only about which image is present at or close to the fovea, but where it is with respect to the fovea. This information is provided by neurons that have firing rates that reflect the relevant information, and stimulus-dependent synchrony is not necessary. Top-down attentional biasing input could thus, by biasing the appropriate neurons, facilitate bottom-up information about objects without any need to alter the time relations between the firing of different neurons. The exact position of the object with respect to the fovea, and effectively thus its spatial position relative to other objects in the scene, would then be made evident by the subset of asymmetric neurons firing.

This is thus the solution that these experiments indicate is used for the representation of multiple objects in a scene (Aggelopoulos and Rolls 2005)—an issue that has previously been difficult to account for in neural systems with distributed representations (Mozier 1991) and for which ‘attention’ has been a proposed solution.

3 Object representation and attention in natural scenes: a computational account

The results described in section 2 and summarised in figure 3 show that the receptive fields of IT neurons were large with a single stimulus in a blank background, and were greatly reduced in size when measured in a complex natural scene. The results also show that there was little difference in receptive-field size or firing rate in the complex background when the effective stimulus was selected for action, and when it was not (Rolls et al 2003).

Trappenberg et al (2002) have suggested what underlying mechanisms could account for these findings, and simulated a model to test the ideas. The model utilises an attractor

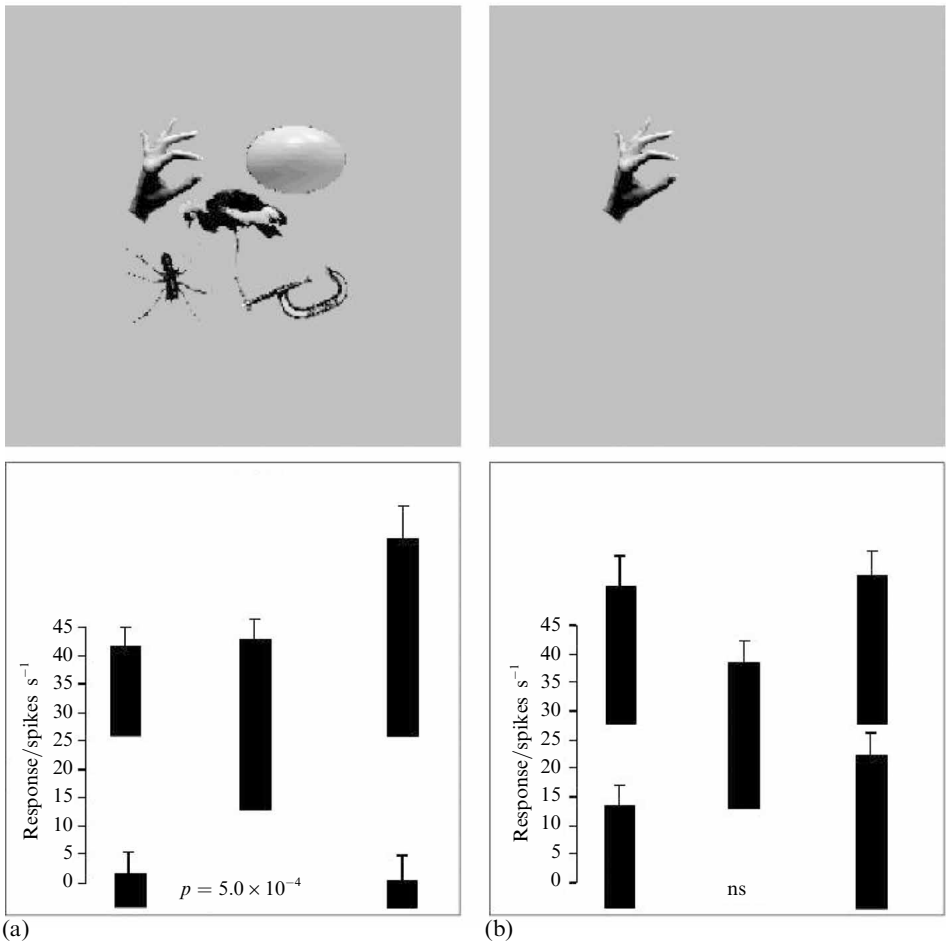


Figure 4. (a) The responses [firing rate with the spontaneous rate subtracted; means (M) \pm 1 SEM] of one neuron when tested with five stimuli simultaneously present in the close (10 deg) configuration with the parafoveal stimuli located 10° from the fovea. (b) The responses of the same neuron when only the effective stimulus was presented in each position. The firing rate for each position is that when the effective stimulus for the neuron was in that position. The p value is that from the ANOVA calculated over the four parafoveal positions. (After Aggelopoulos and Rolls 2005)

network representing the IT cortex (implemented by the recurrent excitatory connections between IT neurons), and a neural input layer with several retinotopically organised modules representing the visual scene in an earlier visual cortical area such as V4 (see figure 5).⁽¹⁾ The attractor-network aspect of the model produces the property that receptive fields of IT neurons can be large in blank scenes by enabling a weak input in the periphery of the visual field to act as a retrieval cue for the object attractor.

⁽¹⁾ An attractor or autoassociation network is implemented by associatively modifiable connections between the neurons in the network. Each vector of neuronal firing rates represents a stimulus or memory, and is stored by associative synaptic modification between the neurons representing that stimulus or memory. When even a partial retrieval cue is provided that is similar to one of the patterns stored in the network, the network is attracted into the state with the neurons in the original pattern active, and thus implements memory retrieval. The neurons continue firing stably with just that set of neurons active, thus implementing short-term memory too. A description of the operation of these and other networks is provided elsewhere (Rolls 2008).

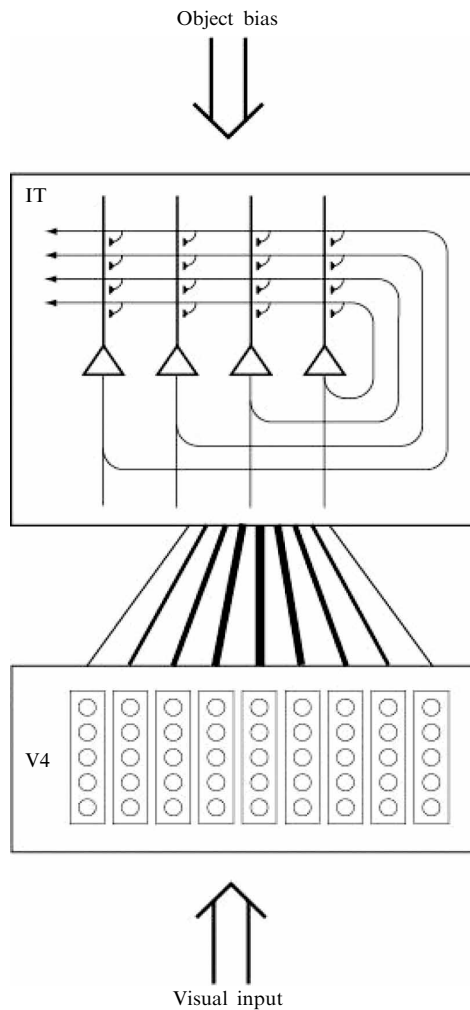


Figure 5. The architecture of the IT model of Trappenberg et al (2002) operating as an attractor network with inputs from the fovea given preferential weighting by the greater magnification factor of the fovea. The model also has a top-down object-selective bias input. The model was used to analyse how object vision and recognition operate in complex natural scenes.

On the other hand, when the object is shown in a complex background, the object closest to the fovea tends to act as the retrieval cue for the attractor, because the fovea is given increased weight in activating the IT module because the magnitude of the input activity from objects at the fovea is greatest owing to the higher cortical magnification factor of the fovea incorporated into the model. [The cortical magnification factor can be expressed as the number of millimetres of cortex representing 1 deg of visual field. The cortical magnification factor decreases rapidly with increasing eccentricity from the fovea (Cowey and Rolls 1975; Rolls and Cowey 1970).] This results in smaller receptive fields of IT neurons in complex scenes, because the objects tend to need to be close to the fovea to trigger the attractor into the state representing that object. (In other words, if the object is far from the fovea in a cluttered scene, then the object will not trigger neurons in IT which represent it, because neurons in IT are preferentially being activated by another object at the fovea.) This may be described as an attractor model in which the competition for which attractor state is retrieved is weighted towards objects at the fovea.

Attentional top-down object-based inputs can bias the competition implemented in this attractor model, but have relatively minor effects (in, for example, increasing receptive-field size) when they are applied in a complex natural scene, as then, as usual, the stronger forward inputs dominate the states reached. In this network, the recurrent collateral connections may be thought of as implementing constraints between the different inputs present, to help arrive at firing in the network which best meets the constraints. In this scenario, the preferential weighting of objects close to the fovea, because of the increased magnification factor at the fovea, is a useful principle in enabling the system to provide useful output. The top-down attentional biasing effect on an object is much more marked in a blank scene, or a scene with only two objects present at similar distances from the fovea, which are conditions in which attentional effects have frequently been examined. The results of the investigation (Trappenberg et al 2002) thus suggest that attention may be a much more limited phenomenon in complex, natural, scenes than in reduced displays with one or two objects present. The results also suggest that the alternative principle, of providing strong weight to whatever is close to the fovea, is an important principle governing the operation of the IT, and in general of the output of the ventral visual system in natural environments. This principle of operation is very important in interfacing the visual system with action systems, because the effective stimulus in making IT neurons fire is in natural scenes usually on, or close to, the fovea. This means that the spatial coordinates of where the object is in the scene do not have to be represented in the IT, nor passed from it to the action selection system, as the latter can assume that the object making IT neurons fire is close to the fovea in natural scenes (Rolls et al 2003; Rolls and Deco 2002).

There may, of course, be, in addition, a mechanism for object selection that takes into account the locus of covert attention when actions are made to locations not being looked at. However, the simulations described in this section suggest that, in any case, covert attention is likely to have a much less significant influence on visual processing in natural scenes than in reduced scenes with one or two objects present.

Given these points, one might question why IT neurons can have such large receptive fields, which show translation invariance (Rolls 2000; Rolls et al 2003). At least part of the answer to this may be that IT neurons must have the capability to have large receptive fields if they are to deal with large objects (Rolls and Deco 2002). A V1 neuron, with its small receptive field, simply could not receive input from all the features necessary to define an object. On the other hand, IT neurons may be able to adjust their size to approximately the size of objects, using in part the interactive attentional effects of bottom-up and top-down effects described elsewhere in this paper.

The implementation of the simulations is described by Trappenberg et al (2002), and some of the results obtained with the architecture, shown in figure 5, follow. In one simulation, only one object was present in the visual scene in a plain background at different eccentricities from the fovea. As shown in figure 6a by the line labelled 'simple background', the receptive fields of the neurons were very large. The value of the object bias, k_{ITbias} , was set to zero in these simulations. Good object retrieval (indicated by large correlations) was found even when the object was far from the fovea, indicating large IT receptive fields with a blank background. The reason that any drop is seen in performance as a function of eccentricity is that some noise was present in the recall process. This demonstrates that the attractor dynamics can support translation-invariant object recognition even though the translation-invariant weight vectors between V4 and IT are explicitly mapped by a modulation factor derived from the cortical magnification factor.

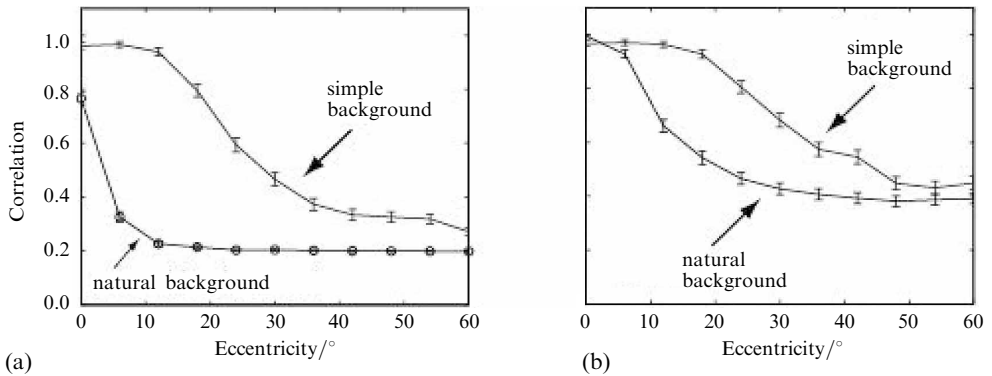


Figure 6. Correlations between the state produced in the IT and the representation of the object in the IT when the object was shown against a blank background, or was shown with other trained objects at all possible locations in the visual scene (natural background). There is no top-down attentional bias in (a), whereas an object-bias facilitating retrieval of the object was present in (b). [In (b), $k_{ITbias} = 0.7$ in the experiments with a natural background, and $k_{ITbias} = 0.1$ in the experiments with a blank background, with details provided in Trappenberg et al (2002).]

In a second simulation, individual objects were placed at all possible locations in a natural and cluttered visual scene. The resulting correlations between the target pattern and the asymptotic IT state are shown in figure 6a with the line labelled 'natural background'. Many objects in the visual scene are now competing for recognition by the attractor network, and the objects around the foveal position are enhanced through the modulation factor derived from the cortical magnification factor. This results in a much smaller size of the receptive field of IT neurons when measured with objects in natural backgrounds.

In addition to this major effect of the background on the size of the receptive field, which parallels, and we suggest may account for, the physiological findings outlined above, there is also a dependence of the size of the receptive fields on the level of object bias provided to the IT network. Examples are shown in figure 6b where an object bias was used. The object bias biases the IT network towards the expected object with a strength determined by the value of k_{ITbias} , and has the effect of increasing the size of the receptive fields in both blank and natural backgrounds (see figure 6b compared to figure 6a). This models the effect found neurophysiologically (Rolls et al 2003).

In natural scenes, the model is able to account for the neurophysiological data that the IT neuronal responses are larger when the object is close to the fovea, by virtue of the fact that objects close to the fovea are weighted by the cortical magnification factor. The model accounts for the larger receptive-field sizes from the fovea of IT neurons in natural backgrounds if the target is the object being selected compared to when it is not selected (Rolls et al 2003). The model accounts for this by an effect of top-down bias which simply biases the neurons towards particular objects, compensating for their decreasing inputs produced by the decreasing magnification factor modulation with increasing distance from the fovea. Such object-based attention signals could originate in the prefrontal cortex and could provide the object bias for the IT (Renart et al 2000, 2001; Rolls and Deco 2002). Important properties of the architecture for obtaining the results just described are the high magnification factor at the fovea and the competition between the effects of different inputs, implemented in the above simulation by the competition inherent in an attractor network.

Similar results have also been obtained in a hierarchical feedforward network where each layer operates as a competitive network (Deco and Rolls 2004). This network thus captures many of the properties of our hierarchical model of invariant visual object recognition in the ventral visual stream (Elliffe et al 2002; Rolls 1992; Rolls and Deco 2002; Rolls and Milward 2000; Rolls and Stringer 2001, 2006a, 2006b; Stringer et al 2006; Stringer and Rolls 2000, 2002; Wallis and Rolls 1997), but incorporates in addition a foveal magnification factor and top-down projections with a dorsal visual stream so that attentional effects can be studied, as shown in figure 7.




Figure 7. Cortical architecture for hierarchical and attention-based visual perception. The system is essentially composed of five modules structured such that they resemble the two known main visual paths of the mammalian visual cortex. Information from the retino-geniculo-striate pathway enters the visual cortex through area V1 in the occipital lobe and proceeds into two processing streams. The occipital-temporal stream leads ventrally through V2–V4 and IT, and is mainly concerned with object recognition. The occipito-parietal stream leads dorsally into MT (posterior parietal complex), and is responsible for maintaining a spatial map of an object's location. The solid lines with arrows between levels show the forward connections, and the dashed lines the top-down backprojections. Short-term-memory systems in the prefrontal cortex (PF 46) apply top-down attentional bias to the object or spatial processing streams. (After Deco and Rolls 2004.)

Deco and Rolls (2004) trained the network described, shown in figure 7 with two objects, and used the trace-learning rule (Rolls and Milward 2000; Wallis and Rolls 1997) in order to achieve translation invariance. In a first experiment only one object was placed on the retina at different distances from the fovea (ie different eccentricities relative to the fovea). This corresponds to the blank background condition. In a second experiment, the object was also displaced at different eccentricities relative to the fovea, but on a cluttered natural background.

Figure 8 shows the average firing activity of the IT neuron specific for the test object as a function of the position of the object on the retina relative to the fovea (eccentricity). In both cases relatively large receptive fields are observed, because of the translation invariance obtained with the trace-learning rule and the competition mechanisms implemented within each layer of the ventral stream. (The receptive-field size is defined as the width of the receptive field at the point where there is a half-maximal response.) However, when the object was in a blank background (solid line), larger receptive fields were observed. The decrease in neuronal response as a function of distance from the fovea is mainly due to the effect of the magnification factor implemented in V1. On the other hand, when the object was in a complex cluttered background (dashed line), the effective size of the receptive field of the same IT neuron shrinks because of competitive effects between the object features and the background features in each layer of the ventral stream. In particular, the global character of the competition expressed in the IT module (due to the large receptive fields and the local character of the inhibition, in our simulations, between the two object-specific pools) is the main cause of the reduction of the receptive fields in the complex scene. The exact reduction in the receptive-field size depends on the amount of competing background information in the scene, and indeed Deco and Rolls (2004) were able to predict the changes in the shapes of the receptive fields depending on where distractors were located relative to the test object and the fovea.




Figure 8. Average firing activity of an IT neuron as a function of eccentricity from the fovea, in the simulation of Deco and Rolls (2004). When the object was in a blank background (solid line), large receptive fields are observed because of the translation invariance of inferior temporal neurons. The decay is mainly due to the magnification factor implemented in V1. When the object was presented in a complex cluttered natural background (dashed line), the effective size of the receptive field of the same IT neuron was reduced because of competitive effect between the object features and the background features within each layer of the ventral stream. (After Deco and Rolls 2004.)

Deco and Rolls (2004) also studied the influence of object-based attentional top-down bias on the effective size of an IT neuron for the case of an object in a blank or a cluttered background. To do this, they repeated the two simulations but now considered a non-zero top-down bias coming from prefrontal area 46v and impinging on the IT neuron specific for the object tested. Figure 9 shows the results. The average firing activity normalised to the maximum value is plotted to compare the neuronal activity as a function of the eccentricity. When no attentional object bias is introduced (a), a shrinkage of the receptive-field size is observed in the complex background (dashed line). When attentional object bias is introduced (b), the shrinkage of the receptive field due to the complex background is slightly reduced (dashed line). Rolls et al (2003) also found that in natural scenes the effect of object-based attention on the response properties of IT neurons was relatively small. They found only a small difference in the receptive-field size or firing rate in the complex background when the effective stimulus was selected for action compared with that when it was not. In the framework of the model (Deco and Rolls 2004), the reduction of the shrinkage of the receptive field is due to the biasing of the competition in the IT layer in favour of the specific IT neuron tested, so that it shows more translation invariance (ie a slightly larger receptive field). The increase of the receptive field of an IT neuron, although small, produced by the external top-down attentional bias offers a mechanism for facilitation of the search for specific objects in complex natural scenes.




Figure 9. Influence of object-based attentional top-down bias from prefrontal area 46v on the effective size of an IT neuron for the case of an object in a blank (solid line) or a cluttered (dashed line) background. The average firing activity was normalised to the maximum value in order to compare the neuronal activity as a function of the eccentricity. When no attentional object bias was introduced (a), a reduction of the receptive field was observed. When attentional object bias was introduced (b), the reduction of the receptive-field size due to the complex background was slightly reduced. (After Deco and Rolls 2004.)

4 A biased competition model of object and spatial attentional effects on the representations in the visual system

In this section I consider a computational account of how top-down influences of attention can have biasing or (mildly) pre-emptive effects in visual perception.

Visual attention exerts top-down influences on the processing of sensory information in the visual cortex, and therefore is intrinsically associated with interactions between cortical areas. Thus, elucidating the neural basis of visual attention is an

excellent paradigm for understanding the basic mechanisms of intercortical neurodynamics. Recent cognitive neuroscience developments allow a more direct study of the neural mechanisms underlying attention in humans and primates. In particular, the work of Chelazzi et al (1993) has led to a promising account of attention termed the biased competition hypothesis (Desimone and Duncan 1995; Reynolds and Desimone 1999). According to this hypothesis, attentional selection operates in parallel by biasing an underlying competitive interaction between multiple stimuli in the visual field toward one stimulus or another, so that behaviourally relevant stimuli are processed in the cortex while irrelevant stimuli are filtered out. Thus, attending to a stimulus at a particular location or with a particular feature biases the underlying neural competition in a certain brain area in favour of neurons that respond to the location, or the features, of the attended stimulus. As a result of the competition, neurons that represent features without a top-down bias have reduced activity.

Neurodynamical models for biased competition have been proposed and successfully applied in the context of attention and working memory. In the context of attention, Usher and Niebur (1996) introduced an early model of biased competition. Deco and Zihl (2001) extended Usher and Niebur's model to simulate the psychophysics of visual attention by visual-search experiments in humans. Their neurodynamical formulation is a large-scale hierarchical model of the visual cortex whose global dynamics is based on biased competition mechanisms at the neural level. Attention then appears as an emergent effect related to the dynamical evolution of the whole network. This large-scale formulation, using a simplified version of the architecture shown in figure 7, has been able to simulate and explain in a unifying framework visual attention in a variety of tasks and at different cognitive neuroscience experimental measurement levels (Deco and Rolls 2005a), namely: single-cells (Deco and Lee 2002; Rolls and Deco 2002), fMRI (Corchs and Deco 2002), psychophysics (Deco and Rolls 2005a; Rolls and Deco 2002), and neuropsychology (Deco and Rolls 2002). In the context of working memory, further developments (Deco and Rolls 2003) managed to model in a unifying form attentional and memory effects in the prefrontal cortex, integrating single-cell and fMRI data, and different paradigms in the framework of biased competition.

In particular Deco and Rolls (2005b) extended previous concepts of the role of biased competition in attention by providing the first analysis at the integrate-and-fire neuronal level, which allows the neuronal nonlinearities in the system to be explicitly modeled, in order to investigate realistically the processes that underlie the apparent gain modulation effect of top-down attentional control.⁽²⁾ In the integrate-and-fire model, the competition is implemented realistically by the effects of the excitatory neurons on the inhibitory neurons, and their return inhibitory synaptic connections. This was also the first integrate-and-fire analysis of top-down attentional influences in vision that explicitly models the interaction of several different brain areas. Part of the originality of the model is that in the form in which it can account for attentional effects in V2 and V4 in the paradigms of Reynolds et al (1999) in the context of biased competition, the model with the same parameters effectively makes predictions which show that the 'contrast gain' effects in MT (Martinez-Trujillo and Treue 2002) can be accounted for by the same model. For example, the top-down attentional modulation effects are most evident when the bottom-up input is weak, eg has low contrast, because the top-down effects themselves must never be so strong that they dominate perception. In addition, the top-down modulation can appear as a nonlinear multiplication effect with the

⁽²⁾In the integrate-and-fire implementation, the synaptic currents that lead to activation of the neuron, and then the generation of a spike when a threshold is reached, are modeled, producing a system that can model many of the nonlinearities in the system, and also the effects of the probabilistic spiking of the neurons in the network on how it operates, as described elsewhere (Rolls 2008; Rolls and Deco 2002).

bottom–up input, though the processes involved include only linear summation within the neurons of bottom–up and top–down synaptic inputs, and the threshold non-linearity of neurons involved in whether an action potential is generated (Deco and Rolls 2005b). These detailed and quantitative analyses of neuronal dynamical systems are an important step towards understanding the operation of complex processes such as top–down attention, which necessarily involve the interaction of several brain areas. They are being extended to provide neurally plausible models of decision-making and action selection (Deco and Rolls 2003, 2005d, 2006; Rolls 2008).

In relation to representation in the brain, the impact of these findings is that they show details of the mechanisms by which representations can be modulated by attention, and, moreover, can account for many phenomena in attention using models in which the firing rate of neurons is represented, and in which stimulus-dependent neuronal synchrony is not involved (Rolls 2007, 2008).

The top–down backprojection pathways between adjacent cortical areas that implement the attentional effects in this model are weak relative to the forward (bottom–up) inputs. Consistent with this, the backprojection synapses end on the apical dendrite of pyramidal cells in the preceding cortical area, quite far from the cell body, where they might be expected to be sufficient to dominate the cell firing when there is no forward input close to the cell body, that is during memory recall, which may be one of the functions of these backprojection pathways, given the very large number of connections and their associative modifiability (Rolls 1989, 2008; Treves and Rolls 1994). In contrast, when there is forward input to the neuron, activating synapses closer to the cell body than the backprojecting inputs, this would tend to electrically shunt the backprojection effects received on the apical dendrites, accounting for their relatively small but useful biasing effect. The associative modifiability is useful for setting up the connectivity required not only for memory recall, but also for top–down attentional effects to influence the correct neurons (Rolls 2008). This is part of the neural architecture of pre-emptive perception.

5 Decision-making in perception

When the surfaces of a Necker cube flip from back to front, it is as if an internal model of the three-dimensional structure of the cube is influencing the representation of the depth of the different surfaces of the cube (Gregory 1970, 1998; Helmholtz 1867/1962; Rolls 2008). When two objects are presented to the visual system in rivalry, first one is seen, and then there is a probabilistic flip to the other object being seen (Maier et al 2005). Again, an internal representation of each object appears to be influencing visual processing so that first the whole of one object, and then of the other object, is seen, and not a combination of the features of both. A recent model of probabilistic decision-making in the brain (Deco and Rolls 2006) contributes to our understanding of these perceptual phenomena, for it is suggested that the underlying computational mechanism may be similar, providing a unifying approach to these aspects of brain processing (Rolls 2008).

The architecture of the model is shown in figure 10, and is within the theoretical framework utilised by Wang (2002), which is based on a neurodynamical model first introduced by Brunel and Wang (2001), and which has been recently extended and successfully applied to explain several experimental paradigms (Deco and Rolls 2002, 2003, 2005b; Deco et al 2004, 2005; Rolls and Deco 2002; Szabo et al 2004). In this framework, probabilistic decision-making is modeled by an attractor network of interacting integrate-and-fire neurons with spiking activity organised into a discrete set of populations, as depicted in figure 10. For a binary decision, there are two populations of neurons, each one corresponding to one of the decisions. Each population has its own biasing input, f_1 and f_2 . The network starts with spontaneous activity, and if the

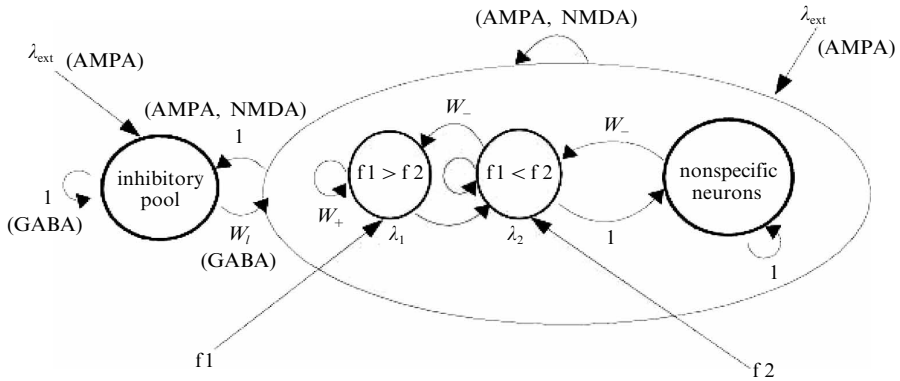


Figure 10. The architecture of the neurodynamical model for a probabilistic decision-making network. The single attractor network has two populations or pools of neurons ($f1 > f2$) and ($f1 < f2$) which represent the decision states. One of these pools becomes active when a decision is made. If pool ($f1 > f2$) is active, this corresponds to the decision that stimulus $f1$ is greater than stimulus $f2$. There is also a population of non-specific excitatory neurons, and a population of inhibitory neurons. Pool ($f1 > f2$) is biased by λ_1 which reflects the strength of stimulus $f1$, and pool ($f2 > f1$) is biased by λ_2 which reflects the strength of stimulus $f2$. The integrate-and-fire network is subject to finite size noise, and therefore probabilistically settles into either an attractor with the population ($f1 > f2$) active, or with the population ($f1 < f2$) active, depending on the biasing inputs λ_1 and λ_2 . The network is thus a biased competition model of decision-making. The weights connecting the different populations of neurons are shown as W_+ , W_- , W_I and 1. All neurons receive a small random Poisson set of input spikes λ_{ext} from other neurons in the system. (Modified from Deco and Rolls 2006.)

biases are equal, what is effectively a biased competition network eventually falls into one of the attractor states (ie with one of the populations of neurons firing with a high rate), with a probability that each attractor wins being 0.5. The probabilistic settling of the network is due to the inherent noise in the finite-size network due to the Poisson-like firing of the neurons (Deco and Rolls 2006). The attractor that wins represents the decision. If one of the biases is stronger than the other, then the probability that the network will reach that decision increases. Because the model is a short-term-memory network, the system can integrate information over long time periods, of hundreds of milliseconds, before a decision is reached. As the biases become more unequal, the reaction time of the decision made by the network decreases. As the biases are both increased, the magnitude of the difference between them for a decision to be reached must be increased in proportion, that is $\Delta I/I = a \text{ constant}$. The network implements Weber's law because as I is increased, the activity of the inhibitory feedback neurons in the integrate-and-fire network increases linearly, and these produce divisive inhibition on the excitatory cells that form the attractor, resulting in a need for ΔI to increase in proportion to the divisive inhibition resulting in $\Delta I/I = a \text{ constant}$.

The decision-making network was tested (Deco and Rolls 2006) against neurophysiological data on decision-making for vibrotactile frequency discrimination. However, it is proposed here that the same type of decision-making network could be implemented in many brain areas, to account for many types of probabilistic decision-making. In the context of visual perception, it is proposed that the two attractor states might represent the two alternative interpretations of which side of a Necker cube is closer, or which binocularly presented image is being seen. Then with some adaptation in the excitatory neurons or the synapses between them, which has been modeled (Deco and Rolls 2005c, 2005d), the firing rate in the currently active attractor would gradually decrease, allowing the other attractor to spontaneously become active in a probabilistic way that depends on the number of spikes that happen to be generated with Poisson-like statistics by the different neurons in the different attractors. In this way, the higher-level representation

of the object, ie the cube or the image, would spontaneously and probabilistically flip, and this higher-level representation would bias the lower-level representations by top-down influences so that first one edge and then the other edge of the cube would be biased to appear close, or the features in one image versus the other would be biased by the top-down competitive influence in pattern (Maier et al 2005) and even binocular rivalry.

It is thus proposed that this model of decision-making (Deco and Rolls 2006) might account for many decision-making processes in the brain, including those involved in the interpretation of visual images. In this sense, what is seen or emphasised at the lower level is biased or pre-empted by the state of the higher-level representations, themselves determined probabilistically. We may note that, in fact, the decision may not be taken only in the high-level network, but could be distributed throughout the system of interconnected networks, with their feedforward and top-down feedback connections all contributing to the decision-making (Deco and Rolls 2006), and all perhaps to the probabilistic change of state according to the extent to which the coupled neurons at different levels of the network show adaptation.

6 Top-down influences on olfactory perception and emotion

Top-down influences can be cognitive, and can operate in other modalities than vision. A dramatic example, in which word-level cognitive influences reach down to bias (or pre-empt) representations, was found recently in the olfactory system, in the first stage at which affect is represented, the human orbitofrontal cortex, as described next (de Araujo et al 2005).

Brie can smell pleasant. However, the same odour taken out of the context of cheese might be unpleasant. There is evidence that the sight (including colour) of a food or wine can influence its flavour. However, what about a more cognitive influence, such as a word? Can this influence the perception and hedonics of food-related stimuli, and if so, how far back down into the sensory system does the cognitive influence reach? To address this, an fMRI investigation was performed in which the delivery of a standard test odour (isovaleric acid combined with cheddar cheese flavour, presented orthonasally with an olfactometer) was paired with a descriptor word on a screen, which on different trials was 'Cheddar cheese' or 'Body odour'. The subjects rated the pleasantness and the intensity of the odour on every trial. Alpha-ionone (pleasant, labelled 'Flowers') and octanol (unpleasant, labelled 'Burned plastic') were used as reference pleasant and unpleasant stimuli for the psychophysics and neuroimaging. Subjects rated the affective value of the test odour as significantly more unpleasant when labelled 'Body odour' than when labelled 'Cheddar cheese'. The medial orbitofrontal cortex (OFC)/rostral anterior cingulate cortex (ACC) was significantly more activated by the test stimulus when labelled 'Cheddar cheese' than when labelled 'Body odour', and these activations were correlated with the pleasantness ratings (de Araujo et al 2005) (see figure 11). This cognitive modulation was also found in the medial amygdala olfactory area, and this extended towards the olfactory tubercle. Thus, cognitive modulation extends in the olfactory system as far down as the secondary olfactory cortex, in the orbitofrontal cortex. The implication is that cognitive factors can have profound effects on our responses to the hedonic and sensory properties of food and other affective stimuli, in that these effects are manifest quite far back into sensory processing, so that at least hedonic representations of odours are affected, and even perceptual representations may be modulated.

An implication of this study and further studies showing cognitive modulation of affective representations in the orbitofrontal cortex of taste and flavour (Grabenhorst et al 2007) and touch (McCabe et al 2008) is that even at the first stage in the processing of visual, taste, olfactory, and somatosensory stimuli at which the emotional or affective value is made explicit in the representation, the orbitofrontal cortex (Rolls 2005b, 2008), top-down cognitive influences can bias representations, so that

the emotion felt can be directly influenced (biased, or mildly 'pre-empted'), at a relatively early level of cortical processing, by the cognitive state. This makes it clear that the types of top-down attentional processing described in this paper for vision apply to other modalities too, and indeed they probably are implemented by the same type of top-down biased competition process (Rolls 2008).

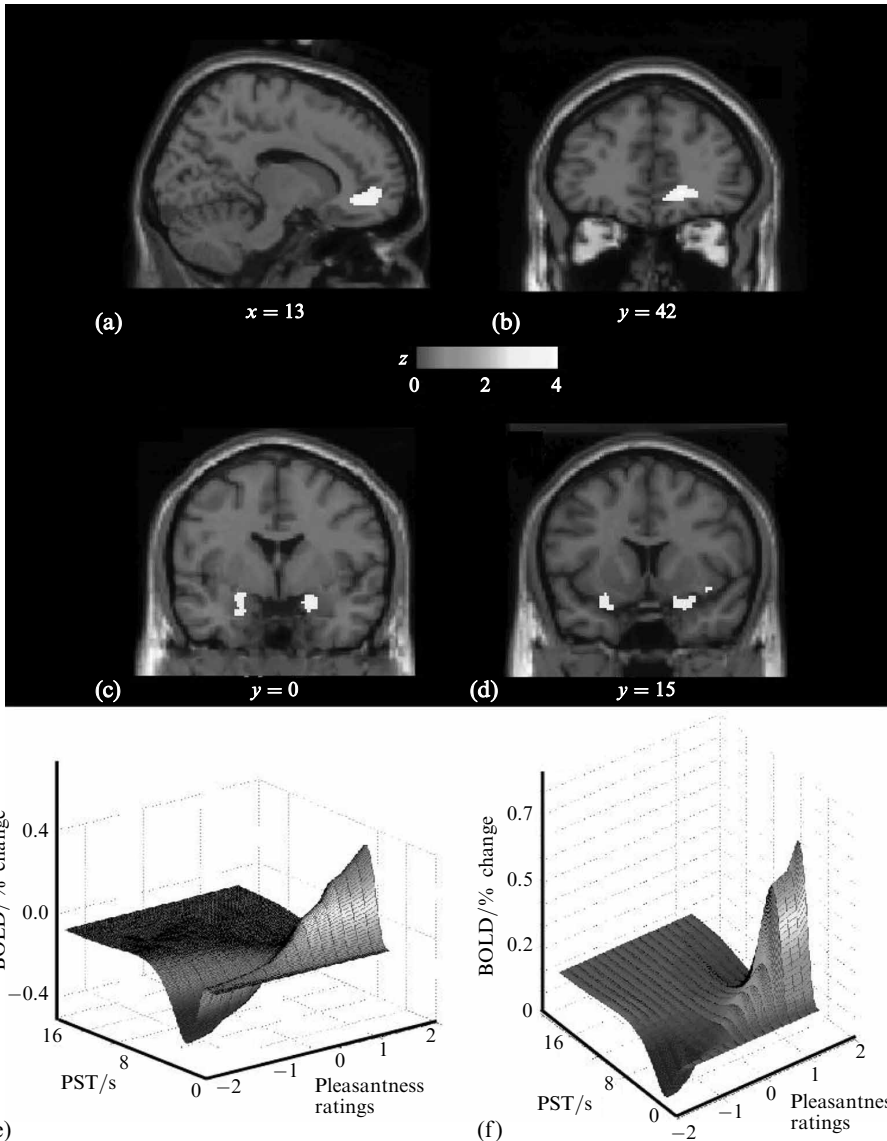


Figure 11. Cognitive influences on olfactory representations in the human brain. Group (random) effects analysis showing the brain regions where the BOLD signal was correlated with pleasantness ratings given to the test odour. The pleasantness ratings were being modulated by the word labels. (a) Activations in the rostral anterior cingulate cortex, in the region adjoining the medial OFC, shown in a sagittal slice. (b) The same activation shown coronally. (c) Bilateral activations in the amygdala. (d) Bilateral activations extended anteriorly to the primary olfactory cortex. The image was thresholded at $p < 0.0001$ uncorrected in order to show the extent of the activation. (e) Parametric plots of the data averaged across all subjects showing that the percentage BOLD change (fitted) correlates with the pleasantness ratings in the region shown in (a) and (b). The parametric plots were very similar for the primary olfactory region shown in (d). PST, post-stimulus time in seconds. (f) Parametric plots for the amygdala region shown in (c). (After de Araujo et al 2005.)

7 Synthesis

It is shown here how top–down attentional processing can interact with bottom–up sensory inputs in a precise and unifying computational-neuroscience approach, which is elaborated elsewhere (Rolls 2008). The computational processes described here provide a way for understanding not only top–down effects in attention, but also top–down effects that are captured by the idea of inference and hypothesis-testing in perception (Gregory 1970, 1998; Helmholtz 1867/1962; Rolls 2008). A representation at a higher level in the sensory processing hierarchy once activated can, via cortico-cortical back-projections, influence earlier representations, and help the earlier levels to settle into an interpretation consistent with higher-level constraints in what is a biased competition process. These attentional and perceptual processes can operate nonlinearly owing to the spiking threshold nonlinearity of neurons, and are most evident when the bottom–up sensory processes are weak or ambiguous, for then the top–down effect, which is necessarily weak so that it does not dominate visual perception, is most evident (Deco and Rolls 2005b; Rolls 2008; Rolls and Deco 2002). The same top–down biased competition process is suggested to account for cognitive effects on perceptual processing not only in vision, but also in other sensory modalities as described in section 6. The same unifying computational-neuroscience approach provides a way to understanding probabilistic effects in perception by a similar biased competition process operating within an attractor network, for in this scenario the inputs for the decision-making can be thought of as biases to the network which influence the retrieval process in the network, with the retrieved attractor state representing the decision taken (Rolls 2008).

A recent model of probabilistic decision-making in the brain (Deco and Rolls 2006) contributes to our understanding of these perceptual phenomena, for it is suggested that the underlying computational mechanism may be similar, providing a unifying approach to these aspects of brain processing (Rolls 2008).

Acknowledgments. The author has worked on some of the investigations described here with Nikolaos Aggelopoulos, Paul Azzopardi, Gordon C Baylis, Hugo Critchley, Gustavo Deco, Peter Földiák, Leonardo Franco, Michael Hasselmo, Julia Hornak, Morten Kringelbach, Christiana M Leonard, Tim J Milward, David I Perrett, Simon M Stringer, Martin J Tovee, Thomas Trappenberg, Alessandro Treves, Guy Wallis, and their collaboration is sincerely acknowledged. Different parts of the research described were supported by the Medical Research Council, PG8513790; by a Human Frontier Science Program grant; by an EC Human Capital and Mobility grant; by the MRC Oxford Interdisciplinary Research Centre in Cognitive Neuroscience; and by the Oxford McDonnell–Pew Centre in Cognitive Neuroscience. This paper is dedicated to two outstanding neuroscientists whom I first met when an undergraduate reading medicine at Cambridge University: Richard Gregory and Colin Blakemore. It is also dedicated to Nick Aggelopoulos, a dedicated neurophysiologist and man of integrity, who worked on some of the experiments described here.

References

- Aggelopoulos N C, Franco L, Rolls E T, 2005 “Object perception in natural scenes: encoding by inferior temporal cortex simultaneously recorded neurons” *Journal of Neurophysiology* **93** 1342–1357
- Aggelopoulos N C, Rolls E T, 2005 “Natural scene perception: inferior temporal cortex neurons encode the positions of different objects in the scene” *European Journal of Neuroscience* **22** 2903–2916
- Araujo I E T de, Rolls E T, Velazco N I, Margot C, Cayeux I, 2005 “Cognitive modulation of olfactory processing” *Neuron* **46** 671–679
- Ballard D H, 1993 “Subsymbolic modelling of hand–eye coordination”, in *The Simulation of Human Intelligence* Ed. D E Broadbent (Oxford: Oxford University Press) pp 71–102
- Biederman I, 1972 “Perceiving real-world scenes” *Science* **177** 77–80
- 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
- Chelazzi L, Miller E, Duncan J, Desimone R, 1993 “A neural basis for visual search in inferior temporal cortex” *Nature* **363** 345–347

- Corchs S, Deco G, 2002 "Large-scale neural model for visual attention: integration of experimental single-cell and fMRI data" *Cerebral Cortex* **12** 339–348
- Cowey A, Rolls E T, 1975 "Human cortical magnification factor and its relation to visual acuity" *Experimental Brain Research* **21** 447–454
- Deco G, Lee T S, 2002 "A unified model of spatial and object attention based on inter-cortical biased competition" *Neurocomputing* **44–46** 775–781
- Deco G, Rolls E T, 2002 "Object-based visual neglect: a computational hypothesis" *European Journal of Neuroscience* **16** 1994–2000
- Deco G, Rolls E T, 2003 "Attention and working memory: a dynamical model of neuronal activity in the prefrontal cortex" *European Journal of Neuroscience* **18** 2374–2390
- Deco G, Rolls E T, 2004 "A neurodynamical cortical model of visual attention and invariant object recognition" *Vision Research* **44** 621–644
- Deco G, Rolls E T, 2005a "Attention, short-term memory, and action selection: a unifying theory" *Progress in Neurobiology* **76** 236–256
- Deco G, Rolls E T, 2005b "Neurodynamics of biased competition and co-operation for attention: a model with spiking neurons" *Journal of Neurophysiology* **94** 295–313
- Deco G, Rolls E T, 2005c "Sequential memory: a putative neural and synaptic dynamical mechanism" *Journal of Cognitive Neuroscience* **17** 294–307
- Deco G, Rolls E T, 2005d "Synaptic and spiking dynamics underlying reward reversal in orbito-frontal cortex" *Cerebral Cortex* **15** 15–30
- Deco G, Rolls E T, 2006 "Decision-making and Weber's Law: a neurophysiological model" *European Journal of Neuroscience* **24** 901–916
- Deco G, Rolls E T, Horwitz B, 2004 " 'What' and 'where' in visual working memory: a computational neurodynamical perspective for integrating fMRI and single-neuron data" *Journal of Cognitive Neuroscience* **16** 683–701
- Deco G, Rolls E T, Zihl J, 2005 "A neurodynamical model of visual attention", in *Neurobiology of Attention* Eds L Itti, G Rees, J Tsotsos (San Diego, CA: Academic Press) pp 593–599, chapter 97
- Deco G, Zihl J, 2001 "Top-down selective visual attention: a neurodynamical approach" *Visual Cognition* **8** 119–140
- Desimone R, Duncan J, 1995 "Neural mechanisms of selective visual attention" *Annual Review of Neuroscience* **18** 193–222
- Elliffe M C M, Rolls E T, Stringer S M, 2002 "Invariant recognition of feature combinations in the visual system" *Biological Cybernetics* **86** 59–71
- Grabenhorst F, Rolls E T, Bilderbeck A, 2007 "How cognition modulates affective responses to taste and flavor: top-down influences on the orbitofrontal and pregenual cingulate cortices" *Cerebral Cortex* doi:10.1093/cercor/bhm185
- Gregory R L, 1970 *The Intelligent Eye* (New York: McGraw-Hill)
- Gregory R L, 1998 *Eye and Brain* (Oxford: Oxford University Press)
- Helmholtz H von, 1867/1962 *Treatise on Physiological Optics* volume 3 (New York: Dover, 1962); English translation by J P C Southall for the Optical Society of America (1925) from the 3rd German edition of *Handbuch der physiologischen Optik* (first published in 1867, Leipzig: Voss)
- McCabe C, Rolls E T, Bilderbeck A, McGlone F, 2008 "Cognitive influences on the affective representation of touch and the sight of touch in the human brain" *Social, Cognitive and Affective Neuroscience* **3** doi:10.1093/scan/nsn005
- Maier A, Logothetis N K, Leopold D A, 2005 "Global competition dictates local suppression in pattern rivalry" *Journal of Vision* **5** 668–677
- Martinez-Trujillo J, Treue S, 2002 "Attentional modulation strength in cortical area MT depends on stimulus contrast" *Neuron* **35** 365–370
- Mozor M, 1991 *The Perception of Multiple Objects: A Connectionist Approach* (Cambridge, MA: MIT Press)
- Renart A, Moreno R, Rocha J de la, Parga N, Rolls E T, 2001 "A model of the IT-PF network in object working memory which includes balanced persistent activity and tuned inhibition" *Neurocomputing* **38–40** 1525–1531
- Renart A, Parga N, Rolls E T, 2000 "A recurrent model of the interaction between the prefrontal cortex and inferior temporal cortex in delay memory tasks", in *Advances in Neural Information Processing Systems 12* Eds S A Solla, T K Leen, K-R Mueller (Cambridge, MA: MIT Press) pp 171–177
- Reynolds J, Desimone R, 1999 "The role of neural mechanisms of attention in solving the binding problem" *Neuron* **24** 19–29

- Reynolds J H, Chelazzi L, Desimone R, 1999 "Competitive mechanisms subserve attention in macaque areas V2 and V4" *Journal of Neuroscience* **19** 1736–1753
- Rolls E T, 1989 "Functions of neuronal networks in the hippocampus and neocortex in memory", in *Neural Models of Plasticity: Experimental and Theoretical Approaches* Eds J H Byrne, W O Berry (San Diego, CA: Academic Press) pp 240–265
- Rolls E T, 1992 "Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas" *Philosophical Transactions of the Royal Society of London B* **335** 11–21
- Rolls E T, 2000 "Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition" *Neuron* **27** 205–218
- Rolls E T, 2003 "Consciousness absent and present: a neurophysiological exploration" *Progress in Brain Research* **144** 95–106
- Rolls E T, 2005a "Consciousness absent or present: a neurophysiological exploration of masking", in *The First Half Second: The Microgenesis and Temporal Dynamics of Unconscious and Conscious Visual Processes* Eds H Ogmen, B G Breitmeyer (Cambridge, MA: MIT Press) pp 89–108, chapter 6
- Rolls E T, 2005b *Emotion Explained* (Oxford: Oxford University Press)
- Rolls E T, 2007 "The representation of information about faces in the temporal and frontal lobes" *Neuropsychologia* **45** 125–143
- Rolls E T, 2008 *Memory, Attention, and Decision-Making: A Unifying Computational Neuroscience Approach* (Oxford: Oxford University Press)
- Rolls E T, Aggelopoulos N C, Zheng F, 2003 "The receptive fields of inferior temporal cortex neurons in natural scenes" *Journal of Neuroscience* **23** 339–348
- Rolls E T, Cowey A, 1970 "Topography of the retina and striate cortex and its relationship to visual acuity in rhesus monkeys and squirrel monkeys" *Experimental Brain Research* **10** 298–310
- Rolls E T, Deco G, 2002 *Computational Neuroscience of vision* (Oxford: Oxford University Press)
- Rolls E T, Deco G, 2006 "Attention in natural scenes: neurophysiological and computational bases" *Neural Networks* **19** 1383–1394
- Rolls E T, Milward T, 2000 "A model of invariant object recognition in the visual system: learning rules, activation functions, lateral inhibition, and information-based performance measures" *Neural Computation* **12** 2547–2572
- Rolls E T, Stringer S M, 2001 "Invariant object recognition in the visual system with error correction and temporal difference learning" *Network: Computation in Neural Systems* **12** 111–129
- Rolls E T, Stringer S M, 2006a "Invariant global motion recognition in the dorsal visual system: a unifying theory" *Neural Computation* **19** 139–169
- Rolls E T, Stringer S M, 2006b "Invariant visual object recognition: a model, with lighting invariance" *Journal of Physiology—Paris* **100** 43–62
- Rolls E T, Tovee M J, 1994 "Processing speed in the cerebral cortex and the neurophysiology of visual masking" *Proceedings of the Royal Society of London B* **257** 9–15
- Rolls E T, Tovee M J, Purcell D G, Stewart A L, Azzopardi P, 1994 "The responses of neurons in the temporal cortex of primates, and face identification and detection" *Experimental Brain Research* **101** 473–484
- Simons D J, Chabris C F, 1999 "Gorillas in our midst: Sustained inattention blindness for dynamic events" *Perception* **28** 1059–1074
- Simons D J, Rensink R A, 2005 "Change blindness: past, present, and future" *Trends in Cognitive Sciences* **9** 16–20
- Stringer S M, Perry G, Rolls E T, Proske J H, 2006 "Learning invariant object recognition in the visual system with continuous transformations" *Biological Cybernetics* **94** 128–142
- Stringer S M, Rolls E T, 2000 "Position invariant recognition in the visual system with cluttered environments" *Neural Networks* **13** 305–315
- Stringer S M, Rolls E T, 2002 "Invariant object recognition in the visual system with novel views of 3D objects" *Neural Computation* **14** 2585–2596
- Szabo M, Almeida R, Deco G, Stetter M, 2004 "Cooperation and biased competition model can explain attentional filtering in the prefrontal cortex" *European Journal of Neuroscience* **19** 1969–1977
- Trappenberg T P, Rolls E T, Stringer S M, 2002 "Effective size of receptive fields of inferior temporal cortex neurons in natural scenes", in *Advances in Neural Information Processing Systems 14* Eds T G Dietterich, S Becker, Z Ghahramani (Cambridge, MA: MIT Press) pp 293–300

-
- Treves A, Rolls E T, 1994 "A computational analysis of the role of the hippocampus in memory" *Hippocampus* **4** 374–391
- Usher M, Niebur E, 1996 "Modelling the temporal dynamics of IT neurons in visual search: a mechanism for top–down selective attention" *Journal of Cognitive Neuroscience* **8** 311–327
- Wallis G, Rolls E T, 1997 "Invariant face and object recognition in the visual system" *Progress in Neurobiology* **51** 167–194
- Wang X J, 2002 "Probabilistic decision making by slow reverberation in cortical circuits" *Neuron* **36** 955–968
- Wurtz R H, Sommer M A, 2004 "Identifying corollary discharges for movement in the primate brain" *Progress in Brain Research* **144** 47–60

ISSN 0301-0066 (print)

ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 37 2008

www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.