

From Knowing What to Knowing Where: Modeling Object-Based Attention with Feedback Disinhibition of Activation

Frank van der Velde and Marc de Kamps

Abstract

■ We propose a neural model of visual object-based attention in which the identity of an object is used to select its location in an array of objects. The model is based on neural activity observed in visual search tasks performed by monkeys. In the model, the identity of the object (target) is selected in the higher areas of the ventral stream by means of a cue. Feedback activation from these higher areas carries information about the identity of the target to the (lower) retinotopic areas of the ventral stream. In these areas, the feedback activation interacts

with feedforward activation produced by the object array. The interaction occurs in local microcircuits, and results in a selective activation on locations in the retinotopic areas of the visual stream that correspond to the location of the target in the object array. The selective activation consists of a form of gain control, produced by disinhibition. Transmitted to the dorsal stream, this activation directs spatial attention to the location of the target. In this way, an action directed at the target can be generated. ■

INTRODUCTION

Active control of behavior often requires selecting an object among other objects as the target for an action. In visual perception, the process of selecting information from the total amount of visual input received at a given time is known as visual selective attention. In this way, an object can be selected as the target for an action by focusing attention on the object, ignoring the other objects in its vicinity.

When an object is selected on the basis of information about its location in the visual field, the selection process is referred to as spatial or location-based attention. In contrast, when an object is selected on the basis of information about its identity (i.e., its shape or color), the selection process is referred to as object-based attention (or visual search). Attention as a whole will likely result from competitive mechanisms within the cortex (e.g., Desimone & Duncan, 1995), in which both spatial attention and object-based attention are important.

Here, we will concentrate on neural mechanisms for object-based attention. In particular, we will present and discuss a neural model that shows how the location of an object can be selected when the identity of that object (target) is known. After all, when an object is selected on the basis of its identity, information about its location is needed as well if the object is to be the target for an action. The model is based on neural activity

observed in monkey studies of object-based attention (e.g., Gottlieb, Kusunoki, & Goldberg, 1998; Motter, 1994a, 1994b; Chelazzi, Miller, Duncan, & Desimone, 1993).

The selection of an object on the basis of its identity was studied by Chelazzi et al. (1993). In the task they used, illustrated in Figure 1a, a cue was presented at the center of gaze and a monkey was trained to hold this object in memory during a delay period. After the delay, an array of two to five choice objects was presented in the periphery. The monkey was trained to make an eye movement (saccade) to the object matching the cue (the target), ignoring the other objects (the distractors) in the array.

During the experiment, the activity of single neurons in the anterior inferotemporal cortex (AIT) was recorded. This area belongs to the ventral stream (or pathway), which is involved in object recognition. Cells in AIT are selectively responsive to specific shapes over a large proportion of the visual field (e.g., Farah, Humphreys, & Rodman, 1999; Oram & Perrett, 1994; Ungerleider & Haxby, 1994). The stimuli used in the experiment were complex objects, selected in such a way that each recorded AIT cell responded strongly to only one of them and weakly to the others. In particular, the same cells that responded to the cue also responded to the matching target, irrespective of the difference in location between the cue and the target.

The cells responsive to the cue maintained activity during the delay period, in line with other 'delayed




Figure 1. Schematic presentation of two visual search tasks. (a) Visual search task recorded in the anterior inferotemporal cortex. First, a cue is presented in the center of gaze (top). Next, in a delay period the cue is removed but the monkey continues to focus on a fixation point presented at the location of the cue (middle). Finally, an array of objects is presented in the periphery (bottom), and the monkey is asked to make a saccade to the object in the array that matches the memorized cue. (b) Visual search task recorded in the lateral intraparietal cortex. First, an array of objects is presented (top). Next, a fixation point is presented in the center of the array (middle), and the monkey is asked to focus on the fixation point. Finally, one of the objects in the array is presented as a cue near the fixation point (bottom). The cue dictates a saccade to the object in the array that matches the cue.

response' studies (e.g., Miyashita, 1988; Yakovlev, Fusi, Berman, & Zohary, 1998). But after presentation of the object array, all cells responsive to the objects in the object array (target and distractors alike) initially showed an increase in activity. This suggests that information about the objects in the object array was processed simultaneously, and not sequentially. However, 90–120 msec before the onset of the saccade a dramatic change in responsiveness had occurred. The cells responsive to the target showed a sustained activity, but the activity of the cells responsive to the distractors declined, even though these objects were still within the receptive field (RF). This

pattern of activity suggests a selection of the target on the basis of its identity in AIT (Chelazzi et al., 1993).

However, although the sustained target activity in AIT indicates the selection of the target on the basis of its identity, information about the location of the target is lost in this area (e.g., Farah et al., 1999; Ito, Tamura, Fujita, & Tanaka, 1996; Oram & Perrett, 1994). In the experiment of Chelazzi et al. (1993), the activity of most target responsive cells in AIT was indeed not affected by the (relative) location of the target in the array, and the cells that responded to the cue also responded to the matching target even though the target was presented at a different location. Yet in the task studied by Chelazzi et al., an eye movement has to be made to the location of the target in the array, which entails that information about the location of the target in the object array is needed. This suggests an involvement of the dorsal stream (or pathway) in this task, because the dorsal stream processes information about the location of objects (Colby & Olson, 1999; Steinmetz & Constantinidis, 1995).

The role of the dorsal stream in object-based attention was studied by Gottlieb et al. (1998), Kusunoki, Colby, Duhamel, and Goldberg (1997) and Goldberg (1996). In their task, illustrated in Figure 1b, a circular array of objects was presented. After a monkey made a saccade to the fixation point in the middle of the array, a cue stimulus was presented near the fixation point. The monkey was trained to make a saccade towards the target object that matched the cue, ignoring the non-matching objects (the distractors). Thus, the target for the eye movement was again selected by the identity of the cue. During the experiments, single neuron activity in the lateral intraparietal cortex (LIP) was recorded. LIP is part of the dorsal stream and it plays a key role in the control of spatial attention related to saccadic eye movements (e.g., Colby, Duhamel, & Goldberg, 1996).

In particular, LIP neurons were studied that had at least one of the objects in the array in their RF. The cue, however, was never in the RF of the recorded LIP neurons. The response of the LIP neurons depended on the location of the target that matched the cue. The neurons did not respond when the cue identified a target outside their RF. However, when the identified target was in their RF, the neurons started to respond some time after the presentation of the cue, but before the saccade. Hence, the presaccadic activity of the tested LIP neurons reflects the direction of spatial attention to the location of the target in the array. Because the target was selected on the basis of its identity match with the cue, this suggests that “the ventral visual stream is telling the dorsal stream where to go” (Goldberg, 1996, p. 91).

The fact that the ventral and the dorsal stream work together in a task as illustrated in Figure 1 is discussed by Goodale and Humphrey (1998). They describe the interaction between both streams in terms of the meta-

phor of teleassistance. This is a process whereby a human operator instructs a robot from a distance to perform a goal-directed action. In this metaphor, the ventral stream (the human in the metaphor) identifies an object, and ‘flags’ the object in the scene. Once the object has been flagged, the dorsal stream (the robot) can direct an action towards the object.

We will show how the ventral stream can ‘flag’ an object so that the information about the location of the target is available for the dorsal stream. In our model, ‘flagging’ the target results from an interaction between feedforward and feedback activation in the ventral stream, in a manner discussed below.

THE MODEL

An outline of the model is illustrated in Figure 2. Spatial information is processed in the dorsal stream, and object

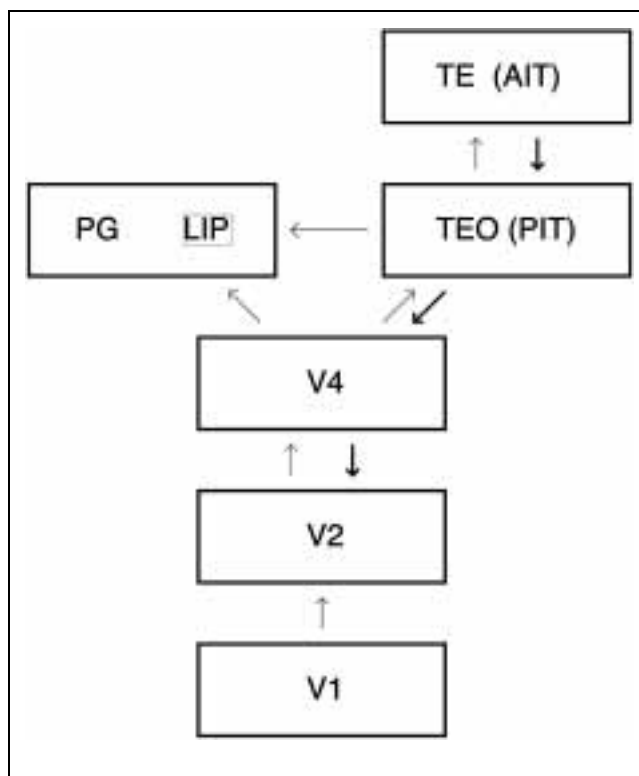


Figure 2. Outline of the model which simulates the tasks illustrated in Figure 1. Information about the location of the objects in the object array is processed in the dorsal (or ‘where’) stream, which runs from V1 to the parietal cortex (area PG). This stream includes the LIP, which is involved in the spatial control of eye movements. Information about the identity of the objects in the object array is processed in a feedforward manner in the ventral (or ‘what’) stream. This stream runs from V1 to TE or AIT, passing through V2, V4, and TEO or posterior inferotemporal cortex (PIT). Feedback activation in the ventral stream carries information about the identity of a selected object (target). In the retinotopic areas of the ventral stream (V2–PIT), an interaction occurs between feedforward activation and feedback activation. On the basis of this interaction, information about the location of the target is retrieved in the ventral stream. Transmitted to the dorsal stream, this information directs spatial attention to the location of the target in the object array.

identity is processed in the ventral stream. The dorsal stream runs from V1 through V2 and V4 to the parietal cortex (area PG), which includes LIP. The ventral stream starts at V1 and goes via V2–V4 and posterior inferotemporal cortex (PIT) to AIT. RF size of the neurons in the ventral stream increases steadily from V1 to PIT, and information is represented in a retinotopic fashion in these areas. Loss of retinotopy occurs in the transition to AIT. The neurons in this layer are selectively responsive to the identity of objects irrespective of their location. As a result, a target can be selected in this layer on the basis of its identity.

However, before an eye movement to the target can be made, information about its location in the object array has to be available. In our model, information about the location of the identified target is retrieved by means of the interaction between a feedforward network and a feedback network. The feedforward network simulates the process of object identification in the ventral stream. The feedback network simulates the communication of top-down information about object identity in the ventral stream. The feedforward network and the feedback network interact within the retinotopic areas of the ventral stream.

Thus, when a target is selected in AIT, as observed by Chelazzi et al. (1993), the feedback network carries information about the identity of the target to the retinotopic areas of the ventral stream. In these areas, the feedback network interacts with the feedforward network that processes information about the objects in the object array (Figure 1). As a result of this interaction, information about the location of the target in the object array is retrieved.

The model illustrated in Figure 2 is based on the nature of the information that is available in the ventral stream. In the higher areas in the ventral stream, information about the location of an object is lost. Hence, from that level onwards, the ventral stream can only communicate information about the identity of the target, but not information about its location. However, information about the location of the target is still available in the lower (retinotopic) areas of the ventral stream. Thus, if the ventral stream can integrate this information with information about the identity of the target, it can ‘flag’ the target in the scene of objects. The interaction between feedback and feedforward activation illustrated in Figure 2 is a direct way to do this.

The model illustrated in Figure 2 gives a prediction about the timing of the ‘flagging’ process. In the metaphor of Goodale and Humphrey, the ventral stream ‘flags’ (thus identifies) the object as a target, but “once a goal object has been selected for goal-directed action, the two systems [i.e., ventral and dorsal stream] process incoming visual information simultaneously” (Goodale & Humphrey, 1998, p. 203). But if the ventral stream has already identified (‘flagged’) the object, it would seem

that there is no need to process that information again simultaneously with the dorsal stream. On the other hand, if the ventral stream and dorsal stream do process the incoming information simultaneously, it is unclear how the ventral stream can instruct the dorsal stream to generate a goal-directed action.

Figure 2 shows that in the model presented here, the ventral and dorsal stream do indeed process the incoming visual information simultaneously, but initially this does not result in selecting ('flagging') the target. The selection of the target occurs when feedback activation starts to interact with feedforward activation in the ventral stream. Because the selection process depends on feedback activation, the model predicts that the effects of object-based attention and the selection of the target will occur in the higher areas of the ventral stream first, and after that in the lower areas. Recently, Mehta, Ulbert, and Schroeder (2000a) have measured the response latencies due to attention in the ventral stream. They found that attention effects do indeed occur in the higher visual areas first, followed later by attention effects in the lower visual areas.

We have implemented the model on two levels. The first level concerns the correlation in activation between the feedforward network and the feedback network. Because the feedforward network carries information about the object display (target and distractors) and the feedback network carries information about the identity of the target, one can expect a (positive) correlation in activation between both networks on the locations in the retinotopic areas where information about the target is processed. In our model, this correlation in activation is the basis for retrieving information about the location of the target. We investigate the conditions under which this correlation in activation between the feedforward network and the feedback network can occur.

The second level concerns the local interaction between the feedforward network and the feedback network. We model this interaction in terms of a local microcircuit. On locations in the retinotopic areas where there is a strong correlation in activation between the feedforward network and the feedback network, the interaction between both networks results in an enhanced activation in the local microcircuit. This enhanced activation represents the retrieval of information about the location of the target. As indicated in Figure 2, it can be transmitted to the dorsal stream to direct spatial attention to the location of the target in LIP, after which the required eye movement can be made (e.g., Dominey & Arbib, 1992).

The pattern of activation simulated with our local microcircuit is based on observations reported by Motter (1994a, 1994b). He presented a circular array of objects (oriented bars) around a fixation spot, similar to the array in Figure 2. The color of the fixation spot was used as a cue. The color of half of the objects in the array matched the color of the cue. Single neuron

activity was recorded in area V4 during a discrimination task based on the color of the cue. The initial activity of V4 cells was the same for cells responding to the matching color and for cells responding to the non-matching color. But 150–200 msec after stimulus onset the cells that responded to the color of the cue showed enhanced activity. In contrast, the activity of cells responding to the nonmatching color was attenuated. In this way, "the topographic representation of the neural activity in area V4 highlights the potential targets in the visual scene at the expense of background objects" (Motter, 1994a, p. 2178).

A dramatic effect was observed when the color of the cue switched in the middle of a trial from one alternative to the other. Following the cue switch, a reversal of activity occurred during the trial. The activity of the initially nonmatching cells increased to the level of the initially matching cells, whereas the activity of the initially matching cells declined to the level of the initially nonmatching cells (Motter, 1994b). We simulated this reversal of activity with our local microcircuit.

SIMULATION OF THE MODEL

Feedforward Identification of Objects

Although the ventral stream processes information about the identity of objects (e.g., shape and color) in a manner not yet fully understood (e.g., Desimone, Miller, & Chelazzi, 1994), a few characteristics of this process have been established.

First, information about object identity is processed rapidly in the ventral stream, which suggests that this information is processed by means of a feedforward network (e.g., Koch & Poggio, 1999; Oram & Perrett, 1994), at least in the case of familiar objects.

Second, the ventral stream consists of a hierarchy of areas (Figure 2), and the RF size of neurons increases with each area by a factor of approximately 2–2.5 (e.g., Kobatake & Tanaka, 1994; Oram & Perrett, 1994). The areas from V1 to PIT are organized in a retinotopic fashion, but neurons in AIT respond to the identity of objects irrespective of their location (e.g., Farah et al., 1999; Ito et al., 1996; Oram & Perrett, 1994).

Third, neurons in each area selectively respond to particular object features, with increasing complexity for neurons in higher areas. However, neurons can respond to more than one object feature, which is evidence of a distributed representation (e.g., Tanaka, 1996; Kobatake & Tanaka, 1994; Oram & Perrett, 1994). For instance, some neurons in V4 observed by Motter (1994a) showed a maximum response for the combination of a particular color and orientation ('primary response'), but they also responded to a different color with the same orientation ('secondary response').

In line with this description, we used a feedforward network in our model to simulate the identification of

the object array presented in Figure 1. The feedforward network consists of a hierarchical sequence of layers, which represent the areas V1, V2, V4, PIT, and AIT of the ventral stream (Figure 2). V1 is the input layer of the network. Each neuron in this layer represents one of four different line orientations (vertical, horizontal, left diagonal, and right diagonal). From V1 to AIT, the RF size of the neurons in the network increases in a ratio as found in the visual cortex, with loss of retinotopic information in AIT. The specific details of the feedforward network are presented in the Appendix.

The network was trained to identify each object in the object array (Figure 1) on each location in the array, using the ‘back-propagation’ training procedure (e.g., Rumelhart, Hinton, & Williams, 1986). As a result, the feedforward network developed distributed representations of the objects in its hidden layers. Figure 3a shows the activity in the layers V2–AIT of the network when one of the objects in the object array (Figure 1) is identified by the network. The object produces location dependent activity in the layers V2–PIT and activates its identification neuron in AIT. Figure 3b shows the activity in the layers V2–AIT when the whole object array in Figure 1 is identified by the network. Each object in the

array produces location-dependent activity in the layers V2–PIT, and activates its identification neuron in AIT.

Feedback Communication of Identity

The feedback communication of identity information in our model occurs by means of a feedback network. The feedback network has the same structure as the layers V2–AIT of the feedforward network (see the Appendix), but its connections are reciprocal. Hence, activation in the feedback network flows from AIT to V2 (Figure 2). The feedback network can be conceived of as lying ‘on top’ of the feedforward network (with the exception of V1).

The feedback network was trained on the basis of the activations in the feedforward network, using the ‘Hebbian learning rule’ (e.g., Brunel, 1996). In the training procedure, an object was presented to the feedforward network for identification, and the resulting activations in the feedforward network (Figure 3a) were used to determine the connections in the feedback network. This procedure was used for all objects on all locations in the object array. As a result, the feedback network developed distributed representations of the objects in its hidden layers as well.

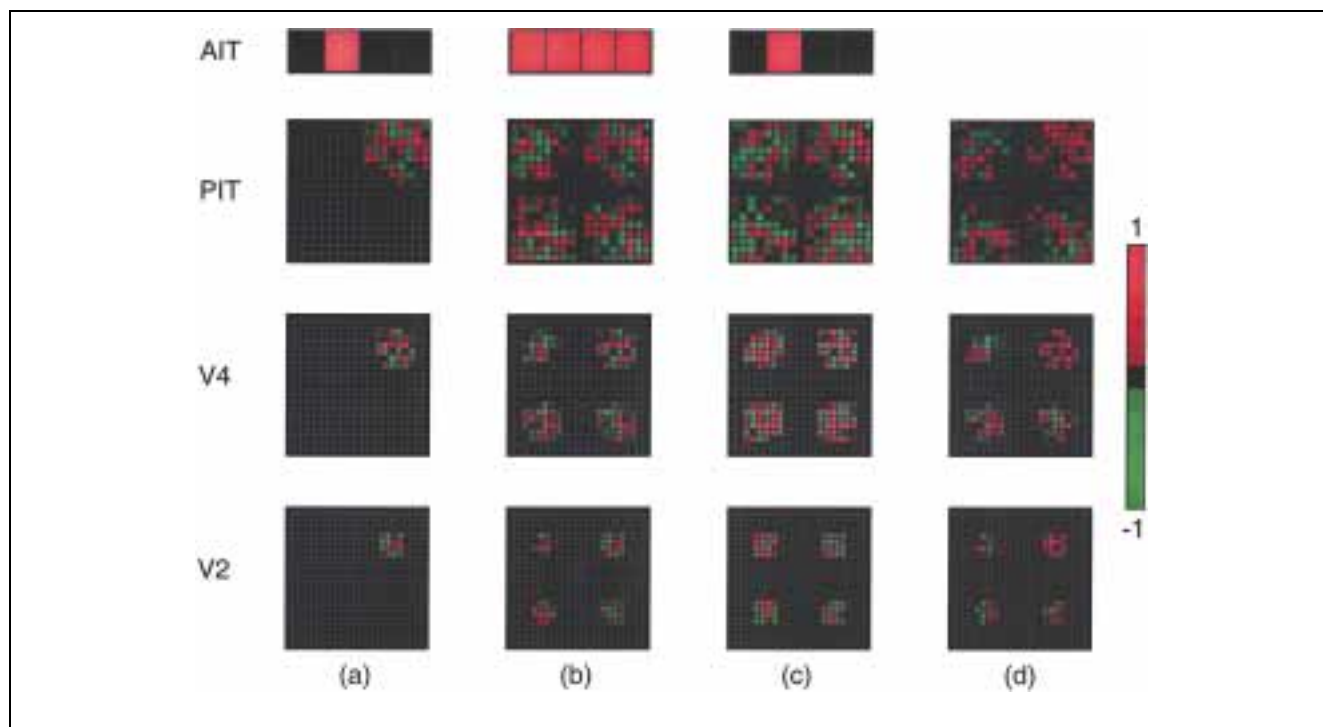


Figure 3. Activity and correlation in activity during visual search. (a) Activity in the feedforward network when the network identifies the square in the object array presented in Figure 1. The object is presented to the network by activating neurons in the vertical and horizontal orientation maps of V1 (not shown). The object produces location-dependent activity in the areas V2–PIT, and it activates the neuron that represents its identity in AIT. (b) Activity in the feedforward network when the four-object array in Figure 1 is identified by the network. The objects in the array produce location-dependent activations in the retinotopic areas V2–PIT. In AIT, each object in the array activates the neuron that represents the identity of the object in the feedforward network. (c) Activity in the feedback network when one object in the object array (the square) is selected as the target for visual search. The target is selected by activating its AIT neuron in the feedback network. This neuron initiates the activation of the target-related neurons in the lower (retinotopic) areas PIT–V2 of the feedback network. Due to the divergent nature of feedback, target-related neurons are activated on all potential object locations in the array. (d) The match (red) and nonmatch (green) in activity between corresponding neurons in the feedforward network (b) and the feedback network (c), given by the product of their activity.

Target selection on the basis of identity in AIT can occur by means of lateral inhibition (Usher & Niebur, 1996). We assume that such a process activates the target-related AIT neuron in the feedback network. Because activation in the feedback network flows from AIT to V2, the target-related AIT neuron initiates the activation of the target-related neurons on all locations in the retinotopic areas (PIT-V2) of the feedback network. This is illustrated in Figure 3c. The fact that neurons on all locations in these areas are activated results from the manner in which the feedback network was trained. It also reflects the divergent nature of feedback activation.

Relation Between Feedforward and Feedback Activation

To identify an object, a feedforward network has to produce activations in its hidden layers that are selective enough to activate the correct identification neuron in AIT. Thus, when different objects (e.g., the target and the distractors in Figure 1) are presented on the same location to the feedforward network, they produce different patterns of activation within the feedforward network during the process of identification. This is true for any feedforward network, irrespective of the manner in which it is trained.

If these selective activations are then used to train the feedback network (using Hebbian learning), one can expect a match in activation between the target-related activity in the feedforward network and the target-related activity in the feedback network. In contrast, the activity produced by the distractors in the feedforward network will not be correlated with the target-related activity in the feedback network.

The relation between the activity in the feedforward network and the activity in the feedback network is illustrated in Figure 3d. This figure presents the individual match in activity (Appendix) between corresponding neurons in the retinotopic areas of the feedforward network (Figure 3b) and the feedback network (Figure 3c). On the locations in the retinotopic areas that correspond to the location of the target in the array, there is a match in activity between (almost) all corresponding neurons in the two networks. In contrast, on the locations in the retinotopic areas that correspond to the locations of the distractors in the array there is a mixture of match and nonmatch in activity between corresponding neurons in the two networks. This mixture results from the distributed representation in both networks.

Figure 3d shows that target-related activity in V2-PIT is distinguished from distractor-related activity in V2-PIT by a high correlation in activity between the feedforward network and the feedback network. This suggests that information about the location of the target can be retrieved in the retinotopic areas of the ventral stream

by means of a local interaction between the feedforward network and the feedback network.

We model this interaction by means of local microcircuits (e.g., Amit & Brunel, 1997; Douglas, Martin, & Witteridge, 1989), which connect the feedforward network with the feedback network on each location in the retinotopic areas of the ventral stream.

A Microcircuit for Local Interaction

Each microcircuit consists of four interconnected populations of neurons. Three of these are presented in Figure 4: two excitatory populations (A and B) and one inhibitory population (I). The behavior of the neurons in the populations in the microcircuit are modeled in terms of average neuron dynamics (see the Appendix).

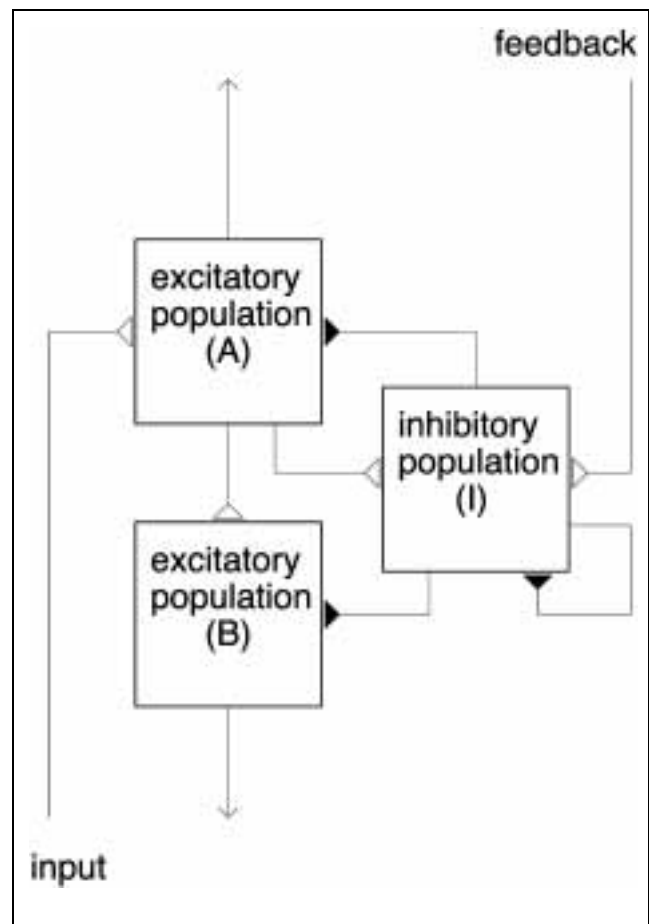


Figure 4. A local microcircuit consisting of interacting populations of neurons. Populations A and B consist of interconnected excitatory neurons, and Population I consists of interconnected inhibitory neurons. Connections between the populations are excitatory (open triangle) or inhibitory (black triangle). The inhibitory connection from I to I represents the existence of disinhibition circuits in Population I. A stimulus produces feedforward activation of stimulus selective neurons in Population A. A cue produces feedback activation of cue selective neurons in Population I. This results in disinhibition of stimulus selective neurons in Population A when the cue matches the stimulus, and inhibition of stimulus selective neurons in Population A when the cue does not match the stimulus.

The neurons in Population A belong to the feedforward network. They receive feedforward input from the lower layers and they transmit their output to the higher layers. The neurons in Population A have overlapping RFs and they are selectively responsive to features of objects (target and/or distractors). The neurons in Population A also receive inhibitory input from neurons in Population I. Population B connects the ventral stream with the dorsal stream. Population B receives excitatory input from Population A and inhibitory input from Population I. The activity in Population B directs spatial attention to the location of an object in the dorsal stream, as observed in LIP (e.g., Gottlieb et al., 1998). The inhibitory neurons in Population I receive their input from Population A in the circuit and from the local population of (excitatory) neurons in the feedback network (the fourth population in the microcircuit, not shown in Figure 4). Population I contains a subset of neurons that are connected with other inhibitory neurons in Population I. This can produce inhibition of inhibition (or disinhibition) in the microcircuit.

In Motter's experiments (Motter, 1994a, 1994b), the activity of nonmatching neurons (not responsive to the cue) was suppressed and the activity of matching neurons (responsive to the cue) was enhanced. We model both effects in terms of a feedback modulation of inhibition in the microcircuit. The suppression of the activity of nonmatching neurons results from the feedback activation of the inhibitory neurons in Population I that are connected to the nonmatching neurons in Population A. The enhancement of the activity of matching neurons results from the feedback activation of the disinhibition circuits in Population I that are connected to the matching neurons in Population A. Thus, the

selectivity of feedback activation in this model consists of a selective activation of inhibitory neurons, resulting in either inhibition (for nonmatching neurons) or disinhibition (for matching neurons).


Disinhibition occurs when inhibitory neurons are inhibited by other inhibitory neurons. The target cells of the first inhibitory neurons then receive less inhibition. Circuits of this kind have been found in the visual cortex (Kisvárdy, 1992; Kisvárdy, Beaulieu, & Eysel, 1993). A specific disinhibition circuit in the primary visual cortex has recently been described by Gonchar and Burkhalter (1999). This circuit operates primarily on the extragranular (output) layers of the cortex, and it receives feedback activation from higher areas.

The effect of disinhibition on the target cells depends on the input that the target cells receive from other sources. If the target cells receive activation from other sources, disinhibition will enhance that activation by reducing the effect of inhibition on the target cells. But if the target cells do not receive activation from other sources, disinhibition will not have a direct effect on the activation of the target cells. In this way, disinhibition provides a form of gain control of the activation of the target cells. Thus, feedback activation of the disinhibition circuits connected to matching neurons in Population A will enhance the activity of these neurons only if these neurons receive feedforward activation.

Behavior of the Microcircuit

We will first investigate the effect of feedback disinhibition, and compare it with the effect of direct feedback activation. The activity of the matching neurons in Population A could be enhanced by direct (cue-related)

Figure 5. Time course of activity in Population A of the local microcircuit (Figure 4) during cue switching. The activity of Population A is high when the cue matches the stimulus on that location, but declines when the cue switches from match to nonmatch (M→NM). Conversely, the activity of Population A is relatively low when the cue does not match the stimulus on that location, but increases when the cue switches from non-match to match (NM→M). The dashed line represents a 10-msec running average, and the solid line represents a 50-msec running average.



feedback activation. However, due to the divergent nature of feedback (Figure 3c), all matching neurons on all locations in the retinotopic areas will be activated in this way, even if there is no stimulus on that location. But in Motter's experiments, the activity of matching neurons was not significantly affected by the presence of the cue when the stimulus was absent (Motter, 1994a). In contrast, the effect of the cue was substantial when the stimulus was present (Motter, 1994a, 1994b). This suggests that the effect of the cue does not result from a direct feedback activation, but instead results from an interaction between feedback activation caused by the cue and feedforward activation caused by the stimulus. Such an interaction is produced by feedback disinhibition. With disinhibition, the cue will not have an effect on locations where a stimulus is absent.

Figure 5 shows the effect of overall disinhibition on the activity of Population A (see the Appendix) when the cue is turned on or off (that is, when cue switches from nonmatch to match or from match to nonmatch, as in Motter, 1994b). When the cue is turned on, the activity of Population A increases as a result of the feedback disinhibition of the microcircuit. When the cue is turned off, feedback disinhibition of the microcircuit stops and the activity of Population A declines to the nonmatch level. This pattern of results is in agreement with Motter's observations in the cue switching task (Motter, 1994b, Figure 7).

The enhancement of activity in Population A that results when the cue is turned on represents the selection of the location of the target. Figure 6a illustrates how this information can be transmitted to the dorsal

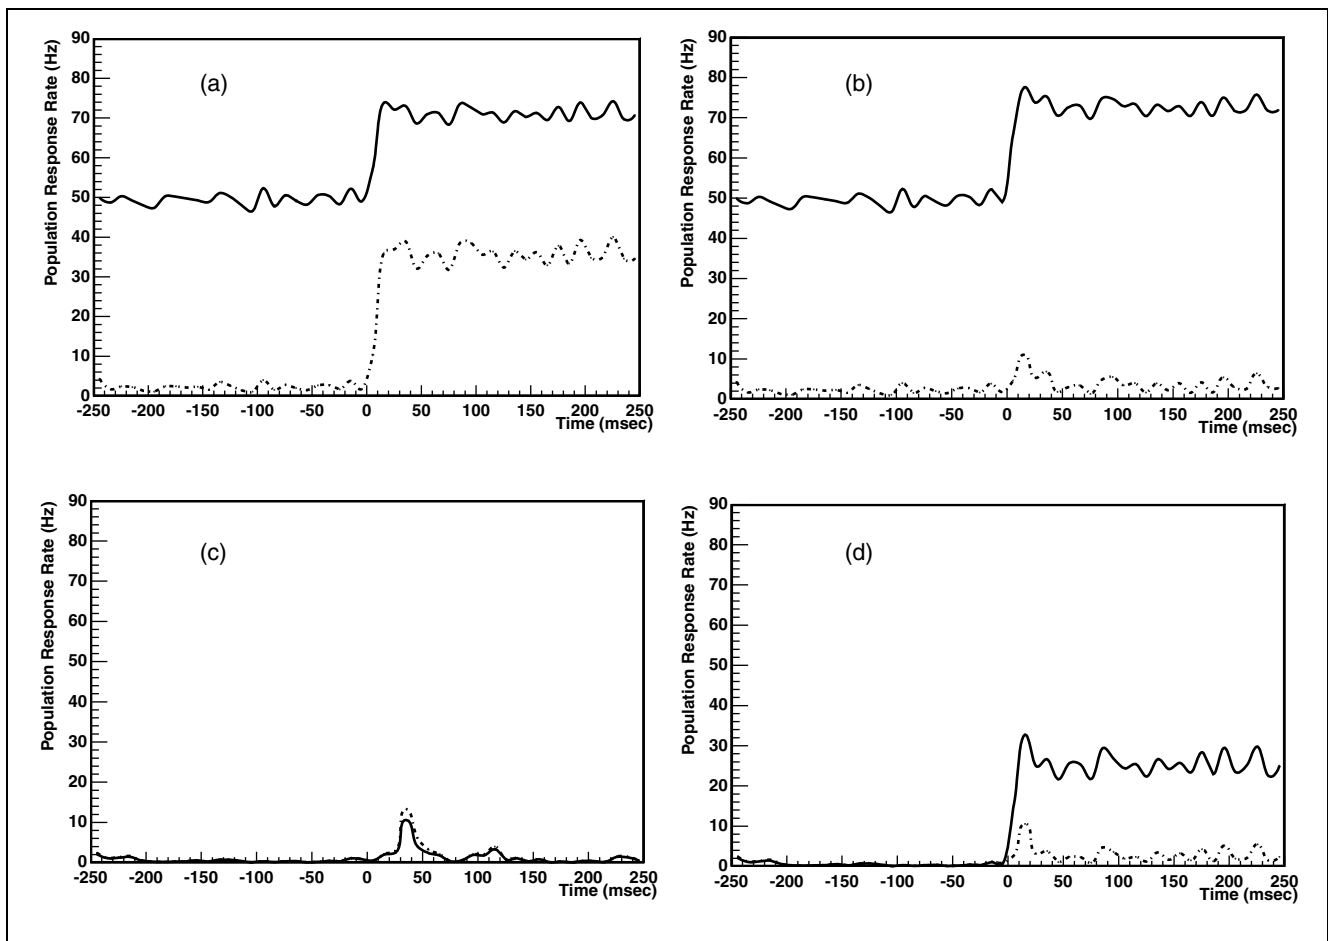


Figure 6. Interaction between feedforward activation and feedback activation in the local microcircuit (Figure 4). The activity of Populations A (solid line) and B (dashed line) is presented with a 10-msec running average. (a) Activity when the cue matches the stimulus on the location of the microcircuit, and feedback activation produces disinhibition in the circuit, starting at 0 msec. Feedback disinhibition increases the activity in Population A from nonmatch to match (as in Figure 5), and initiates activity in Population B. The increased activity in Population A represents the selection of the location of the target. The initiated activity in Population B can be transmitted to the dorsal stream, to direct spatial attention to the location of the target, as in LIP. (b) Activity when the cue matches the stimulus on the location of the microcircuit, and feedback activation results in direct activation of Population A, starting at 0 msec. Direct feedback activation increases the activity in Population A, but does not result in a significant activation of Population B, due to the increased inhibition in the circuit. (c) Activity when the stimulus is absent and the cue produces disinhibition in the circuit, starting at 0 msec. Feedback disinhibition has no significant effect on the activity in both populations when the stimulus is absent. (d) Activity when the stimulus is absent and the cue produces direct feedback activation of Population A, starting at 0 msec. Feedback activation produces significant activity in Population A, even though the stimulus is absent on the location of the microcircuit.

stream. This figure compares the activity of the excitatory Populations A and B in the microcircuit. When the cue is off, the activity of Population A is on the non-match level (Figure 5) and the activity of Population B is on a resting level. When the cue is turned on, the resulting disinhibition increases the activity of Population A to the match level (Figure 5) and it initiates the activation of Population B. The activation of Population B is caused by the combined effect of the increased activity of Population A and the reduced effect of inhibition on Population B.

To compare disinhibition with direct feedback activation, Figure 6b illustrates the effect on Populations A and B when the presence of the cue results in a direct feedback activation of Population A (see the Appendix). The effect of direct feedback activation on the activity of Population A is similar to that of disinhibition (Figure 6a). However, the activity of Population B is strongly reduced when the presence of the cue results in direct feedback activation instead of disinhibition. This results from the fact that a direct feedback activation of Population A also increases the activity of the inhibitory neurons, which reduces the effect on Population B caused by the enhanced activity of Population A.

Figure 6c and d illustrate the effect of disinhibition and direct feedback activation on Populations A and B when the cue is on but the stimulus is absent. Disinhibition (Figure 6c) has no significant effect on the activity of both populations in this case. But the direct feedback activation of Population A initiates a substantial activation of Population A (Figure 6d). A substantial activity of cue-related neurons when the stimulus is absent is in contrast with the observations by Motter (1994a).

Selecting the Location of the Target

In our model, Population B in the microcircuit connects the ventral stream with the dorsal stream. As illustrated in Figure 6a, Population B is activated when feedback results in an overall disinhibition in the microcircuit, which occurs when there is an overall match between feedforward and feedback activation. As illustrated in Figure 3d, an overall match in activation between the feedforward and the feedback network is found on locations in the retinotopic areas that correspond to the location of the target in the array. Thus, the activation of Population B in the microcircuit can be used to select the location of the target in the array.

Figure 7 shows the activation of Population B in the microcircuit, before and after the onset of the cue, on each location in each of the areas in the ventral stream. On each location, the activation of Population B in the microcircuit is determined by the match and nonmatch in activity between the feedforward and the feedback network, as illustrated in Figure 3d. In




Figure 7. Activation of Population B in the microcircuit (Figure 4) on each location in each of the areas in the ventral stream (Figure 3d), before and after onset of the cue (at 0 msec).

particular, disinhibition in the microcircuit will result from matching neurons, whereas inhibition in the microcircuit will result from nonmatching neurons (see the Appendix).

Figure 7 illustrates the selection of the location of the target by means of the activation of Population B in the microcircuit. There is a clear difference in activation of Population B between the location of the target and the locations of the distractors in each of the retinotopic areas of the ventral stream. The onset of the cue does not result in a substantial activation of Population B on the locations of the distractors in each area. As illustrated in Figure 3d, this results from the fact that disinhibition and inhibition nearly cancel on these locations. In contrast, a substantial activation of Population B occurs after onset of the cue on the location of the target in each area. As illustrated in Figure 3d, this activation results from the overall disinhibition on the location of the target.

Thus, the activation of Population B in the microcircuit selects the location of the target in the array. In turn, this activation can be used to activate location-selective neurons in the dorsal stream, as observed in LIP (e.g., Gottlieb et al., 1998). The first area in the ventral stream where the location of the target can be selected is PIT (Figure 7). It is suggested by Webster, Bachevalier, and Ungerleider (1994) that the connection between PIT (TEO) and LIP (Figure 2) is involved in the visual search task illustrated in Figure 1. A second area where the location of the target can be selected is V4 (Figure 7). This area is involved in the generation of saccades (Weber & Fisher, 1990; Fischer & Boch, 1981). The location of the target could also be selected in V2. As illustrated in Figure 2, this area could transmit location information to the dorsal stream as well.

DISCUSSION

The simulations of the model presented here show that information about the location of an identified object (target) in an object array can be retrieved by an interaction in the retinotopic areas of the ventral stream between a feedback network that carries information about the identity of the target and a feedforward network that carries information about the object array (stimulus). In this way, the ventral stream can select an object so that the dorsal stream can generate a goal-directed action (e.g., see Goodale & Humphrey, 1998).

The interaction between the feedback network and the feedforward network in our model is based on the fact that the target-related activity in the feedback network is highly correlated with the target-related activity in the feedforward network. In contrast, the target-related activity in the feedback network is not correlated with the distractor-related activity in the feedforward network. We showed that this difference in correlation results when the connections in the feedback network are developed (using Hebbian learning) on the basis of the activations that occur in the feedforward network during object identification.

The difference in correlation between target-related activity and distractor-related activity can be used to select information about the location of the target in the object array. We presented a model of a local microcircuit in which this selection occurs. The local microcircuits connect the feedforward and the feedback network on each location in the retinotopic areas of the ventral stream. When activity in both networks is highly correlated, a feedback modulation of activity, based on disinhibition, occurs in the microcircuit. This modulation of activity enhances target-related activity in the microcircuit and activates a gating circuit that transmits activation to the dorsal stream. In the dorsal stream, the transmitted activation directs spatial attention to the

location of the target. In this way, an eye movement directed at the target can be generated.

However, the microcircuit we presented is not only a gating circuit, but it also models the data obtained by Motter (1994a, 1994b). In this way, our model integrates two seemingly different processes: activation modulation due to attention in the ventral stream and the selection of the location of an identified target. The microcircuit we presented makes a specific prediction about the nature of the gating and attention process: gating and activation modulation result from disinhibition, which affects in particular the output layers in the cortex. As illustrated in Figure 6b, feedback activation instead of feedback disinhibition produces activation modulation, but does not produce gating of activation, which shows that the combination of activation modulation due to attention and gating is not trivial. Recently, Mehta, Ulbert, and Schroeder (2000b) showed that attentional modulation of activation in the ventral pathway is concentrated in the extragranular laminae (the output laminae of the cortex), and occurs in a manner consistent with disinhibition.

The importance of the presence of the stimulus in the interaction between the feedforward network and the feedback network can be illustrated with a variation of the task presented in Figure 1b. In this variation (Gottlieb et al., 1998), LIP neurons were tested with blocks of trials in which the target was always present in the RF of the neuron ('target' condition), interleaved with blocks of trials in which the same target was removed from the object display ('no-target' condition). In the no-target condition, the monkeys were rewarded to make the same saccade as in the target condition, this time to an empty space. The removal of the target greatly reduced the response of the observed LIP neurons (Gottlieb et al., 1998). This result is in line with the fact that the interaction (based on disinhibition) between the feedforward network and the feedback network has no effect when the stimulus is absent (Figure 6c).

The fact that the ventral and dorsal streams process different attributes of a stimulus (identity vs. location, or 'what' vs. 'where') has raised the questions of where and how this information is integrated (e.g., Rao, Rainer, & Miller, 1997; Caminiti, 1996; Jeannerod, 1996). A likely area of integration is the prefrontal cortex, because both streams project to this area, although to separate regions (e.g., Rao et al., 1997; Caminiti, 1996). The integration of information from the ventral and dorsal streams in the prefrontal cortex was studied by observing the delay activity of neurons in this area in a variation of the task presented in Figure 1a (Rao et al., 1997). In this variation, a second delay period was inserted after the presentation of the object array (Figure 1a, bottom), and the monkey was rewarded to make a saccade to the memorized location of the target in the array. As expected, some neurons showed either identity-related or location-related delay activity. How-

ever, neurons were also found that showed delay activity related to a combination of identity and location. In particular, a preferred location cued by a preferred object elicited more activity in these neurons than a preferred location cued by a nonpreferred object. Least activity resulted when a nonpreferred location was cued. This result suggests that information from the ventral and dorsal streams are indeed integrated in the prefrontal cortex.

How identity and location can be linked in the prefrontal cortex can be explained in terms of the model simulated here. Because object identity and object location are processed separately in the ventral and dorsal streams, the link between the identity of the cue (target) and the location of the target in the array is ambiguous when the information in both streams is conveyed in a feedforward manner to the prefrontal cortex. However, information about the location of the target can be retrieved in the ventral stream by means of feedback modulation, as simulated here. This information can be used to direct spatial attention to the location of the target in the dorsal stream, which resolves the ambiguity of the location information conveyed by the dorsal stream to the prefrontal cortex. As a result, the information about the location of the target can be linked there with the information about the identity of the cue, which was already available in the prefrontal cortex.

The ability to select the location of an identified target is not only important for making eye movements, as in the tasks illustrated in Figure 1, but is also important for other forms of action, such as reaching. The question then arises how the ventral stream can influence the dorsal stream for these other forms of action. The answer to this question depends on the way in which spatial information is represented for these other forms of action.

Batista, Buneo, Snyder, and Andersen (1999) and Andersen, Snyder, Batista, Buneo, and Cohen (1998) have studied spatial representation in the parietal reach region (PRR), which is involved in the execution and planning of limb movements (reaching), and compared it with spatial representation in LIP. They showed that spatial information in both areas is first represented in eye-centered coordinates. This is even true in the case of auditory-triggered responses in an eye movement task, although the information in that case is initially received in head-centered coordinates. Hence, it seems that all information concerning an action is first represented in eye-centered coordinates.

In the model presented here, the ventral stream selects the location of a target in terms of its retinotopic representation. In turn, the selected retinotopic representation can influence the eye-centered representation of the target in both LIP and PRR. Eye-centered representation can be transformed into other forms of spatial representation (e.g., Pouget & Sejnowski, 1997). By means of these transformations, the selection of a target in eye-centered coordinates will result in the selection of

the target in other spatial coordinates as well. In this way, the ventral stream can influence the early representation stages of all forms of action. This would be much more difficult if the ventral stream could not influence the dorsal stream in terms of eye-centered coordinates. In that case, it seems, the ventral stream would have to be able to influence action in each of the other forms of spatial representation, which puts a heavy burden on the capacity for spatial representation in the ventral stream.

A similar argument can be made for the planning of actions. For instance, information about the location of a target could first be selected in eye-centered coordinates and then be transformed into other forms of spatial representation, as described above. The information could then be retained in these forms of spatial representation, so that it could play a role in the planning of an action. Conversely, it could also be the case that a planned action can only be executed after it has first been transformed into eye-centered coordinates. Then, the ventral pathway could influence the planned action during the execution stage, for instance, by directing the action to the object at its present location (which is information that might not be available during the planning stage). Both forms of influence on planning are possible in terms of the model presented here.

The functional role of the feedback connections in the cortex is currently a main issue in neuroscience (e.g., Crick & Koch, 1998; Oram & Perrett, 1994). Perhaps something of the mystery (Oram & Perrett, 1994) of feedback connections would be resolved if feedback connections turn out to be an important mechanism for attentional control of information analyzed in different streams and areas (e.g., Duncan, Humphreys, & Ward, 1997). In visual processing this would be true both for attention directed by spatial cues (e.g., Müller, Teder-Sälejärvi, & Hillyard, 1998), and for attention directed by identity cues, as demonstrated by the simulations presented here.

APPENDIX

Networks and Correlations

The input layer for the feedforward network is V1. The whole layer consists of a 24×24 matrix in which each element represents a V1 RF. For each RF in V1 there are four input neurons that represent four line orientations (vertical, horizontal, left diagonal, and right diagonal), which give 2304 neurons in V1. An object is presented to the network by activating these neurons in V1.

The layers V2–AIT and the connections between the layers in the feedforward network and the feedback network are organized in terms of the RFs in V1. Each neuron in V2 has a RF that covers a 2×2 submatrix of RFs in V1. Each 2×2 submatrix in V1 (529 in all) is

covered by one V2 neuron, which gives 529 neurons in V2. A neuron in V2 is connected to all (16) V1 neurons in its RF. Each neuron in V4 has a RF that covers a 4×4 submatrix of RFs in V1. Each 4×4 submatrix in V1 (441 in all) is covered by one V4 neuron, which gives 441 neurons in V4. A neuron in V4 is connected to all (9) V2 neurons that have RFs that are fully covered by the RF of the V4 neuron. Each neuron in PIT has a RF that covers a 8×8 submatrix of RFs in V1. Each 8×8 submatrix in V1 (289 in all) is covered by one PIT neuron, which gives 289 neurons in PIT. A neuron in PIT is connected to all (25) V4 neurons that have RFs that are fully covered by the RF of the PIT neuron. Finally, each AIT neuron has a RF that covers the whole 24×24 matrix of RFs in V1. Hence, each AIT neuron is connected to all PIT neurons.

The activity of the neurons in the layers V2–AIT is given by the function $2/(1 + e^{-\text{net}}) - 1$, where net = (weighted sum of inputs) + bias. In contrast with the logistic function (e.g., Zipser & Andersen, 1988), this function gives zero activity for a zero input. In this way, the match in activity between corresponding neurons in the feedforward and feedback network can be given by the product of the activity of both neurons.

Microcircuit

The excitatory populations (A and B) and the inhibitory population (I) are modeled in terms of average neuron activity, representing the overall population activity (as described in Usher & Niebur, 1996). The activities are given by current equations, representing the input to a population, and a response function that transforms currents into rates.

The current equations for Populations A, B and I are:

$$\begin{aligned}\tau_s \frac{dI_A}{dt} &= -I_A + J_{A \text{ to } A} F(I_A) - J_{I \text{ to } A} F(I_I) + I_{\text{ff}} + I_{\text{bg}} \\ \tau_s \frac{dI_B}{dt} &= -I_B + J_{B \text{ to } B} F(I_B) - J_{I \text{ to } B} F(I_I) + J_{A \text{ to } B} F(I_A) \\ &\quad + I_{\text{bg}} \\ \tau_s \frac{dI_I}{dt} &= -I_I - J_{I \text{ to } I} F(I_I) + J_{A \text{ to } I} F(I_A) + (-m + nm) I_{\text{fb}} \\ &\quad + I_{\text{bg}}\end{aligned}$$

In these equations, I_i is the current in population i , and τ_s is the synaptic time constant (taken as 5 msec). The parameters J_i , representing the synaptic weights, are

$$\begin{aligned}J_{A \text{ to } A}, J_{A \text{ to } B}, J_{A \text{ to } I} &= 1.0 \\ J_{B \text{ to } B} &= 0.3 \\ J_{I \text{ to } I} &= 0.1 \\ J_{I \text{ to } A} &= 1.1 \\ J_{I \text{ to } B} &= 1.1.\end{aligned}$$

The sensory input (stimulus) is represented by a feedforward input current I_{ff} , injected in Population A.

A Gaussian background current I_{bg} is injected in all populations, with mean \bar{I}_{bg} and standard deviation σ . Feedback is represented by a current I_{fb} , divided in the fractions m (match) and nm (nonmatch). In the case of overall disinhibition (Figure 6a,c), $m = 1$ and $nm = 0$. In the case of direct feedback activation (Figure 6b,d), I_{fb} is not injected in Population I (as in the current equations above), but I_{fb} is injected in Population A. In Figure 7, m and nm are determined on the basis of Figure 3d, with $m + nm = 1$. The values (in unit of the current) for the currents are:

$$\begin{aligned}I_{\text{ff}} &= 0.05 \\ \bar{I}_{\text{bg}} &= 0.025 \\ \sigma &= 0.03 \\ I_{\text{fb}} &= 0.025.\end{aligned}$$

The function $F(I_i)$ represents the response function that transforms currents into discharge rates. Here, it is given by the noise-corrected response function $F(I, \sigma) = 1 / [T_r + T_{\text{sp}}(I, \sigma)]$. In this equation, T_r represents the absolute refractory period of a neuron (taken as 1 msec), and T_{sp} represents the mean interspike interval, as given by Ricciardi (1977):

$$T_{\text{sp}}(I, \sigma) = \tau \sqrt{\pi} \int_{x(I, \sigma)}^{y(I, \sigma)} \exp(z^2) [1 + \text{erf}(z)] dz.$$

Here, erf is the error function and the limits of integration are given by:

$$x(I, \sigma) = -\frac{I\tau}{\sigma\sqrt{\tau}}, y(I, \sigma) = \frac{1 - I\tau}{\sigma\sqrt{\tau}}.$$

Acknowledgments

We wish to thank two anonymous reviewers for their valuable comments on this article.

Reprint requests should be sent to: Frank van der Velde, Unit of Experimental and Theoretical Psychology, University of Leiden, Wassenaarseweg 52, 2333 AK Leiden, The Netherlands. E-mail: vdvelde@fsw.leidenuniv.nl.

REFERENCES

- Amit, D. J., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7, 237–252.
- Andersen, R. A., Snyder, L. H., Batista, A. P., Buneo, C. A., & Cohen Y. E. (1998). Posterior parietal areas specialized for eye movements (LIP) and reach (PRR) using a common coordinate frame. *Novartis Foundation Symposium*, 218, 109–128.
- Batista, A. P., Buneo, C. A., Snyder, L. H., & Andersen, R. A. (1999). Reach plans in eye-centered coordinates. *Science*, 285, 257–260.
- Brunel, N. (1996). Hebbian learning of context in recurrent neural networks. *Neural Computation*, 8, 1677–1710.
- Caminiti, R. (1996). Attentional and spatial mechanisms in the parietal cortex. In R. Caminiti, K.-P. Hoffmann, F. Lacquaniti, & J. Altman (Eds.), *Vision and movement mechanisms in the cerebral cortex* (pp. 89–97). Strasbourg: HFSP.

- Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, *363*, 345–347.
- Colby, C. L., & Olson, C. R. (1999). Spatial cognition. In M. J. Zigmond, F. E. Bloom, S. C. Landis, J. L. Roberts, & L. R. Squire (Eds.), *Fundamental neuroscience* (pp. 1363–1383). San Diego: Academic Press.
- Colby, C. L., Duhamel, J.-R., & Goldberg, M. E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology*, *76*, 2841–2852.
- Crick, F., & Koch, C. (1998). Feature article: Consciousness and neuroscience. *Cerebral Cortex*, *8*, 97–107.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.
- Desimone, R., Miller, E. K., & Chelazzi, L. (1994). The interaction of neural systems for attention and memory. In C. Koch & J. L. Davies (Eds.), *Large-scale neuronal theories of the brain* (pp. 75–91). Cambridge: MIT Press.
- Dominey, P. F., & Arbib, M. A. (1992). A cortico-subcortical model for generation of spatially accurate sequential saccades. *Cerebral Cortex*, *2*, 153–175.
- Douglas, R. J., Martin, K. A. C., & Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural Computation*, *1*, 480–488.
- Duncan, J., Humphreys, G., & Ward, R. (1997). Competitive brain activity in visual attention. *Current Opinion in Neurobiology*, *7*, 255–261.
- Farah, M., Humphreys, G. W., & Rodman, H. R. (1999). Object and face recognition. In M. J. Zigmond, F. E. Bloom, S. C. Landis, J. L. Roberts, & L. R. Squire (Eds.), *Fundamental neuroscience* (pp. 1339–1361). San Diego: Academic Press.
- Fischer, B., & Boch, R. (1981). Enhanced activation of neurons in prelunate cortex before visually guided saccades of trained rhesus monkeys. *Experimental Brain Research*, *44*, 129–137.
- Goldberg, M. E. (1996). Attentional and spatial mechanisms in the parietal cortex. In R. Caminiti, K.-P. Hoffmann, F. Lacquaniti, & J. Altman (Eds.), *Vision and movement mechanisms in the cerebral cortex* (pp. 89–97). Strasbourg: HFSP.
- Gonchar, Y., & Burkhalter, A. (1999). Connectivity of GABAergic calretinin-immunoreactive neurons in rat primary visual cortex. *Cerebral Cortex*, *9*, 683–696.
- Goodale, M. A., & Humphrey, G. K. (1998). The objects of action and perception. *Cognition*, *67*, 181–207.
- Gottlieb, J. P., Kusunoki, M., & Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, *391*, 481–484.
- Ito, M., Tamura, H., Fujita, I., & Tanaka, K. (1996). Size and position invariance of neuronal responses in monkey inferotemporal cortex. *Journal of Neurophysiology*, *73*, 218–226.
- Jeannerod, M. (1996). Cortical coding of visual object attributes during object-oriented behaviour. In R. Caminiti, K.-P. Hoffmann, F. Lacquaniti, & J. Altman (Eds.), *Vision and movement mechanisms in the cerebral cortex* (pp. 15–22). Strasbourg: HFSP.
- Kisvardáy, Z. F. (1992). GABAergic networks of basket cells in the visual cortex. In R. R. Mize, R. E. Marc, & A. M. Sillito (Eds.), *Progress in brain research*, *90*, 385–405.
- Kisvardáy, Z. F., Beaulieu, C., & Eysel, U. T. (1993). Network of GABAergic large basket cells in cat visual cortex (area 18): Implication for lateral disinhibition. *Journal of Comparative Neurology*, *327*, 398–415.
- Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *Journal of Neurophysiology*, *71*, 856–867.
- Koch, C., & Poggio, T. (1999). Predicting the visual world: Silence is golden. *Nature Neuroscience*, *2*, 9–10.
- Kusunoki, M., Colby, C. L., Duhamel, J.-R., & Goldberg, M. E. (1997). The role of the lateral intraparietal area in the control of visuospatial attention. In H. Sakata, A. Mikami, & J. M. Fuster (Eds.), *The association cortex: Structure and function* (pp. 191–205). Amsterdam: Harwood.
- Mehta, S. D., Ulbert, L., & Schroeder, C. E. (2000a). Intermodal selective attention in monkeys: I. Distribution and timing of effects across visual areas. *Cerebral Cortex*, *10*, 343–358.
- Mehta, S. D., Ulbert, L., & Schroeder, C. E. (2000b). Intermodal selective attention in monkeys: II. Physiological mechanisms of modulation. *Cerebral Cortex*, *10*, 359–370.
- Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, *335*, 817–820.
- Motter, B. C. (1994a). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of Neuroscience*, *14*, 2178–2189.
- Motter, B. C. (1994b). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *Journal of Neuroscience*, *14*, 2190–2199.
- Müller, M. M., Teder-Sälejärvi, W., & Hillyard, S. A. (1998). The time course of cortical facilitation during cued shifts of spatial attention. *Nature Neuroscience*, *1*, 631–634.
- Oram, M. W., & Perrett, D. I. (1994). Modeling visual recognition from neurobiological constraints. *Neural Networks*, *7*, 945–972.
- Pouget, A., & Sejnowski, T. J. (1997). Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, *9*, 222–237.
- Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science*, *276*, 821–824.
- Ricciardi, L. (1977). *Diffusion processes and related problems*. Berlin: Springer.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*, 533–536.
- Steinmetz, M. A., & Constantinidis, C. (1995). Neurophysiological evidence for a role of posterior parietal cortex in re-directing visual attention. *Cerebral Cortex*, *5*, 448–456.
- Tanaka, K. (1996). Representation of visual features of objects in the inferotemporal cortex. *Neural Networks*, *9*, 1459–1476.
- Ungerleider, L. G., & Haxby, J. V. (1994). ‘What’ and ‘where’ in the human brain. *Current Opinion in Neurobiology*, *4*, 157–165.
- Usher, M., & 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*, 311–327.
- Weber, H., & Fisher, B. (1990). Cortical mechanisms for vision revealed by saccadic eye movements in trained rhesus monkeys. In E. Iwai & M. Mishkin (Eds.), *Vision, memory and the temporal lobe* (pp. 323–330). New York: Elsevier.
- Webster, M. J., Bachevalier, J., & Ungerleider, L. S. (1994). Connections of inferior temporal areas TEO and TE with parietal and frontal cortex in macaque monkeys. *Cerebral Cortex*, *4*, 470–483.
- Yakovlev, V., Fusi, S., Berman, E., & Zohary, E. (1998). Inter-trial neuronal activity in inferior temporal cortex: A putative vehicle to generate long-term visual associations. *Nature Neuroscience*, *1*, 310–317.
- Zipser, D., & Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, *331*, 679–684.