King's Research Portal

DOI:

[10.1162/089892904322984526](https://doi.org/10.1162/089892904322984526)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Spratling, M. W., & Johnson, M. H. (2004). A feedback model of visual attention. *JOURNAL OF COGNITIVE NEUROSCIENCE*, 16(2), 219 - 237. <https://doi.org/10.1162/089892904322984526>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Feedback Model of Visual Attention

M. W. Spratling and M. H. Johnson

Centre for Brain and Cognitive Development, Birkbeck College, London. UK.

Abstract

Feedback connections are a prominent feature of cortical anatomy and are likely to have a significant functional role in neural information processing. We present a neural network model of cortical feedback that successfully simulates neurophysiological data associated with attention. In this domain our model can be considered a more detailed, and biologically plausible, implementation of the biased competition model of attention. However, our model is more general as it can also explain a variety of other top-down processes in vision, such as figure/ground segmentation and contextual cueing. This model thus suggests that a common mechanism, involving cortical feedback pathways, is responsible for a range of phenomena and provides a unified account of currently disparate areas of research.

1 Introduction

Top-down effects play an important role in sensory information processing (Siegel et al., 2000). For example, during visual perception, information propagates through the visual processing hierarchy from primary sensory areas to higher cortical regions. In addition to this feedforward transmission of information, feedback connections convey information in the reverse direction and lateral connections integrate information across the visual field (Lamme et al., 1998; Lamme and Roelfsema, 2000).

Feedforward and feedback connections can be distinguished by the cortical layers from which they originate and in which they terminate, as illustrated in figure 1 (Felleman and Van Essen, 1991; Crick and Koch, 1998; Mountcastle, 1998; Budd, 1998; Ebdon, 1996; Mumford, 1992; Lamme et al., 1998; Johnson and Burkhalter, 1997; Barbas and Rempel-Clower, 1997). Feedforward connections are provided by the axon projections of pyramidal cells in layers II and III. These projections terminate predominantly in layer IV of the higher region (as do inputs from the thalamus). The main target for feedforward projections are spiny-stellate cells which, in turn, target the basal dendrites of pyramidal cells in layers II and III. Feedback connections are provided by the axon projections from pyramidal cells in layers V and VI and terminate mainly in layers I and VI of the lower region (or are sent to sub-cortical structures). The main targets of the feedback projections terminating in layer I are the apical dendrites of pyramidal cells with somata in layers II, III and V (Cauller, 1995; Cauller et al., 1998; Rockland, 1998; Budd, 1998; Rolls and Treves, 1998). In general, cortical regions tend to be reciprocally connected (Felleman and Van Essen, 1991; Crick and Asanuma, 1986; Mountcastle, 1998; Lamme et al., 1998).

Attention is one top-down process that operates via the cortical feedback projections (Treue, 2001; Schroeder et al., 2001; Mehta et al., 2000; Desimone and Duncan, 1995; Pollen, 1999) targeting the apical dendrites in layer I (Cauller, 1995; Olson et al., 2001). By manipulating attention or expectation, it is possible to explore the affects of feedback on the response properties of cortical pyramidal cells. Attention modulates the sensory-driven activation of cells (Kanwisher and Wojciulik, 2000; Luck et al., 1997; Kastner et al., 1999; McAdams and Maunsell, 2000), such that activity in response to an attended stimulus is increased in amplitude and duration (Schroeder et al., 2001; Kastner and Ungerleider, 2000). Such top-down modulation also affects the on-going competition between cells (Luck et al., 1997; Itti and Koch, 2001; Reynolds et al., 1999). Competition can be biased both by stimulus saliency and by attention (Kastner and Ungerleider, 2000; De Weerd et al., 1999; Olson, 2001). Hence, increased attention has effects similar to increasing the contrast or saliency of the stimulus (Kastner and Ungerleider, 2000; Olson, 2001; Itti and Koch, 2001; Reynolds et al., 2000). Feedback thus serves “to amplify and focus activity” (Hupé et al., 1998).

All connectionist models of perception investigate how feedforward, sensory-driven, information is processed and represented. However, despite the apparent importance of top-down processes in perception, relatively few models have investigated the role of feedback. In this article we present a neural network model of visual processing which does incorporate feedback connections. Since feedback plays a prominent role in attention we evaluate the performance of our model by comparison with the response properties of cortical pyramidal cells during attentional tasks. Similar effects of feedback, or recurrent, activity on neural responses can also be observed in non-attentional tasks (Lamme et al., 1998; Zipser et al., 1996; Lamme and Roelfsema, 2000; Lee et al., 1998). We thus propose that the same underlying mechanisms are responsible for a variety of other top-down processes in vision. We demonstrate this claim by using the same neural network model to explain empirical data associated with

Figure 1: Cortical layers and regions. A schematic showing pyramidal cells within the six layers of the cortical sheet. Pyramidal cell bodies are shown as filled triangles, dendrites as solid lines and axons as dashed lines. All other cell types have been omitted, as has the intra-regional connectivity. The cortical sheet is shown divided into two regions at different levels in an information processing hierarchy. Axon projections connecting these two regions are illustrated.

figure/ground segmentation, feature binding, and contextual cueing. Hence, our model suggests that a common mechanism, involving cortical feedback pathways, is responsible for a range of phenomena in visual perception and could potentially provide a unified account of currently disparate areas of research.

1.1 The Biased Competition Model

The biased competition model (Duncan, 1998; Desimone and Duncan, 1995; Reynolds et al., 1999) is a leading account of the neurophysiological data associated with attention (Frith, 2001; Kastner and Ungerleider, 2000). This theory proposes that visual stimuli compete to be represented by cortical activity. Competition may occur at each stage along the visual information processing pathway. The outcome of the competition is influenced not only by bottom-up, sensory driven, processes but also by top-down, attention-dependent, biases towards the information that is most relevant to current behavioural goals.

A neural network architecture, to implement this model, has been proposed (Reynolds and Desimone, 1999; Reynolds et al., 1999), and is shown in figure 2(a). In this network, two “input” neurons, with distinct stimulus selectivities, project both excitatory and inhibitory connections to a single “output” neuron in a subsequent stage of the cortical hierarchy. The response of the output node is dependent upon the activation of each of the input nodes and the strengths of the excitatory and inhibitory afferents. Attention acts by increasing the synaptic efficiencies of the connections originating from the input neuron selective to the attended stimulus.

This neural network architecture was “not intended to be an account of the actual neural circuitry underlying visual attention” (Reynolds et al., 1999, p.1752). It is therefore unsurprising that while this implementation succeeds as a descriptive model of the observed behaviour it has shortcomings as a detailed, functional, account. Firstly, the time-varying response of this model is a poor match to that of real cells (see section 2). Secondly, attention acts by multiplicatively modulating the strengths of afferent synapses, for which no biologically plausible mechanism is specified. Finally, while this network is proposed as an implementation of the biased competition model there are actually no competitive processes occurring between nodes within this neural architecture.

Reynolds and Desimone (1999) suggest that there may be many other possible implementations of the biased competition model. Therefore, we shall refer to the particular neural network architecture described above as the ‘Reynolds and Desimone model’ to distinguish it from other implementations and from the theory of biased competition itself. In the following section we propose an alternative implementation, which is illustrated in figure 2(b). In this neural network, nodes within a cortical region compete via lateral inhibitory connections and top-down signals modulate the activations of nodes, rather than the efficiencies of synapses. As with the Reynolds

Figure 2: Models of attention. Nodes are shown as large circles, excitatory synapses as small open circles and inhibitory synapses as small filled circles. The x values represent top-down biases that vary in strength depending on the attentional state. These top-down signals are assumed to arise from neural generators outside the modelled circuits. (a) The Reynolds and Desimone implementation of the biased competition model. In this model, the top-down signals modulate the strength of the feedforward connections. (b) The proposed feedback model. In this model, the top-down signals modulate the strength of activation of the nodes. Feedback connections reciprocating the feedforward connections also exist, but have been omitted from this figure for clarity.

and Desimone model, the top-down signals vary in strength depending on the attentional state. These attention-dependent, top-down, signals are assumed to originate in higher cortical regions which are not explicitly modelled. However, this should not be taken to imply that these signals arise from a dedicated “attentional system”, a point we return to in the Discussion.

The nodes employed by the proposed model have more complex behaviour than those used in the Reynolds and Desimone model. Hence, the feedback model is capable of simulating experimental data in greater detail. More importantly, the model proposes a biologically plausible mechanism by which biased competition can operate, and can explain other top-down, perceptual, processes that have not previously been thought to share common mechanisms.

1.2 The Feedback Model

The Reynolds and Desimone model, in common with others (*e.g.*, Olshausen et al., 1993), uses top-down signals to multiplicatively modulate the synaptic strengths of inter-regional connections so that attended information can be selectively routed to higher cortical regions. Equivalent results can be achieved by using top-down signals to modulate the activity of neurons rather than weights of synapses (Salinas and Thier, 2000; Salinas and Abbott, 1997). This mechanism has previously been modelled by allowing the activity generated by stimulation of the receptive field to be multiplicatively modulated by the response to a separate set of inputs applied to a “gain field” (Salinas and Abbott, 1996; Salinas and Thier, 2000; Salinas and Sejnowski, 2001) or a “contextual field” (Phillips and Singer, 1997; Phillips et al., 1995). These algorithms thus require inputs from different sources to be integrated separately and to have dissimilar effects on activity.

Feedforward and feedback connections preferentially target distinct regions of pyramidal cell dendrites (Spratling, 2002). For example, pyramidal cells in layers II and III predominantly receive feedforward information at the basal dendrites and feedback information at the apical dendrites. The apical dendrite appears to act as a functionally distinct dendritic compartment since activation applied to the apical dendrites is integrated prior to transmission to the soma (Larkum et al., 1999; Körding and König, 2000, 2001). The distal (*i.e.*, apical) and proximal (mainly basal) dendrites of pyramidal cells thus appear to act as separate dendritic compartments (Yuste et al., 1994) capable of independently integrating the feedforward and feedback information that they receive (Spratling, 2002). The axon initial segment acts as the final site of integration as it is here that action potential initialization occurs (Stuart et al., 1997). However, inputs to the different dendritic regions contribute to this output in different ways (Larkum et al., 2001). Stimulation of the apical dendrite causes smaller, but more protracted, excitatory post-synaptic potentials (EPSPs) at the soma than does equivalent stimulation of the basal dendrites (Rockland, 1998; Budd, 1998). Hence, apical inputs have weaker effects on output activity than basal inputs. Such findings are consistent with the suggestion that feedback acts to modulate responses that are primarily driven by feedforward inputs (Koch and Segev, 2000; Crick and Koch, 1998; Hupé et al., 1998; Friston and Büchel, 2000).

We propose a model in which feedback stimulation is integrated in the apical dendrite, and feedforward information is separately integrated in the basal dendrite. The total strength of the top-down activation is then used to multiplicatively modulate the total strength of feedforward activation to determine the final response of the

Figure 3: The feedback model in greater detail. A schematic of the proposed model showing two interacting cortical regions. Each region contains one layer of pyramidal cells (somata shown as filled triangles). These model neurons each have two, independent, dendritic regions: the basal dendrites, which receive feedforward connections (shown as solid lines), and the apical dendrites, which receive feedback connections (shown as dashed lines). Nodes in different regions are reciprocally connected by feedforward and feedback projections (only the connections between one pair of nodes are shown) and may also be targeted by connections originating from other regions or by sensory inputs. Nodes within the same region compete via lateral inhibitory connections which target the basal dendrites (only selected connections between neighbouring nodes are shown). Excitatory and inhibitory synapses are shown as open and filled circles respectively.

node. This formulation enables bottom-up, sensory-driven, stimulation to drive the response of the node even in the absence of top-down activity. In contrast, feedback activation can not drive the node’s activity in the absence of feedforward activation, but it can amplify any response to feedforward stimulation. The activations of nodes in the model are also affected by lateral inhibitory connections targeting the basal dendrites, via which neurons in the same region compete to represent stimuli (Spratling and Johnson, 2001, 2002). A schematic of the model is shown in figure 3. For comparison with the Reynolds and Desimone model, a simplified illustration of the feedback model with two input nodes and two output nodes, is also shown in figure 2(b).

For simplicity the model only contains one layer of cells in each region. Hence, in contrast with the cortex, feedback connections originate from the same neurons that receive feedforward connections. This simplification is justified since pyramidal cells within the same cortical column are believed to have similar response properties. Furthermore, there is insufficient data about the differences between the layers to meaningfully model distinct superficial and deep layer neurons.

2 Results

Both electro-physiological and psychophysical experiments on attention can be classified into those which explore spatial attention (the effects of directing attention to one location rather than another) and those which explore featural, or object-based, attention (the effects of directing attention to one object, or stimulus feature, rather than another). In addition, experiments may be classified into those which address attentional selection (the role of attention in the selection of one item out of many) and those which address attentional facilitation (the enhancement of processing that occurs for an individual stimulus at the focus of attention). In sections 2.1 and 2.2 we demonstrate that the feedback model, described above, can successfully simulate results from these different experimental paradigms. In addition, we demonstrate in section 2.3 that the same model can explain experimental data associated with figure/ground segmentation and contextual cueing.

Figure 4: The effect of spatial attention on the response of a neuron. Responses are shown for different combinations of stimuli appearing within the receptive field of a single neuron. (a) The response of a cell in V2 (adapted from [Reynolds et al., 1999](#)). (b) Simulation results from the Reynolds and Desimone model. (c) Simulation results from the feedback model. For (a) the response was measured in spikes per second and time in milliseconds, for (b) and (c) both response and time are in arbitrary units and have been scaled to resemble (a).

For clarity, in each simulation we have employed the simplest network which can successfully model the physiological data. Hence, appropriate numbers of nodes and synaptic weight values have been chosen for each experiment. However, it should be noted that the architecture of the model remains constant and that every node in every reported simulation operates in the same way (as described in the Methods section). Furthermore, the same parameters have been used throughout all the simulations reported here. Hence, each simulation can be considered to be a specific example of the behaviour generated by of a single computational model.

2.1 Spatial Attention

2.1.1 Attentional Selection

For neurons in the ventral pathway the response to a stimulus, that generates strong neural activation when presented in isolation, is reduced by the introduction of a second, non-preferred, stimulus within the receptive field ([Reynolds et al., 1999](#)). Hence, rather than being processed independently, multiple stimuli, within the same receptive field, appear to compete in a mutually suppressive manner ([Kastner and Ungerleider, 2000](#)). If attention is directed toward one stimulus then the response becomes more similar to the response that would be generated by that stimulus in isolation ([Moran and Desimone, 1985](#); [Luck et al., 1997](#); [Reynolds et al., 1999](#)). Hence, attention appears to bias the competition in favour of the attended stimulus. These effects are illustrated in figure 4(a), which shows the response of a single cell recorded in area V2 of a rhesus monkey. Similar results have been demonstrated for cells in area V4, inferior temporal cortex, in area MT of the dorsal pathway, and in prefrontal cortex ([Treue and Martinez-Trujillo, 1999](#); [Reynolds et al., 1999](#); [Reynolds and Desimone, 1999](#); [Everling et al., 2002](#)).

The firing rate of the cell is shown in response to a preferred stimulus, to a poor stimulus, and to both stimuli when attention is directed to a location outside the receptive field. When attention is directed to the location occupied by the preferred stimulus in the pair, the response of the cell becomes more similar to that elicited by the preferred stimulus in isolation. The neural network architecture for the Reynolds and Desimone model, shown in figure 5(a), was used to simulate these results. For this simulation one input node was activated by the preferred stimulus and the other input node was activated by the poor stimulus. The input nodes had different connection strengths with the output node to reflect the stimulus selectivity of the recorded neuron. This architecture generated the results shown in figure 4(b). It can be seen that the steady-state responses from the model are in good agreement with the experimental data. When both stimuli are present, the output node receives strong excitatory input from the preferred stimulus and strong inhibitory input from the poor stimulus, resulting in a response that is

Figure 5: Details of the neural network architectures used to simulate the data shown in figure 4. Nodes are shown as large circles, excitatory synapses as small open circles and inhibitory synapses as small filled circles. The x values represent top-down biases that have an activation value of either zero or one depending on the attentional state. These top-down signals are assumed to arise from neural generators outside the modelled circuits. Feedforward (and lateral) connections are shown as solid lines and feedback connections are shown as dashed lines. Values for the synaptic strengths are indicated, these weights were chosen to provide the best fit between the behaviour of each model and the experimental data (see the Methods section). (a) The Reynolds and Desimone model. (b) The feedback model.

intermediate between the responses generated by each stimulus in isolation. When attention is directed to the preferred stimulus this has the effect of multiplying the synaptic efficiencies from the attended input by a factor of five, and causes the response of the output node to become dominated by the attended stimulus.

This experiment was also simulated with the feedback model, using the network shown in figure 5(b). For this simulation, each output node had a preference to a different one of the two inputs. The input nodes received feedback connections from the output nodes which had weights proportional to the corresponding feedforward connections. In addition, the input nodes each received a feedback connection, of weight 0.5, from different attention-dependent sources of feedback. The results of the simulation, showing the response of one output node, are shown in figure 4(c). When only one stimulus is presented to the network, the output node with the preference to this input quickly wins the competition. Hence, the recorded node is strongly active in response to its preferred stimulus, but has a weak and brief response to its non-preferred stimulus. When both stimuli are present, both output nodes are strongly activated and there is on-going competition between them. Hence, the response of the recorded node to the pair of stimuli is less than its response to its preferred stimulus in isolation. When attention is directed to the preferred stimulus, one of the top-down signals is active, modulating the activity of one input node so that the feedforward activation received by the output nodes is stronger for one input than the other. The output of the node with the preference for this amplified-input is therefore enhanced.

Note that for both models a similar effect would result if the activities in the two input nodes were made unequal, not by differences in the strength of the attention-dependent biases, but by a difference in the strength of the feedforward activation each node received, for instance, due to the stimuli having unequal contrasts. Hence, a strong bottom-up signal can bias competition in just the same way that top-down modulation can. Such effects have been observed experimentally (Reynolds and Desimone, 1999; De Weerd et al., 1999; Kastner and Ungerleider, 2000; Vecera, 2000). Furthermore, training has been observed to increase the apparent saliency of a stimulus (Jagadeesh et al., 2001). This effect is explained since training is likely to enhance the neuronal representation of the training stimulus and hence provide that stimulus with a competitive advantage (Jagadeesh et al., 2001).

2.1.2 Attentional Facilitation

When attention is directed to a particular location, the processing of a stimulus appearing at that location is enhanced (Reynolds et al., 2000; Kastner et al., 1999; Kirschfeld and Kammer, 2000). Increased attention has effects similar to increasing the contrast or saliency of the stimulus (Kastner and Ungerleider, 2000; Olson, 2001; Itti and Koch, 2001; Reynolds et al., 2000). Results from one experiment illustrating this effect are shown in figure 6(a). In this experiment only one stimulus appeared within the receptive field of the recorded cell. The response of the cell was measured when the location of this stimulus was attended and when it was not attended for varying stimulus contrasts. The data shown was measured from area V4 of the rhesus monkey.

The Reynolds and Desimone model (implemented by the network architecture of figure 7(a)) accounts for attentional facilitation by modulating the synaptic efficiencies of the afferents projecting to the recorded neuron. Simulation results for this model are shown in figure 6(b). The feedback model can also simulate the effects of facilitation by amplifying the activation of the attended input node. Results for the feedback model, generated

Figure 6: The effect of changing stimulus contrast on spatial attention. Responses to a single stimulus within the receptive field of the neuron are shown with and without attention. The contrast of the stimulus increases from left to right. (a) The averaged response for a population of cells in V4 (adapted from Reynolds et al., 2000). (b) Simulation results from the Reynolds and Desimone model. (c) Simulation results from the feedback model. For (a) the response was measured in spikes per second and time in milliseconds, for (b) and (c) both response and time are in arbitrary units and have been scaled to resemble (a).

Figure 7: Details of the neural network architectures used to simulate the data shown in figure 6. Nodes are shown as large circles, excitatory synapses as small open circles and inhibitory synapses as small filled circles. The x values represent top-down biases that have an activation value of either zero or one depending on the attentional state. These top-down signals are assumed to arise from neural generators outside the modelled circuits. Feedforward connections are shown as solid lines and feedback connections are shown as dashed lines. Values for the synaptic strengths are indicated, these weights were chosen to provide the best fit between the behaviour of each model and the experimental data (see the Methods section). (a) The Reynolds and Desimone model. (b) The feedback model.

using the neural network architecture shown in figure 7(b), are plotted in figure 6(c).

Both neural network architectures successfully simulate the change in magnitude of attentional facilitation with stimulus contrast. Attention produces the largest change in response at intermediate stimulus contrasts and has only a weak effect at high contrast. Since sub-optimal stimuli were used to generate the physiological data it appears that weak attentional modulation at high contrast is not due to the firing rate of the recorded cell reaching saturation. The two models explain this effect in different ways. For the Reynolds and Desimone model, as the activation of the input node increases, the response of the output node tends towards the maximum response that is possible given the balance between the excitatory and inhibitory afferents. Further increasing the effective strength of the input via attention has a diminishing effect as this upper limit is reached. In contrast, in the feedback model it is saturation of the response of the node in the lower region that result in attention having less effect at high contrast.

While both models account equally well for the the general effects of attentional facilitation, the models can be distinguished when compared to the finer detail of the physiological data. It can be seen that for the biological data, response latencies reduce with increased stimulus contrast, but are not affected by attention (in this respect increasing attention is unlike increasing the stimulus contrast). For the Reynolds and Desimone model, response latencies correctly reduce with increased stimulus contrast, but also with attention. For the feedback model, response latencies are marginally reduced by stimulus contrast but not by attention. The feedback model thus provides a better fit to the detailed experimental data.

It can also be seen from the physiological data that attentional modulation occurs at longer latencies with increasing contrast. Reynolds et al. (2000) state that this effect can be explained by the the biased competition model, however, the implementation they propose fails to do so. Instead, with their neural network architecture the effect of attention at high contrast is at least as strong at short latencies as it is at longer latencies. In contrast, the feedback model does successfully simulate this aspect of the experimental data. It does so due to the input node reaching saturation during the early part of its response at high contrast. The top-down activation is therefore unable to amplify the early activity of the input node. However, at longer latencies the response of the input node is attenuated (see equation 5 in the Methods section) enabling attention to modulate activity at longer latencies.

In contrast with the above experiment, other physiological data more directly illustrates that attention multiplicatively modulates neural response. A cell will generate a range of different firing rates when its response is measured across a range of different stimuli. For example, a cell with selectivity for stimulus orientation will generate its maximum response at one stimulus orientation and progressively weaker responses to stimuli at orientations that increasingly deviate from the preferred one. Attention has a purely multiplicative effect changing the height, but not the width, of such a tuning curve, such that the percentage increase in response is approximately the same for each stimulus (Treue, 2001; Treue and Martinez-Trujillo, 1999; McAdams and Maunsell, 1999). Hence, for these experiments the largest, absolute, increase in response is for the stimulus that generates the strongest activation. These results initially appear to be at odds with those described above. However, both models can explain this data. For the Reynolds and Desimone model, afferents from strong and weak stimuli will have a different balance of excitatory and inhibitory weights. For a given stimulus contrast, attentional modulation of these weights will have the same relative effect on the response of the output node to each stimulus. Similarly, for the

Figure 8: The effect of featural attention on the response of a neuron. Responses are shown for different combinations of stimuli appearing within the receptive field of a single neuron. (a) The averaged response for a population of cells in V4 (adapted from Chelazzi et al., 2001). (b) Simulation results from the Reynolds and Desimone model. (c) Simulation results from the feedback model. For (a) the response was measured in spikes per second and time in milliseconds, for (b) and (c) both response and time are in arbitrary units and have been scaled to resemble (a).

feedback model, the activation of the input nodes corresponding to weak and strong stimuli would be modulated equally by attention, but these activation values would drive the output neuron via differently weighted afferents and hence cause attention to have a proportionate effect on the response of the output node. Alternatively, with the feedback model, the activation of the recorded cell might be directly modulated by an attention-dependent, top-down, bias.

2.2 Featural Attention

Rather than specifying the spatial location that is to be attended, it is also possible to experiment with the effects of cueing the target object that is to be attended. In one such experiment (Chelazzi et al., 2001), rhesus monkeys were presented with an array containing one or two objects, one of which might have previously been cued as the target for a saccade. Responses were measured from cells in area V4, with receptive fields sufficiently large to encompass the stimulus array. Different responses were generated when the target object was the preferred stimulus of the recorded cell compared to when the target was a non-optimal stimulus (see figure 8(a)). Results are similar to those for attentional selection using spatial cues (presented above) in that when the stimulus array contains a pair of objects the response of the cell becomes more similar to the response that would be generated by the attended stimulus in isolation.

Figure 8(b) shows the simulation of this data generated using the Reynolds and Desimone model. The details of the neural network architecture used are shown in figure 9(a). It can be seen that this architecture is identical (except for the synaptic weight values) to that used previously to simulate the effects of spatial selection (see figure 5(a)). The only difference between these simulations are the combinations of sensory and top-down signals used to generate the results. However, it might be expected that the source of feedback should differ between featural and spatial attention. For spatial attention we would expect that feedback to ventral areas is transmitted via the dorsal stream. Differential feedback would be received by ventral regions with receptive fields at the appropriate scale to define the attended region. In contrast, for featural attention we would expect feedback signals to be transmitted via the ventral pathway. Hence, featural attention should modulate nodes in the higher region and be transmitted down to the lower region via the feedback connections between these regions. While this can not be simulated using the Reynolds and Desimone model, it can be using the feedback model. Figure 9(b) shows the neural network architecture used to simulate this data using the feedback model, and the results of this simulation are shown in figure 8(c). In contrast to the network used to simulate the spatial attention experiments, presented previously, in this implementation, top-down signals are employed which provide biases to nodes in the

Figure 9: Details of the neural network architectures used to simulate the data shown in figure 8. Nodes are shown as large circles, excitatory synapses as small open circles and inhibitory synapses as small filled circles. The x values represent top-down biases that have an activation value of either zero or one depending on the attentional state. These top-down signals are assumed to arise from neural generators outside the modelled circuits. Feedforward (and lateral) connections are shown as solid lines and feed-back connections are shown as dashed lines. Values for the synaptic strengths are indicated, these weights were chosen to provide the best fit between the behaviour of the model and the experimental data (see the Methods section). (a) The Reynolds and Desimone model. (b) The feedback model.

higher cortical region¹. When a single stimulus is presented to the network the node with the preference to that stimulus wins the competition and inhibits the output of the other node. When both stimuli are presented, both nodes in the higher region receive equal feedforward activation. In this condition even a small top-down signal to one node, or the other, can bias the competition in favour of that node.

2.3 Familiarity and Context

In all the tasks described above, the attended object or location was cued prior to stimulus onset. Top-down activity thus corresponded to an expectation about the content of the test stimulus. In other tasks the required focus of attention may not be defined *a priori*, but may be influenced entirely by the content of the stimulus itself. Similarly, in non-attentional tasks, the content of the visual scene may influence the processing of specific stimuli. Perception will be influenced both by bottom-up biases, such as stimulus saliency, and by top-down influence such as familiarity (Olson et al., 2001; Chun, 2002).

In this section we present various experiments on the effect of familiarity on neural activity in the feedback model. In these simulations two regions of neurons are reciprocally connected by feedforward and feedback connections, as illustrated in figure 10(a). The lower region receives sensory input in the form of an eight-by-eight pixel image. Nodes in this region have small receptive fields and are selective to simple patterns within the input image (short horizontal and vertical bars, examples of which are shown in figure 10(c)). Nodes in the higher region receive feedforward connections from all the nodes in the lower region. These nodes thus have larger receptive fields and can learn to become selective to larger, more complex, patterns within the input (stimuli composed of several individual bars, as shown in figure 10(b)). The selectivity of the nodes in the upper region is determined by prior experience. Nodes learn to become selective to frequently reoccurring patterns within the input data. Hence, unlike previous simulations, where appropriate synaptic weights were set by hand, in these simulations synaptic weights are learnt². Furthermore, in contrast to the simulations presented above, where the top-down signals were generated by sources outside the modelled circuits, in this section the simulations do not make use of external sources of feedback. Instead, the effects of feedback processes are generated by the recurrent connectivity between the two model cortical regions.

Figure 11 shows the response of a single node in the lower region under different conditions. The recorded node is selective to the bar outlined in the images shown at the top of the figure. The response of this node when different patterns are presented to the network is shown. Note that in each case the recorded neuron receives an identical sensory input, however, the context within which this image feature appears changes. Before training,

¹Rather than considering this a change in the network architecture, we can consider that the same architecture was used in previous experiments, but that the top-down signals targeting nodes in the upper region were silent, and hence could be ignored. Similarly, in the current experiment, top-down signals targeting the lower region are irrelevant and have been ignored.

²See (Spratling and Johnson, 2002) for details of the learning algorithm used. Feedback connections from the upper to the lower region were also learnt. An activity-dependent learning rule was used which was equivalent to the one employed for training the feedforward connections. This resulted in corresponding feedforward and feedback weights having similar strengths.

Figure 10: (a) The neural network architecture used for experiments on familiarity and context. Nodes in the lower region receive feedforward connections from an eight-by-eight pixel input image. Nodes in the upper region receive feedforward connections from all the nodes in the lower region. Reciprocal feedback connections also exist between nodes in the upper region and nodes in the lower region. (c) Examples of image features to which nodes in the lower region are selective. (b) The patterns to which nodes in the upper region become selective after training.

Figure 11: The effect of familiarity on the response of a simulated neuron in the feedback model. The two patterns shown at the top of the figure were presented to the network and the response of a node in the lower region, which was selective to the outlined bar, was recorded. Responses were recorded (a) before, and (b) after a node in the upper region was trained to respond to the familiar pattern. Before training, the recorded node generates equal responses to both stimuli. After training, the node generates an enhanced response to the familiar stimulus compared to the novel stimulus. The insets in both figures show the activity, in response to both patterns, of the node in the upper region which becomes selective to the familiar pattern after training. Before training, this upper region node generates a weak response to both stimuli. After training, the upper region node produces a much stronger response to the familiar stimulus compared to the novel stimulus. This difference in activity in the upper region results in the modulation observed in the response of the lower region node. Both response and time are measured in arbitrary units, but the same scale has been used in each plot.

Figure 12: The effect of familiarity on the responses of two simulated neurons in the feedback model. The pattern shown at the top of the figure was presented to the network and the responses of two nodes in the lower region, were recorded. The recorded nodes were selective to the two outlined bars indicated, each of which formed part of a different familiar pattern (shown in figure 10(b)). The presented pattern provided equal sensory-driven activity to both nodes but contained only one of the familiar patterns. (a) Before training, each node generates equal responses to the stimulus. (b) After training, one node generates an enhanced response to the bar that is part of the presented familiar stimulus, compared to the other node. The insets in both figures show the responses of the two nodes in the upper region which become selective to the familiar patterns after training. Before training, both upper region nodes generate a weak response to the stimulus. After training, the upper region node that is selective to the presented familiar pattern produces a much stronger response to the stimulus compared to the other node. This difference in activity in the upper region results in the modulation observed in the responses of the lower region nodes. Both response and time are measured in arbitrary units, but the same scale has been used in each plot.

nodes in the upper region are equally unselective to all patterns, and hence, the recorded node receives similar top-down activation in each case and the response of the recorded node is the same for each stimulus. However, after training, one node in the higher region becomes selective to the familiar pattern. The feedback from this node modulates the activation of the recorded node, such that after training the recorded node shows an enhanced response to a familiar pattern compared to a novel pattern. A similar enhancement to the response of the node would be expected in a task exploring feature-based attentional facilitation.

Note that in this experiment, response enhancement is driven by the stimulus and so it could be described as resulting from a bottom-up process. However, the modulation results from feedback to the recorded node, and is determined by the response properties of neurons in a higher-level region, so that, in this sense, it is a top-down process. For the recorded neuron, the sensory-driven stimulation is identical in each case, and changes in the response are due to stimulus features that occur outside its receptive field. The response of a cortical cell to a particular stimulus is known to be highly dependent on the context in which that stimulus appears (Gilbert et al., 2000). While stimuli outside the receptive field are incapable of generating a response when presented in isolation, they can modulate the activation generated by stimuli appearing within the receptive field. Many of these effects are due to long-range, horizontal, connections intrinsic to a cortical region (Li, 2000, 1999; Somers et al., 1998; Stemmler et al., 1995; Gilbert et al., 1996; Lamme et al., 1998). However, similar modulatory, contextual, effects also result from feedback (Zipser et al., 1996; Hupé et al., 1998; Gilbert et al., 2000; Olson et al., 2001).

One role for such feedback effects is figure/ground segmentation (Roelfsema et al., 1998, 2000; Hupé et al., 1998; Zipser et al., 1996; Lamme, 2000). Spatially distributed image features may be grouped together into a whole and segmented from the background by feedback modulation which provides the neurons representing these elements with a competitive advantage over neurons representing other image features (Reynolds and Desimone, 1999). Physiological data shows that the response of a node is enhanced when its receptive field is within the figure compared to when it receives identical stimulation from a location on the ground (Lee et al., 1998; Lamme,

Figure 13: The effect of familiarity on noise suppression in the feedback model. The pattern shown at the top of the figure was presented to the network and the responses of two nodes in the lower region, were recorded. The recorded nodes were selective to the two outlined bars indicated, each of which formed part of a different familiar pattern. The presented pattern was a corrupted version of one of the familiar patterns, such that each node received equal sensory-driven activity. (a) Before training each node generates equal responses to the stimulus. (b) After training one node generates an enhanced response to the bar that is part of the presented familiar stimulus, compared to the other node. Both response and time are measured in arbitrary units, but the same scale has been used in each plot.

1995; Zipser et al., 1996; Lamme et al., 1998, 2000; Lamme and Spekreijse, 2000). The enhancement to activity occurs during the sustained response of the neuron at longer latencies. This late onset of modulation could be due either to response saturation during the initial burst of activity (as postulated to explain the results shown in figure 6) or it could result from a delay in receiving feedback from higher regions (Treue, 2001).

Behavioural results show that figure/ground segmentation is influenced by experience such that a familiar shape is more likely to be perceived as the figure (Peterson et al., 1991). We thus modified the previous task to simulate figure/ground segmentation. In this case the responses of two different nodes in the lower region are plotted when a single stimulus is presented. The stimulus is the superposition of a familiar and a novel pattern (see figure 12). Both recorded nodes are selective to a bar which is contained within different familiar patterns. However, the presented pattern contains only one of those familiar patterns (together with the novel pattern). Before training, nodes in the higher region are unselective and hence both nodes receive equally weak feedback and respond similarly to the stimulus. After training the stimulus evokes a strong response from the node in the higher region which is selective to the embedded familiar pattern. This node provides feedback to one of the recorded nodes which enhances the response of this node in comparison to the other recorded node. The result in this case is similar to the effect that would be expected in an attentional selection task.

Similar results of training on a search task have been described by Lee et al. (2002). In this experiment the responses of cells in primary visual cortex were found to be dependent on experience and the behavioural relevance of stimuli. These cells only became sensitive to certain odd-ball stimuli defined by shape-from-shading after training. Perceptual pop-out saliency depended on experience and appeared to result from feedback from area V2. Higher-order stimulus attributes were thus shown to influence lower-level processing.

Figure 13 shows results for an experiment identical to the one above, except that a different stimulus was used. Again, with this stimulus, both recorded neurons receive the same feedforward stimulation. However, in this case the stimulus contains nearly all the pixels that make up a familiar pattern, plus a few extra active pixels. The missing and additional pixels can be interpreted as noise in the image. The responses of two neurons in the lower region are shown. One node receives input from a bar that forms part of the familiar pattern, but in which a pixel is missing. The other node receives input from the additional pixels. Prior to training, both nodes generate the same response. After training, the node that responds to the bar which is part of the familiar pattern has an enhanced response compared to the activation evoked in the other node by the noise. Hence, higher-level knowledge about previous events, stored in the feedforward synaptic weights, can provide top-down information to enable familiar

stimuli to be represented more strongly than unfamiliar stimuli or background noise.

Attention has been suggested as one possible mechanism for solving the ‘binding’ problem (Treisman, 1998; Reynolds and Desimone, 1999; Roelfsema et al., 2000; Luck and Ford, 1998). One proposal is that spatial attention acts to enhance the processing of stimulus features, across multiple dimensions, and hence provide a cue for grouping together all the features originating from a single object at one location (Treisman, 1998). More generally, top-down information could act to disambiguate sensory data for which there would otherwise be multiple interpretations. Such a proposal is compatible with the feedback model since the competition between nodes representing incompatible interpretations of the sensory data could be biased by top-down signals. This would enhance the responses of nodes compatible with the biased interpretation and enable these nodes to suppress the activity of nodes representing other possible interpretations. Similarly, sources of feedback, other than those generated by attentional demands, could also act to bias the interpretation of feedforward information. For example, figure 14 illustrates how contextual information from outside a neuron’s receptive field could bias the interpretation of images presented to the network described throughout this section. We introduce an ambiguous situation into the previous task by making all the pixels in one quadrant of the image active. Given the selectivities of the nodes in the lower region, this input might be interpreted either as four horizontal bars or four vertical bars. Hence, due to the presence of multiple objects in the image, there is a need to group together features that belong to a single object, *i.e.*, to determine whether a particular pixel should be grouped with other pixels in the same row, or with other pixels in the same column.

Without any feedback, the competition between the nodes representing horizontal bars and those representing vertical bars is resolved entirely by the relative strengths of the feedforward connections, or in the event of all weights being equal, the competition is resolved randomly (due to the noise in the node activations). However, if contextual information is provided, by presenting the ambiguous features in combination with a familiar pattern, then the competition is biased in favour of one interpretation or the other (see figure 14). Feedback from the node in the higher region, activated by the familiar pattern, modulates the activity of all nodes which are consistent with this familiar pattern. One node with a receptive field within the lower-left quadrant thus receives feedback activation. This additional, top-down information is sufficient to bias the competition in favour of either the horizontal or vertical bars, as appropriate to the context. Similar results would be expected if feedback was provided by an attention-dependent top-down signal rather than being generated by a familiar stimulus context.

3 Discussion

We have presented a model of visual attention in which neurons compete to respond to stimuli. The outcome of this competition is influenced not only by bottom-up, sensory driven, processes but also by top-down, attention-dependent, biases. Several other models of attention are based on the same principles of competition and cooperation (*e.g.*, Rolls and Deco, 2002; Deco et al., 2002; Corchs and Deco, 2002; Hamker, 1999, 2002; Usher and Niebur, 1996; Grossberg and Raizada, 2000; Phaf et al., 1990) and some of these models have been applied to simulating single-cell data of the type that we have used to evaluate our model. For example, experimental data for selective spatial attention has been modelled by Corchs and Deco (2002) and Grossberg and Raizada (2000), and data concerning object-based selective attention has been modelled by Usher and Niebur (1996) and Hamker (2002). However, we believe that our model advances work in this area in several ways. Firstly, we have demonstrated that a single model can account for physiological data associated with both featural and spatial attention, as well as the complex interaction between attention and stimulus contrast. Despite the relative simplicity of our model it simulates this data in detail. Secondly, our model employs novel mechanisms to implement the competition between nodes and to model the modulatory effects of feedback pathways. Both these mechanisms are biologically plausible (Spratling and Johnson, 2001, 2002; Spratling, 2002), and the latter mechanism provides an account for the anatomically observed asymmetry between ascending and descending inter-regional cortical connections. Finally, we have demonstrated that the same model can also account for other top-down processes in visual perception. The proposed model thus suggests that each of these phenomena, which are currently considered to be distinct, result from common mechanisms. Hence, our model extends the applicability of the biased competition theory by suggesting that this principle is involved in a wide range of different perceptual phenomena. Competition between neural representations is a common feature of cortical information processing (Keysers and Perrett, 2002; O’Reilly, 1998). Furthermore, modulatory effects on neural activity have been proposed as a common computational mechanism employed throughout the cortex (Phillips and Singer, 1997; Salinas and Thier, 2000; Salinas and Sejnowski, 2001). Our model combines these two mechanisms to suggest that top-down information originating from a wide range of different sources can modulate neural activity and bias competition.

Many other neural network models of visual attention have been proposed which do not employ the principles of biased competition. For example, in many models attentional selection and object representation occur in

Figure 14: The effect of familiarity on parsing ambiguous patterns in the feedback model. The stimulus presented to the network consisted of all pixels in the lower-left quadrant being active together with the remainder of one of the two familiar patterns. The responses of eight nodes in the lower region, were recorded. The recorded nodes were selective to the horizontal or vertical bars within the lower-left quadrant, as indicated. (a) and (c) Before training, all the nodes representing vertical bars are active, irrespective of the contextual information. This is due to the nodes representing the vertical bars having slightly stronger feedforward weights in the particular network tested. (b) After training, when the familiar pattern is consistent with there being a horizontal bar in the lower-left quadrant, all the nodes representing horizontal bars are active. (d) After training, when the familiar pattern is consistent with there being a vertical bar in the lower-left quadrant, all the nodes representing vertical bars are active. Both response and time are measured in arbitrary units, but the same scale has been used in each plot.

separate, interacting, systems (e.g., Koch and Ullman, 1985; Itti and Koch, 2001; Heinke et al., 2002; Heinke and Humphreys, 2003; Mozer, 1988). These models typically employ a “saliency map” or “selection network” in which competition occurs between different spatial locations. This module is responsible for selecting the focus of attention and is used to gate the perceptual input to a separate neural system which performs object recognition. Rather than postulating a separate cortical system for the selection and control of the focus of attention, the biased competition model suggests that “attention is an emergent property of many neural mechanisms working together to resolve competition for visual processing” (Desimone and Duncan, 1995, p.194). Since competition can occur at each stage in the visual processing hierarchy, this model resolves the dichotomy between early (Broadbent, 1958; Treisman, 1969) and late (Deutsch and Deutsch, 1963; Shiffrin and Schneider, 1977) selection theories of attention by challenging the traditional view that there are distinct preattentive and attentive stages in perceptual processing. Rather than proposing that attention always operates to select objects at an early stage of visual processing, or that attentional selection operates at a late stage in visual processing, the biased competition model suggests that selection occurs at different stages depending on the stimuli and the task.

In many of the experiments we have reported, top-down information is provided by sources that are external to the simulated neural network. The biased competition model proposes that these top-down signals originate in cortical regions that respond distinctively in each of the attentional task conditions. Physiological data suggest that such signals might originate in a number of frontal and parietal regions of the cortex (Kanwisher and Wojciulik, 2000; Kastner and Ungerleider, 2000). These regions thus seem to be responsible for the selection and control of attention in these tasks. A very similar set of cortical regions are implicated in the provision working memory (Kastner and Ungerleider, 2000; Luck et al., 1997). The sustained activity of cells in these regions provides working memory and generates top-down signals that can affect the responses of cells throughout the ventral hierarchy (Supér et al., 2001; Duncan, 1998). These effects on neural activity, recorded during working memory tasks, are very similar to those recorded during attention tasks in the absence of visual stimulation (Kastner and Ungerleider, 2000). Hence, the top-down signals from working memory are likely to bias processing of subsequently presented stimuli (Desimone, 1998; Reynolds and Desimone, 1999; de Fockert et al., 2001). It therefore seems likely that the feedback signals in the attentional tasks we have simulated originate in working memory since all these tasks required a memory to be maintained of the cued spatial location or the cued target object.

While top-down bias results from working memory in certain experiments, it results from long-term memory in the experiments on familiarity and context. In these tasks, top-down bias was generated by the selectivity of the modelled neurons rather than being provided by an external source. The reciprocal connectivity of the modelled neural network then results in the modulation of neural response properties. Training causes a node to become selective to a certain stimulus. This node subsequently generates a stronger response to this stimulus which can bias processing of perceptual data via feedback. Such bias can modulate the responses of nodes in the previous region. It could also be transmitted via feedforward connections, to increase the strength (or saliency) of the bottom-up signal corresponding to the familiar item in subsequent processing stages. In this way, the processing of the perceptual data is biased by the selectivity of each node. Since the selectivities of cortical cells, at all stages of the visual processing hierarchy, can be modified by past experience (Kobatake et al., 1998; Karni, 1996; Logothetis, 1998; Sigala and Logothetis, 2001; Desimone, 1996) sensory-driven activity is biased to be interpreted in terms of stored knowledge about previously encountered situations (Lee et al., 2002; Siegel et al., 2000).

In conclusion, feedback information can originate from a variety of sources and be generated at a variety of times under different task conditions. Such top-down information has thus been given a variety of names, such as attention, expectation, context and familiarity. We suggest that an identical feedback mechanism is responsible for all these top-down effects on perceptual processing. We have demonstrated this claim by presenting a single neural network model that can account for a number of perceptual processes which are currently considered to be distinct. The model proposes that cued attention to a target location, or object, results in top-down activity that operates via cortical feedback projections to modulate sensory-driven neural activations and affect the on-going competition between cells (Duncan, 1998; Desimone and Duncan, 1995; Reynolds et al., 1999). Similarly, our model suggests that during figure/ground segmentation, feedback may enhance the activities of neurons representing spatially distributed image features, providing these elements with a competitive advantage over neurons representing other image features (Reynolds and Desimone, 1999). Furthermore, with this model feedback may also act to resolve ambiguities in sensory data by providing bias for one possible interpretation over all others and hence serve as a mechanism to solve the ‘binding’ problem (Treisman, 1998; Reynolds and Desimone, 1999; Roelfsema et al., 2000; Luck and Ford, 1998). A final phenomenon which can be accounted for by this model is contextual cueing, the process by which high-level knowledge generates top-down biases for perceptual processing (Chun, 2002; Olson et al., 2001).

4 Methods

4.1 The Reynolds and Desimone Model

The total excitatory and inhibitory activation received by the output node is calculated as:

$$y_{12,excit} = \sum_{i=1}^2 w_{i12}^+ y_{i1} x_{i1}$$

$$y_{12,inhib} = \sum_{i=1}^2 w_{i12}^- y_{i1} x_{i1}$$

Where y_{i1} is the activation of node i in the input region, w_{i12}^+ is the strength of the excitatory weight from input i , w_{i12}^- is the strength of the inhibitory weight from input i , and x_{i1} is the strength of the attention-dependent signal directed at input i . Reynolds et al. (1999) suggest that x_{i1} should take a value of 5 for attended stimuli, and 1 for unattended stimuli.

The time-varying output of the node can be calculated as:

$$y_{12}^t = y_{12}^{t-1} + \gamma \{ (\beta - y_{12}^{t-1}) y_{12,excit} - (\alpha + y_{12,inhib}) y_{12}^{t-1} \}$$

Where α is a parameter controlling the rate of decay of a node's activity, β is the maximum response of a node, and γ is a parameters that updates the activity slowly so as to dampen oscillatory behaviour. Reynolds et al. (1999) specify parameter values of $\alpha = 0.2$ and $\beta = 1$ and we use a value of $\gamma = 0.1$ in our simulations. If the input activity remains constant for a sufficient time, the output reaches a steady-state value given by:

$$y_{12}^t = \frac{\beta y_{12,excit}}{y_{12,excit} + y_{12,inhib} + \alpha}$$

4.2 The Feedback Model

For each node the activations of the apical and basal dendrites were calculated as:

$$y_{jk,apical}^t = \sum_{i=1}^{m_a} v_{ijk} x_{ijk}^{t-1} \quad (1)$$

$$y_{jk,basal}^t = \sum_{i=1}^{m_b} w_{ijk} X_{ijk}^{t-1} \quad (2)$$

Where $y_{jk,apical}^t$ is the activation of the apical dendrite of node j in region k at time t , $y_{jk,basal}^t$ is the activation of the basal dendrite of the same node, m_a is the total number of synapses on the apical dendrite, m_b is the total number of synapses on the basal dendrite, v_{ijk} is the synaptic weight from input i to the apical dendrite of node j in region k , w_{ijk} is the synaptic weight from input i to the basal dendrite of node j in region k , x_{ijk}^t is the activation of input i to the apical dendrite of node j in region k , and X_{ijk}^t is the activation received by the basal dendrite of node j in region k from input i after pre-integration lateral inhibition³:

$$X_{ijk}^t = x_{ijk}^t \left(1 - \alpha^t \max_{\substack{p=1 \\ (p \neq j)}}^n \left\{ \frac{w_{ipk}}{\max_{q=1}^m \{w_{qp k}\}} \frac{y_{pk}^{t-1}}{\max_{q=1}^n \{y_{qk}^{t-1}\}} \right\} \right)^+ \quad (3)$$

Where α^t is a scale factor controlling the strength of lateral inhibition, y_{pk}^{t-1} is the activation of node p in region k at time $t - 1$ (defined below in equation 4), and $(z)^+$ is the positive half-rectified value of z . In the reported simulations, the value of α^t was gradually increased at each time-step, from an initial value of zero to a maximum value of six in steps of 0.1.

The values of the apical inputs (x_{ijk}^t) are the activations of nodes in a higher cortical region at the previous time step (*i.e.*, $y_{j'k+1}^{t-1}$), or are top-down signals that arise from neural generators outside the modelled circuits. The values of the basal inputs (X_{ijk}^t) are the inhibited activations from nodes in a lower cortical region at the

³For full details of the implementation of pre-integration lateral inhibition, and for a justification of this scheme on both biological and computational grounds, see Spratling and Johnson (2001, 2002).

previous time step (*i.e.*, $y_{j'k-1}^{t-1}$). Note that the activations of every node are determined in an identical way and that attention-dependent top-down signals are treated in the same way as any other source of feedback: they provide another activation value, x , that contributes to the apical activation via a synaptic weight, v .

The activation of the apical dendrite multiplicatively modulates the activation of the basal dendrite in order to determine the output activation of each node:

$$y_{jk}^t = y_{jk,basal}^t (1 + y_{jk,apical}^t) \quad (4)$$

Where y_{jk}^t is the activation of node j in region k at time t . This formulation enables bottom-up, sensory-driven, stimulation to drive the response of the node even in the absence of top-down activity. In contrast, feedback activation can not drive the node's activity in the absence of feedforward activation, but it can amplify any response to feedforward stimulation.

The presence of reciprocal excitatory connections can lead to positive feedback effects resulting in run-away activation values. To prevent this, the activity of each node is attenuated in proportion to the cumulative strength its previous activity (C_{jk}^{t-1}), and this attenuated activation is clipped to be in the range [0,1]:

$$y_{jk}^t = \left[\frac{y_{jk}^t}{1 + C_{jk}^{t-1}} \right]_0^1 \quad (5)$$

This attenuated activity has a time-varying profile which resembles that of biological neurons; having an initial burst of activity followed by a sustained response at a lower firing rate. To make this profile more realistic the change in activity of the node is smoothed by taking into account the previous activity of the node:

$$y_{jk}^t = \tau_1 y_{jk}^t + (1 - \tau_1) y_{jk}^{t-1} \quad (6)$$

The cumulative activity, C_{jk}^t , of the node is calculated as:

$$C_{jk}^t = \tau_2 y_{jk}^t + (1 - \tau_2) C_{jk}^{t-1} \quad (7)$$

Where τ_1 and τ_2 are time constants, which take the values $\tau_1 = 0.5$ and $\tau_2 = 0.25$.

Finally, the activity of each node was modified by a small amount of noise, such that:

$$y_{jk}^t = y_{jk}^t (1 + \rho) \quad (8)$$

The noise values, ρ , were logarithmically distributed positive real numbers in the range [0,0.01]. Since the magnitude of the noise is small it has very little effect on neural activity except when multiple nodes have virtually identical synaptic weights. When this occurs, the noise causes one of these nodes to win the competition to be active in response to the current stimulus.

4.3 Synaptic Weight Values

In the simulations of the attentional data reported in sections 2.1 and 2.2, appropriate values for the synaptic weights were set by hand. Since the model was being used to simulate neurophysiological data collected from different cells with distinct selectivities, different sets of weight values were used in different simulations. Weights were chosen to provide the best fit between the behaviour of each model and the experimental data, with the following restrictions. Weight values were limited to the range [0,1], and were adjusted in steps of 0.1. The total sum of the synaptic weights received at each basal dendrite was made equal to one (*i.e.*, $\sum_{i=1}^m (w_{ijk})^+ = 1$). Where a network was symmetrical, corresponding weights were given the same value (*e.g.*, the weight from input one to node one was made equal to the weight from input two to node two). Feedback weights were given a value of half the corresponding feedforward weight. Given these constraints, there were 11 possible sets weights for the interconnections between the input and output regions, and 11 possible sets of weights for the external sources of feedback. Rather than exhaustively searching through these possibilities, weights were adjusted incrementally to improve the fit with the experimental data until satisfactory results were obtained.

Acknowledgements

This work was funded by MRC Research Fellowship number G81/512.

Reprint requests should be sent to Dr M. W. Spratling, Centre for Brain and Cognitive Development, Birkbeck College, 32 Torrington Square, London, WC1E 7JL, UK, or via email to m.spratling@bbk.ac.uk.

References

- Barbas, H. and Rempel-Clower, N. (1997). Cortical structure predicts the pattern of corticocortical connections. *Cerebral Cortex*, 7:635–46.
- Broadbent, D. E. (1958). *Perception and communication*. Pergamon, London, UK.
- Budd, J. M. L. (1998). Extrastriate feedback to primary visual cortex in primates: a quantitative analysis of connectivity. *Proceedings of the Royal Society of London. Series B*, 265(1400):1037–44.
- Cauler, L. J. (1995). Layer I of primary sensory neocortex: where top-down converges upon bottom-up. *Behavioural Brain Research*, 71(1-2):163–70.
- Cauler, L. J., Clancy, B., and Connors, B. W. (1998). Backward cortical projections to primary somatosensory cortex in rats extend long horizontal axons in layer I. *Journal of Comparative Neurology*, 390(2):297–310.
- Chelazzi, L., Miller, E. K., Duncan, J., and Desimone, R. (2001). Responses of neurons in macaque area V4 during memory-guided visual search. *Cerebral Cortex*, 11(8):761–72.
- Chun, M. M. (2002). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, 4(5):170–8.
- Corchs, S. and Deco, G. (2002). Large-scale neural model for visual attention: integration of experimental single cell and fMRI data. *Cerebral Cortex*, 12(4):339–48.
- Crick, F. and Asanuma, C. (1986). Certain aspects of the anatomy and physiology of the cerebral cortex. In Rumelhart, D. E., McClelland, J. L., and The PDP Research Group, editors, *Parallel Distributed Processing: Explorations in the Microstructures of Cognition. Volume 2: Psychological and Biological Models*, pages 333–71. MIT Press, Cambridge, MA.
- Crick, F. and Koch, C. (1998). Constraints on cortical and thalamic projections: the no-strong-loops hypothesis. *Nature*, 391:245–50.
- de Fockert, J. W., Rees, G., Frith, C. D., and Lavie, N. (2001). The role of working memory in visual selective attention. *Science*, 291(5509):1803–6.
- De Weerd, P., Peralta, M. R., Desimone, R., and Ungerleider, L. G. (1999). Loss of attentional stimulus selection after extrastriate cortical lesions in macaques. *Nature Neuroscience*, 2(8):753–8.
- Deco, G., Pollatos, O., and Zihl, J. (2002). The time course of selective visual attention: theory and experiments. *Vision Research*, 42(27):2925–45.
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Science USA*, 93:13494–9.
- Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of London. Series B*, 353(1373):1245–55.
- Desimone, R. and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18:193–222.
- Deutsch, J. A. and Deutsch, D. (1963). Attention: some theoretical considerations. *Psychological Review*, 70:80–90.
- Duncan, J. (1998). Converging levels of analysis in the cognitive neuroscience of visual attention. *Philosophical Transactions of the Royal Society of London. Series B*, 353(1373):1307–17.
- Ebdon, M. (1996). *Towards a General Theory of Cerebral Neocortex*. PhD thesis, University of Sussex.
- Everling, S., Tinsley, C. J., Gaffan, D., and Duncan, J. (2002). Filtering of neural signals by focused attention in the monkey prefrontal cortex. *Nature Neuroscience*, 5(7):671–6.
- Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in primate cerebral cortex. *Cerebral Cortex*, 1:1–47.
- Friston, K. J. and Büchel, C. (2000). Attentional modulations of effective connectivity from V2 to V5/MT in humans. *Proceedings of the National Academy of Science USA*, 97(13):7591–6.
- Frith, C. (2001). A framework for studying the neural basis of attention. *Neuropsychologia*, 39(12):1367–71.
- Gilbert, C., Ito, M., Kapadia, M., and Westheimer, G. (2000). Interactions between attention, context and learning in primary visual cortex. *Vision Research*, 40(10-12):1217–26.
- Gilbert, C. D., Das, A., Ito, M., Kapadia, M., and Westheimer, G. (1996). Spatial integration and cortical dynamics. *Proceedings of the National Academy of Science USA*, 93(2):615–22.
- Grossberg, S. and Raizada, R. (2000). Contrast-sensitive perceptual grouping and object-based attention in the laminar circuits of primary visual cortex. *Vision Research*, 40(10-12):1413–32.
- Hamker, F. H. (1999). The role of feedback connections in task-driven visual search. In Heinke, D., Humphreys, G. W., and Olson, A., editors, *Connectionist Models in Cognitive Neuroscience: Proceedings of the 5th Neural Computation and Psychology Workshop (NCPW98)*, pages 252–61, London, UK. Springer.
- Hamker, F. H. (2002). How does the ventral pathway contribute to spatial attention and the planning of eye movements? In Würtz, R. P. and Lappe, M., editors, *Proceedings of the 4th Workshop on Dynamic Perception*, pages 83–8, St. Augustin, Germany. Infix Verlag.

- Heinke, D. and Humphreys, G. W. (2003). Attention, spatial representation and visual neglect: simulating emergent attention and spatial memory in the selective attention for identification model (SAIM). *Psychological Review*, 110(1):29–87.
- Heinke, D., Humphreys, G. W., and di Virgilio, G. (2002). Modeling visual search experiments: selective attention for identification model (SAIM). *Neurocomputing*, 44-46:817–22.
- Hupé, J. M., James, A. C., Payne, B. R., Lomber, S. G., Girard, P., and Bullier, J. (1998). Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature*, 394(6695):784–7.
- Itti, L. and Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, 2(3):194–202.
- Jagadeesh, B., Chelazzi, L., Mishkin, M., and Desimone, R. (2001). Learning increases stimulus salience in anterior inferior temporal cortex of the macaque. *Journal of Neurophysiology*, 86(1):290–303.
- Johnson, R. R. and Burkhalter, A. (1997). A polysynaptic feedback circuit in rat visual cortex. *Journal of Neuroscience*, 17(18):7129–40.
- Kanwisher, N. and Wojciulik, E. (2000). Visual attention: insights from brain imaging. *Nature Reviews Neuroscience*, 1(2):91–100.
- Karni, A. (1996). The acquisition of perceptual and motor skills: a memory system in the adult human cortex. *Cognitive Brain Research*, 5:39–48.
- Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., and Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, 22(4):751–61.
- Kastner, S. and Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, 23:315–41.
- Keysers, C. and Perrett, D. I. (2002). Visual masking and RSVP reveal neural competition. *Trends in Cognitive Sciences*, 6(3):120–5.
- Kirschfeld, K. and Kammer, T. (2000). Visual attention and metacontrast modify latency to perception in opposite directions. *Vision Research*, 40:1027–33.
- Kobatake, E., Wang, G., and Tanaka, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *Journal of Neurophysiology*, 80(1):324–30.
- Koch, C. and Segev, I. (2000). The role of single neurons in information processing. *Nature Neuroscience*, 3(supplement):1171–7.
- Koch, C. and Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, 4(4):219–27.
- Körding, K. P. and König, P. (2000). Learning with two sites of synaptic integration. *Network: Computation in Neural Systems*, 11(1):25–39.
- Körding, K. P. and König, P. (2001). Supervised and unsupervised learning with two sites of synaptic integration. *Journal of Computational Neuroscience*, 11(3):207–15.
- Lamme, V. A. F. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *Journal of Neuroscience*, 15(2):1605–15.
- Lamme, V. A. F. (2000). Neural mechanisms of visual awareness: a linking proposition. *Brain and Mind*, 1:385–406.
- Lamme, V. A. F. and Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11):571–9.
- Lamme, V. A. F. and Spekreijse, H. (2000). Modulations of primary visual cortex activity representing attentive and conscious scene perception. *Frontiers in Bioscience*, 5:232–43.
- Lamme, V. A. F., Supér, H., Landman, R., Roelfsema, P. R., and Spekreijse, H. (2000). The role of primary visual cortex (V1) in visual awareness. *Vision Research*, 40(10-12):1507–21.
- Lamme, V. A. F., Supér, H., and Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology*, 8(4):529–35.
- Larkum, M. E., Zhu, J. J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, 398(6725):338–41.
- Larkum, M. E., Zhu, J. J., and Sakmann, B. (2001). Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. *Journal of Physiology*, 533(2):447–66.
- Lee, T. S., Mumford, D., Romero, R., and Lamme, V. A. F. (1998). The role of primary visual cortex in higher level vision. *Vision Research*, 38(15-16):2429–54.
- Lee, T. S., Yang, C. F., Romero, R. D., and Mumford, D. (2002). Neural activity in early visual cortex reflects behavioral experience and higher-order perceptual saliency. *Nature Neuroscience*, 5(6):589–97.
- Li, Z. (1999). Visual segmentation by contextual influences via intracortical interactions in primary visual cortex. *Network: Computation in Neural Systems*, 10(2):187–212.

- Li, Z. (2000). Pre-attentive segmentation in the primary visual cortex. *Spatial Vision*, 13(1):25–50.
- Logothetis, N. (1998). Object vision and visual awareness. *Current Opinion in Neurobiology*, 8(4):536–44.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, 77:24–42.
- Luck, S. J. and Ford, M. A. (1998). On the role of selective attention in visual perception. *Proceedings of the National Academy of Science USA*, 95:825–30.
- McAdams, C. J. and Maunsell, J. H. R. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, 19(1):431–41.
- McAdams, C. J. and Maunsell, J. H. R. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *Journal of Neurophysiology*, 83(3):1751–5.
- Mehta, A. D., Ulbert, I., and Schroeder, C. E. (2000). Intermodal selective attention in monkeys. II: Physiological mechanisms of modulation. *Cerebral Cortex*, 10:359–70.
- Moran, J. and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229(4715):782–4.
- Mountcastle, V. B. (1998). *Perceptual Neuroscience: The Cerebral Cortex*. Harvard University Press, Cambridge, MA.
- Mozar, M. C. (1988). A connectionist model of selective attention in visual perception. In *Proceedings of the 10th Annual Conference of the Cognitive Science Society*, pages 195–201, Hillsdale, NJ. Erlbaum.
- Mumford, D. (1992). On the computational architecture of the neocortex II: the role of cortico-cortical loops. *Biological Cybernetics*, 66:241–51.
- Olshausen, B. A., Anderson, C. H., and Van Essen, D. C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *The Journal of Neuroscience*, 13(11):4700–19.
- Olson, C. R. (2001). Object-based vision and attention in primates. *Current Opinion in Neurobiology*, 11(2):171–9.
- Olson, I. R., Chun, M. M., and Allison, T. (2001). The contextual guidance of attention: human intracranial event-related potential evidence for feedback modulation in anatomically early temporally late stages of visual processing. *Brain*, 124(7):1417–25.
- O'Reilly, R. C. (1998). Six principles for biologically based computational models of cortical cognition. *Trends in Cognitive Sciences*, 2(11):455–62.
- Peterson, M. A., Harvey, E. M., and Weidenbacher, H. J. (1991). Shape recognition contributions to figure-ground reversal: which route counts? *Journal of Experimental Psychology, Human Perception and Performance*, 17(4):1075–89.
- Phaf, R. H., Van der Heijden, A. H. C., and Hudson, P. T. W. (1990). SLAM: a connectionist model for attention in visual selection tasks. *Cognitive Psychology*, 22(3):273–341.
- Phillips, W. A., Kay, J., and Smyth, D. (1995). The discovery of structure by multi-stream networks of local processors with contextual guidance. *Network: Computation in Neural Systems*, 6(2):225–46.
- Phillips, W. A. and Singer, W. (1997). In search of common foundations for cortical computation. *Behavioural and Brain Sciences*, 20(4):657–722.
- Pollen, D. A. (1999). On the neural correlates of visual perception. *Cerebral Cortex*, 9(1):4–19.
- Reynolds, J. H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19:1736–53.
- Reynolds, J. H. and Desimone, R. (1999). The role of neural mechanisms of attention in solving the binding problem. *Neuron*, 24(1):19–29.
- Reynolds, J. H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, 26:703–14.
- Rockland, K. S. (1998). Complex microstructures of sensory cortical connections. *Current Opinion in Neurobiology*, 8:545–51.
- Roelfsema, P. R., Lamme, V. A. F., and Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. *Nature*, 395:376–81.
- Roelfsema, P. R., Lamme, V. A. F., and Spekreijse, H. (2000). The implementation of visual routines. *Vision Research*, 40:1385–411.
- Rolls, E. T. and Deco, G. (2002). *Computational Neuroscience of Vision*. Oxford University Press, Oxford, UK.
- Rolls, R. T. and Treves, A. (1998). *Neural Networks and Brain Function*. Oxford University Press, Oxford, UK.
- Salinas, E. and Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Science USA*, 93:11956–61.
- Salinas, E. and Abbott, L. F. (1997). Invariant visual perception from attentional gain fields. *Journal of Neurophysiology*, 77(6):3267–72.

- Salinas, E. and Sejnowski, T. J. (2001). Gain modulation in the central nervous system: where behavior, neurophysiology and computation meet. *The Neuroscientist*, 7(5):430–40.
- Salinas, E. and Thier, P. (2000). Gain modulation: a major computational principle of the central nervous system. *Neuron*, 27:15–21.
- Schroeder, C. E., Mehta, A. D., and Foxe, J. J. (2001). Determinants of attentional control in cortical processing: evidence from human and monkey electrophysiologic investigations. *Frontiers in Bioscience*, 6:d672–84.
- Shiffrin, R. M. and Schneider, W. (1977). Controlled and automatic human information processing: II perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84(2):127–90.
- Siegel, M., Kording, K. P., and König, P. (2000). Integrating top-down and bottom-up sensory processing by somato-dendritic interactions. *Journal of Computational Neuroscience*, 8:161–73.
- Sigala, N. and Logothetis, N. K. (2001). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, 415:318–20.
- Somers, D. C., Todorov, E. V., Siapas, A. G., Toth, L. J., Kim, D. S., and Sur, M. (1998). A local circuit approach to understanding integration of long-range inputs in primary visual cortex. *Cerebral Cortex*, 8(3).
- Spratling, M. W. (2002). Cortical region interactions and the functional role of apical dendrites. *Behavioral and Cognitive Neuroscience Reviews*, 1(3):219–28.
- Spratling, M. W. and Johnson, M. H. (2001). Dendritic inhibition enhances neural coding properties. *Cerebral Cortex*, 11(12):1144–9.
- Spratling, M. W. and Johnson, M. H. (2002). Pre-integration lateral inhibition enhances unsupervised learning. *Neural Computation*, 14(9):2157–79.
- Stemmler, M., Usher, M., and Niebur, E. (1995). Lateral interactions in primary visual cortex: a model bridging physiology and psychophysics. *Science*, 269:1877–80.
- Stuart, G., Spruston, N., Sakmann, B., and Häusser, M. (1997). Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends in Neurosciences*, 20(3):125–31.
- Supér, H., Spekreijse, H., and Lamme, V. A. F. (2001). A neural correlate of working memory in the monkey primary visual cortex. *Science*, 293(5527):120–4.
- Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London. Series B*, 353(1373):1295–1306.
- Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review*, 76:282–99.
- Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, 24(5):295–300.
- Treue, S. and Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399(6736):575–9.
- Usher, M. and Niebur, E. (1996). Modeling the temporal dynamics of IT neurons in visual search: a mechanism for top-down selective attention. *Journal of Cognitive Neuroscience*, 8(4):311–27.
- Vecera, S. P. (2000). Toward a biased competition account of object-based segregation and attention. *Brain and Mind*, 1:353–84.
- Yuste, R., Gutnick, M. J., Saar, D., Delaney, K. R., and Tank, D. W. (1994). Ca^{2+} accumulations in dendrites of neocortical pyramidal neurons: an apical band and evidence for two functional compartments. *Neuron*, 13(1):23–43.
- Zipser, K., Lamme, V. A. F., and Schiller, P. H. (1996). Contextual modulation in primary visual cortex. *Journal of Neuroscience*, 16(22):7376–89.