

Attention, Habituation and Conditioning: Toward a Computational Model¹

Christian Balkenius

Lund University, Sweden²

Is attention a purely perceptual process or is it in any way related to motor control? The aim of this article is to show that attention puts similar demands on a cognitive system as motor control and present evidence supporting the view that similar mechanisms operate in the two processes. A computational model of attention is presented that uses habituation as well as classical and instrumental conditioning to explain a number of attentional processes. Evidence from neurophysiology is reviewed that suggest that attention is controlled in a way similar to actions. This view makes it possible to adapt traditional learning theoretical mechanisms to the control of attention. Computer simulations are presented that illustrates the operation of the model.

Keywords: attention, learning theory, brain systems.

Introduction

What is attention for? Is it a purely perceptual mechanism that prevents cognitive overload by keeping irrelevant sensory information away, or is it part of the machinery controlling actions? Allport (1990) reviews evidence against theories of attention that assume that sensory information must be selected at an early stage of processing. Instead he suggests that the function of selective attention is to specify the parameters for goal-directed actions. In grasping a particular object, it is necessary that the spatial coordinates of that object are selected, and not those of some other object in the scene (Allport, 1987, 1990). Many recent empirical studies support this view of attention (Tipper, Howard, Houghton, 1998, Castiello, 1999). It appears that one important role of attention is to open up sensory channels with information relevant to a particular action. Before an object can be reached or grasped, the focus of attention must first select information from the environment that should guide the action. This implies that attentional fixations must be coordinated with motor actions.

¹ This research was supported in part by The Swedish Council for Research in the Humanities and Social Sciences.

² Address for correspondence: Dr. Christian Balkenius, Lund University Cognitive Science, Kungshuset, Lundagård, S-222 22 LUND, Sweden. E-mail: christian.balkenius@lucs.lu.se (Fax: +46-46-222-97-58). Home page: <http://www.lucs.lu.se/People/Christian.Balkenius/>

To focus attention is a way to probe the environment for further information. In a series of experiments, Ballard, Hayhoe, Pook and Rao (1997) have shown that subjects appear to look in the external world even when they could in principle remember the external state. They report that even for the simple task of copying the pattern of four colored blocks, subjects made as many as eighteen fixations on the model even though the pattern can easily be memorized. This shows that humans prefer to use a deictic strategy where information is obtained from the environment even when it is slower than a memory-based strategy.

Apart from guiding actions, the focus of attention may also control the redirection of attention itself. For example, a visual cue can suggest where to look for more information. In many experiments, a cue, such as an arrow, is used to point in the direction of a target stimulus. Instead of guiding action, this fixation sets the parameters for an attentional shift.

The close connection between attention and action suggests that it would be useful if both could be controlled in similar ways. In fact, the control of attention has many similarities with motor control. This is apparent when overt attentional shifts in the form of eye-movement are considered, but attentional shifts can be seen as actions more generally although motor activity need not be present. Instead, the action consists of a covert attentional shift that only influences the way sensory information is processed. The evidence for the view of attention-as-action is mainly circumstantial and is suggested by the parts of the brain apparently involved in both attention and motor control. These parallels will be developed further below, but here I only want to stress that this view of attention opens up for an exchange of mechanisms between theories of motor control, classical learning theory and attention.

Below, traditional response generating mechanisms such as classical and instrumental will be put to use in the control of attention. These learning mechanisms have both an activating and an inhibiting component and this is also true of attention. While we often think of attention as primarily directed toward a stimulus, an equally important aspect of attention is that it directs attention away from irrelevant stimuli. In patients with parietal injury, one of the most striking effects is their inability to change the focus of attention away from a stimulus (Rafal & Robertsson, 1995). After unilateral damage to the parietal lobe, patients are severely impaired at moving their attention. Posner (1988) has explained this phenomenon in terms of a parietal attention system whose function is to disengage attention from the current target and allow a shift toward a new location. When this system is damaged, a part of the visual scene will maintain attention despite attempts to refocus.

The prefrontal cortex effectuates another type of inhibition. Patients with prefrontal injury, especially in the orbital regions, are easily distracted by external stimuli. Fuster (1997) describes this as a defect of the exclusionary aspect of attention, that is, an inability to suppress attention to inappropriate stimuli. Animals with lesions in the orbital regions are unable to suppress the orienting reaction and are generally hyperactive. "[...] many have interpreted the hyperactivity of frontal animals as the expression of an underlying tendency to overreact to external stimuli; in other words, the hyperactivity seems secondary and reducible to hyperreactivity.

The relevance of this interference to issues of attention is obvious...” (Fuster 1997, p. 71).

The above evidence suggests a view of attention as consisting of at least two levels. The first directs attention toward stimuli in the environment and the second level is responsible for the inhibitory control of the lower subsystem. Different types of learning are in principle possible at both levels. Below, I will describe how classical and instrumental conditioning can function as activating mechanisms for attention, while contextual control can play the inhibitory role.

Figure 1 illustrates the components of an attentionally controlled action. First attention to object B is inhibited. This will disengage attention to B and make a shift of attention possible. In the next step, an attentional shift s is performed that results in the change of gaze g from B to A. The focus of attention opens up sensory channels that are used as implicit argument for the motor action a that directs the hand to object A.




Figure 1: Four principles for attention. Attention to object B is inhibited which disengages the focus of attention (attention-as-inhibition). Shifting the attention and gaze g from object B to A is an action s (attention-as-action). Object B is selected for the action a by directing the focus of attention, and gaze, toward it (selection-for-action). The focus of attention is used as an implicit argument for the action. The focus of attention refers to the object B without explicitly representing all of its properties (deictic reference).

Learning Mechanisms

The central claim of this article is that attention can be controlled in the same way as actions using similar learning mechanisms and by related areas of the brain. The principles described above allow learning mechanisms such as habituation and conditioning to play a natural role in the control of attention. Although theories of habituation and conditioning have originally been developed to explain the learning of motor and emotional responses, the view of attention as action opens up for a marriage between theories of learning and attention. This does not mean that the concept of attention has not made its way into learning theories. On the contrary, it is well

represented in a number of modern learning theories (Mackintosh, 1975, Pearce & Hall, 1980, Wagner, 1981, Schmajuk, Lam & Gray, 1996). However, there has been very little interest in learning mechanisms in the study of attention itself although many of the experiments in the field can be considered instances of conditioning.

This section reviews a number of learning mechanisms that may play a role in attention. These mechanisms will be used below in a computational model that shows how they can be used in the control of attention. A number of different learning systems in the brain are described and related to the learning theoretical concepts of habituation, classical and instrumental conditioning. However, these categories are operationally defined to fit certain experimental procedures and one cannot hope for a perfect match with brain systems. As we will see, it becomes necessary to distinguish between more than these three types of learning.

Habituation

Habituation can be defined as a learning process where an animal learns to ignore a stimulus that does not predict anything of value to it. Usually this decreasing interest in a stimulus is studied through its effect on the orienting response toward the stimulus. This reaction can be operationally defined as any response that (1) is elicited by novel stimuli of any modality, and (2) habituates upon repetition of the stimulus (Gray, 1975). This definition makes clear that the orienting reaction is directed toward novelty and thus plays an important role in the exploration of the environment. The other part of the definition describes the connection between the orienting reaction and habituation.

Habituation in its most basic form is deceptively simple to describe. Each time a novel stimulus is presented, the orienting response will become weaker (Wang, 1995). This could result from either an increased threshold for response generation or decreased sensitivity to the eliciting stimulus.

Many brain regions have been implicated in the control of habituation. In the simplest case, reference is made to the properties of many neurons to habituate upon repeated activation (Wang, 1995). Depending on whether these neurons are found in the sensory or motor regions of the brain, the effect can be conceived of as either sensory or motor related habituation. Habituation of this type is non-associative since it only influences a single signal without reference to the situation in which it occurs or other stimuli that are present.

The existence of neurons with this type of habituating properties has secluded the fact that many types of habituation can not be explained in this simple way (Gray, 1975). The simplicity of the basic habituation situation should not be taken as evidence for the lack of complexity of the phenomenon. To end any such misconceptions, let us look at the events that will cause the orienting reaction toward a certain stimulus to be reinstated (Sokolov, 1960, Gray, 1975).

First, when a stimulus is changed along some stimulus dimension, the orienting reaction will typically reappear. For a sound stimulus, this happens when its pitch, loudness, or duration changes. The orienting reaction not only reappears when the loudness increases, it may as well be caused by decreased loudness. In fact, the orienting reaction can be elicited by the omission of an expected sound. This is a

limiting case when the orienting reaction is caused by nothing when something was expected (O'Keefe & Nadel, 1978).

Second, the passage of time can also reinstate the orienting reaction (Thompson & Spencer, 1966). A stimulus to which an animal has habituated will elicit the orienting reaction again if it has not been encountered for a while. This phenomenon is similar to spontaneous recovery in conditioning (Pavlov, 1927).

Third, the orienting reaction also reappears when a new novel stimulus is presented (Thompson & Spencer, 1966). This new stimulus will make the organism more likely to attend to the original stimulus again. This situation is called dishabituation (Gray, 1975, Wang, 1995). A possible explanation for this effect is that there is a direct link between the detection of novelty and the temporary shut down of the habituation system.

Fourth, a possibly related fact is that a new context will also cause the orienting reaction to reappear (Gray, 1975, O'Keefe & Nadel, 1978). Since this is similar to the situation above it may be explained by the same mechanisms. However, it does not seem entirely unreasonable that a novel object in an old context should be distinguished from a known object in a novel context.

Fifth, drowsiness can also cause the orienting reaction to appear again. This is called overhabituation (Sokolov, 1960, Gray, 1975). Quite surprisingly, it has been shown that caffeine removes the overhabituated orienting reaction. Apparently, caffeine activates a system in the brain that will then inhibit the systems responsible for the orienting reaction.

This suggests an alternative view of habituation as a process that learns what to expect in a certain situation or context. Habituation occurs when a stimulus is entirely expected in terms of previous events (Gray, 1982, Balkenius, 1995). This view is similar to the idea that habituation is the result of a match of incoming stimuli to a neural model of that stimulus (Sokolov, 1960). It is more general, however, since it allows expectations to be the result of complex processes at a cognitive level and not simply the formation of a template for the stimulus.

This type of associative habituation could result from associating the context with the expected stimuli in that context. Alternatively, the context could acquire inhibitory influences on the orienting reaction toward expected stimuli. Both these mechanisms are possible and there is reason to believe that they both exist in the brain. The contextual inhibition of attention to an irrelevant stimulus is one of the main mechanisms used in the model below to control attention.

The septohippocampal system has been implicated in the matching between expected and actual states of the world (Gray, 1982). According to Gray's theory, a mismatch in the septohippocampal system results in the activation of a behavioral inhibitions system that shuts off any ongoing behavior and prepares the organism for action as well as increasing attention to the environment. This comparator function is a recurrent theme of many theories of the hippocampus and its related areas (Smythies, 1966, O'Keefe and Nadel, 1978, Gray, 1982).

As mentioned above, the inhibitory influence on the orienting reaction by prefrontal cortex is well documented (Fuster, 1997). Animals with prefrontal lesions generally have problems habituating the orienting reaction (Fuster, 1997, Kimble, Bagshaw & Pribram, 1965).

To summarize, I have described a number of habituation mechanisms. Two are non-associative and influence either the sensory side by limiting susceptibility to a stimulus or the motor side by increasing the threshold for response activation. The other type of mechanisms are associative and allow the context to influence the orienting reaction either by generating expectations that are compared to the environment or by directly inhibiting the orienting reaction. Since there exist evidence for each of these mechanisms, it would be too hasty to conclude that habituation is a unitary phenomenon. What I am promoting is a view of habituation as a collection of different mechanisms that operate according to different principles in different areas of the brain. What must be stressed for each type of habituation, however, is that habituation only occurs when a stimulus predicts nothing of importance and should not be attended. If a rewarding event follows the presentation of the neutral stimulus, a different type of learning will occur: classical conditioning.

Classical Conditioning

Classical conditioning is defined as a learning process that is controlled by the pairing of two stimuli (Mackintosh, 1983). Typically the stimulus presented first is neutral and the second is of some significance to the animal. The first stimulus is called the conditioned stimulus and the second is called the unconditioned stimulus. The result of learning can be seen in the animal's reaction to the neutral stimulus when it is presented on its own, the conditioned response.

While the initial reaction to a neutral stimulus is orientation, it will be replaced by a conditioned response when it has been paired with the unconditioned stimulus. This conditioned response typically resembles the initial unconditioned response that is produced by the unconditioned stimulus.

For example in eye-blink conditioning, the onset of a tone may signal that the administration of an electric current to the eye is imminent. The unconditioned response to the current is to blink and after a number of pairings of the tone and the current, animals will blink after the tone is presented, and slightly before the onset of the current. It is possible to distinguish between at least three types of classical conditioning depending on whether the animal learns about the significance of a stimulus, the appropriate response to it, or about its relation to other stimuli.

In the first type of classical conditioning, what is learned is the significance of an initially neutral stimulus. This was the type of conditioning studied by Pavlov (1927). When his dog recognized that a particular sound signaled the delivery of food, it learned about the significance of the sound for food related activities which consequently made it salivate when the sound was presented.

The likely site for this type of learning in the brain is in the amygdala (Rolls, 1992, Holland & Gallagher, 1999). The amygdala is perfectly placed for associating highly analyzed stimulus information with value since it receives projections both from the higher sensory areas and gustatory and olfactory inputs. The latter may represent the primary value of a food stimulus (Amaral, Price, Pitkänen & Camichael, 1992). The amygdala projects to a number of lower brain including the hypothalamus, which controls innate reactions to significant stimuli. It should be noted that the role of the amygdala in appetitive conditioning is controversial (LeDoux, 1995, Rolls, 1992). Most studies of this area has concentrated on its role in

aversive conditioning, especially the acquisition of fear reactions (LeDoux, 1995), but the evidence for an appetitive role as well is accumulating (Rolls, 1992, 1995, Ono & Nishijo, 1992, Holland & Gallagher, 1999).

Conditioning of this type can be viewed as a valuation of a stimulus rather than as stimulus response learning (Rolls, 1986, 1995). In the context of attention, the importance of this type of conditioning is that it lets an animal learn about what stimuli are of significance to it and should thus be attended (Rolls, 1986). This applies whether or not the stimulus signals a desirable or undesirable event.

There exists a second type of classical conditioning that is more accurately seen as stimulus-response learning. A paradigmatic example is the eye-blink conditioning described above where the onset of the tone causes a motor response that is precisely timed to just before the onset of the electric current. This type of learning has been shown to take place in the cerebellum (Thompson, 1988, Yeo & Hesslow, 1998). This type of conditioning differs from the first type since it is not driven by stimulus value but by the response it generates. In eye blink conditioning, the electric current is certainly aversive but it is not the aversive properties that control learning but instead the fact that a blink response is activated. Learning in the cerebellum can be seen as an automation of a response that could already be performed in the form of an unconditioned response (Thach, 1999). This could be the mechanism behind automatic attention.

A final type of classical conditioning occurs when an animal learns about the relation between two neutral stimuli. A situation which is called sensory preconditioning (Mackintosh, 1983). This may be the process that is responsible for the learning of the stored regularities, or expectations, required for a comparator-based explanation of habituation (Gray, 1995).

It is clear that classical conditioning can not be explained by a single learning mechanism. Instead, a number of interacting systems are needed. In more complicated cases of conditioning many more mechanisms must interact with the ones described above.

When a stimulus that has previously acquired the ability to generate a conditioned response is presented on its own without the unconditioned stimulus, the conditioned response will gradually disappear again. This process is called extinction and has a number of interesting properties.

First, when an unexpected stimulus is presented it can cause the extinguished response to reappear. Pavlov (1927) called this phenomenon disinhibition. It shows that extinction cannot be explained by the removal of the earlier acquired association. In Pavlov's (1927) account for extinction, he assumed that the initial learning was counteracted by the acquisition of an inhibitory association from the representation of the conditioned stimulus. Extinction was thus not seen as the removal of the initial association, but as the addition of yet another one. There is now ample evidence to support this view of extinction (Mackintosh, 1983).

Second, the passage of time will reinstate the extinguished response. This is called spontaneous recovery (Pavlov, 1927). This has been explained either as a passive process or as an effect of a change in the temporal context of the extinction that would render the extinction less effective.

Third, unlike initial acquisition of a response, extinction does not generalize to new situations or contexts. What has been extinguished in one situation will reappear in another. While the initial learning of a response is associated mainly with the conditioned stimulus, the extinction appears to be controlled by the context (Bouton & Nelson, 1998).

It is interesting to note that these properties are very similar to the ones described for habituation above. It suggests that both processes can be explained in terms of a contextually controlled learning system that acquires an inhibitory influence on either the innate orienting reactions or on classically conditioned responses. Such a contextual system is an important component of the model presented below.

This view is supported by studies of the orbitofrontal cortex in animals that appears to react to “frustrative nonreward”, that is, a situation when an expected reward is omitted (Amsel, 1962). Rolls (1995) has advanced the view that the orbitofrontal cortex is responsible for the detection of an omitted reward or punishment. When expectations are not met, the orbitofrontal system learns to inhibit the initially learned association in the amygdala.

This view is also supported by lesions studies where animals with frontal lesions have been reported to be frustration resistant, ignoring the omission of expected reward (Fuster, 1997). Since omission of an expected reward is what causes extinction, one would expect that extinction is impaired by frontal lesions and this is in fact the case (Tanaka, 1973, LeDoux et al., 1989). When frontal animals are requested to extinguish a previously conditioned behavior, they persevere in their learned behavior.

Instrumental Conditioning

In instrumental (or operant) conditioning an animal learns about the relation between a response performed in a certain situation and a reward or punishment that follows it (Mackintosh, 1983). In contrast with classical conditioning, the presentation of the rewarding or punishing stimulus depends on what the animal does.

There are essentially four primary situations in which an animal can learn by instrumental conditioning: presentation of a reward, omission of expected reward, presentation of punishment, and omission of expected punishment (Gray, 1975). Here, I will only describe the first two cases since they are most interesting for a theory of attention.

For learning by reward, a simple trial-and-error based mechanism is conceivable. Initially, responses are generated at random but as rewards are presented, the rewarded responses will become more and more likely in the current situation until the animal has learned the correct response.

There are number of reasons why this simple explanation is not sufficient to explain instrumental learning in animals. In any realistic situation, there exist a very large number of potential responses. To test them all at random would be intractable. If a sequence of responses must be performed before the reward is received, the situation becomes even worse. The number of possible sequences grows exponentially with the number of steps in the sequence. Even for very short sequences of responses, the possible number of combinations becomes astronomical.

There are three interrelated solutions to this problem. The first is to use an innate selection mechanism that tries out responses using a better strategy than random selection. The second possibility is to generalize from prior situation and to initially try responses that have yielded rewards in similar situations before. Such a generalization mechanism can drastically alter the learning times for a sequence of instrumental responses (Balkenius, 1996). A final possibility is to use knowledge of what responses are possible (or reasonable) in a given situation to limit the search to these responses.

A possible site for instrumental learning in the brain is in the basal ganglia and the related frontal cortex (Wang, Aigner & Mishkin, 1990, Houk et al., 1995, Knowlton, Mangels & Squire, 1996). Damage to the basal ganglia such as in Parkinson's disease is known to interfere with effective motor control (Kolb & Whishaw, 1990). A striking result of this disorder is the inability to initiate movements.

Neurons have been found in the basal ganglia that respond to reward related events, both primary and conditioned rewards (Schultz, Romo, Ljungberg, Mirenowicz, Jollerman & Dickinson, 1995). The basal ganglia receive the required input signals to perform an instrumental learning function. Information about primary reward and punishment enters in the pathway from the amygdala (Amaral et al., 1992). Information about the current sensory situation is received from all of sensory cortex (Schultz et al., 1995), while contextual information and signals related to the current step in a motor program comes from frontal cortex (Gray, 1995). Motor cortex can also inform the basal ganglia about the selected response, which is needed to correlate the rewards with it.

The main outputs of the basal ganglia are directed toward primary motor cortex and premotor cortex as well as prefrontal cortex through the thalamus (Fuster, 1993). These outputs do not activate actions directly but appear to represent high-level instructions that are implemented by more posterior motor systems (Goldman-Rakic, 1987). The basal ganglia are also able to influence the superior colliculus, which is responsible for the generation of eye-movements and orienting reactions (Nauta & Domesick, 1984).

Like classically conditioned responses, instrumental responses can be extinguished by the omission of reward. The omitted reward serves to limit the generalization of the situation in which the reward is given. However, it is not initially obvious whether the reward was omitted because the situation is different or is simply a sign of an inherent randomness of the reward. Another complicating factor is that the omission of reward activates innate "frustration" systems that may trigger aggressive behaviors which will interfere with learning (Amsel, 1962).

I have already described the role of prefrontal cortex in the detection of omitted reward in classical conditioning and it is tempting to extend this role also to the instrumental case. Fortunately, there are empirical data to support this view. There exist neurons in the prefrontal cortex that project to the basal ganglia, which could support the required suppression of non-rewarded instrumental behavior (Fuster, 1997, Maurice, Deniau, Glowinski & Thierry, 1998). These connections could also govern discrimination learning between the contexts where a response produces a reward and contexts where it does not.

In humans, frontal lesions result in an inability to change behavior that is no longer appropriate (Shimamura 1995, Kolb and Whishaw 1990). For example, in the Wisconsin card-sorting test, subjects are asked to first figure out how to sort cards according to a simple criterion such as color. When the subjects have succeeded, the criterion is changed and the subjects have to figure out the new sorting-rule. Patients with frontal injury are often unable to do this. They may be able to verbalize that the rules have changed but they will persevere in their incorrect behavior.

A final learning mechanism that may be involved in instrumental learning is the automatication described above. If the cerebellum learns to automatically perform extensively rehearsed actions in their appropriate contexts, then instrumentally controlled responses, as well as classically controlled, will eventually become automatic. This is often referred to as over-learning (Mackintosh, 1983).

A Computational Model

In an attempt to synthesize the different ideas presented above, a computational model of visual attention has been developed. The model attempts to make use of the different learning mechanisms described above for the control of visual attention. The goal is to show that a view of attention as action allows the various learning systems of the brain to make sense as parts of an attentional system.

The model is computational in two senses. First, it has a computational realization that allows it to be implemented in a computer or a robot. Second, it is the computational goal of the system rather than a replication of physiological detail that has guided the development of the model. The model should be understood at a functional level and not as a physiological model. However, the model was developed with specific brain regions in mind and this makes it possible to compare the operation of the different parts of the model with that of specific brain areas.

The aim of the model is to make clear the architecture of attention. It does not try to explain the outcome of one particular experiment. It is thus a model of a whole organism and not a single experiment. However, it is not possible to design models that incorporate only fully detail subsystems. Instead, a reasonable aim is to place more emphasis on certain parts of the system and less on others. In the model, the role of each subsystem in attention is stressed while their functions in, for example, motor control are mainly left out.

To simplify the description, the difference between shift of attention and shift of gaze is ignored and the terms “attentional shift”, “saccades” and “orientation” will be used interchangeably. This is certainly a simplification but emphasizes their close relation (Shimojo, Tanaka, Watanabe, 1996). A consequence of this design decision is that the important class of overt attention is not yet included in the model.

A further simplification is that only external stimuli will be discussed. The ability to focus attention on internally generated stimuli or events separated from the present external situation will not be considered here.

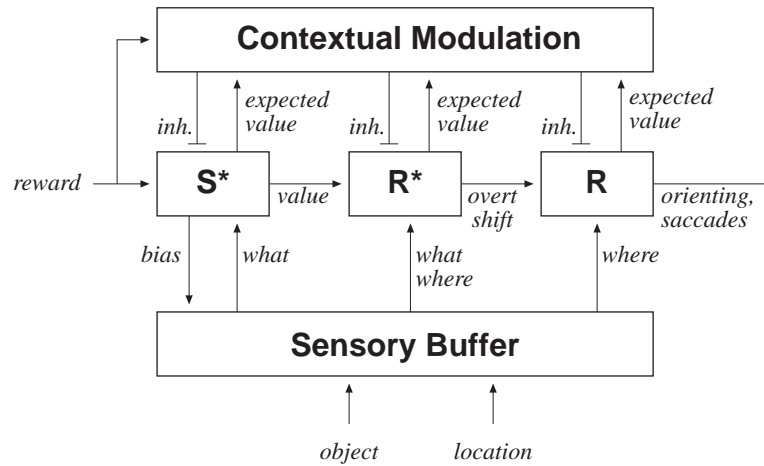


Figure 2: Overview of the components of the computational model and their main interactions. The current scene is stored in an sensory buffer containing both spatial and feature data. A location cue can directly trigger a saccade through the response system R. Learned attentional shifts are controlled by the response learning system R* which can execute an overt attentional shift through R.

Figure 2 shows an overview of the subsystems of the computational model and their main interactions. There are three main levels of control in the model. On the lowest level, a sensory buffer stores the current sensory state including both the sensory properties of objects and their locations in different coordinate systems. On the second level, there are three subsystems that control the shift of attention from one location to another. They all use different principles of control and operate mainly independently of each other. The first system S* evaluates stimuli and uses classical conditioning to assign a value to each fixated stimulus. This subsystem can enhance storage of a stimulus in the sensory buffer and indirectly cause a saccade to the most highly valued stimulus. The second system R* evaluates responses and uses instrumental conditioning to learn when to perform an endogenous saccade. Finally, a system R executes saccades independently of any learning directly based on the activity levels in the sensory buffer. Common to the three systems subsystems, S*, R* and R, is that they are all excitatory in their influence on attention. They are also modulated by the contextual system on the third level. The context system is involved in habituation in R, extinction of value in S*, and extinction instrumental saccades in R*, that is, it has an inhibitory influence on attention.

The Sensory Buffer

A sensory buffer is a central structure of many models of visual processing (Marr, 1982, Treisman & Gelade, 1980, Kosslyn, 1994). The sensory buffer is where the attentional selection is assumed to operate through enhancing or suppressing visual representations. The direction of visual attention can be controlled either by visual features or by spatial location. It is well known that simple features of the visual

scene can be used directly to attract attention to a location (Treisman, 1988). This phenomenon has been extensively studied in the context of visual search. Subjects can almost instantaneously find, for example, a single red circle among many blue circles, or a square surrounded by circles. When they have to find an object characterized by a conjunction of properties, such as color and shape, search becomes much slower and appears to require sequential processing. The conclusion of these studies is that attention can be directed toward a target that can be easily distinguished from its surrounding. This implies that a feature cue such as color can easily be converted to a location cue. Attention can also be object based. In this case, a specific object in the scene is selected and processing of that object is enhanced.

In the implementation of the model, five parallel feature channels are used for the sensory buffer. The first four channels code the color of each stimulus in red/green and yellow/blue components. The last channel includes a saliency measure for each location on the retina independently of its color. The representations in the different feature channels all use the same retinotopic code. The various learning systems are assumed to work directly on these different channels and higher level processes, such as object recognition, are not included in the model. Taken together, the output from all the feature channels forms a distributed representation of the stimuli in view.

Whether attention is directed by low-level features or a whole object, it always has a spatial component (Kosslyn, 1994). It is clear that this spatial component of attention does not use a common frame of reference for all locations. Instead there appear to exist a number of coordinate systems that are used for different purposes. One line of evidence bearing on this view comes from the studies of patients with neglect (Allport, 1990). Persons with unilateral neglect appear to be totally unaware of space on one side of the body. In many patients however, the neglect appears to only involve parts of space. Spatial attention may be impaired toward distant objects or parts of their own body (Bisiach, Peani Vallar & Berti, 1986). In other patients, neglect affects one side of an object. Patients with this condition may only eat the food to the left on their plate. Another variant is attentional deficits that relate to the location of an intended movement. The existence of multiple coordinate systems indicates the need for mechanisms for coordinate transformations. Furthermore, it becomes necessary to learn in which coordinate system a certain action or expectation should be represented.

In the model, the role of the sensory buffer is to represent the currently viewed part of the environment using a distributed representation. This representation contains both stimulus identification and localization. The identification part consists of representations at different levels. There are specific representations of the identity of a fixated stimulus, but also more general representations of the color and salience of stimuli at different locations on the retina. These more general representations are also partially spatial since they are represented in a retinotopic map and can be used for localization. The location of the fixated stimulus is also multiply represented in different coordinate systems. In the current implementation of the model there are two coordinate systems: a relative coordinate system anchored at the retina and an absolute coordinate system fixed to the environment.

The different parts of the sensory buffer communicate in such a way that activity in one part of the buffer will enhance processing in other related areas. For example, the activation of a location cue will enhance processing of color at that location. Conversely, the activation of a color cue will increase the activity for a stimulus of that color at that specific location. In a similar way, a more salient stimulus will produce an increased representation of its corresponding location.

Fixed Responses

The salience of the different locations activated in the sensory buffer is used by the response system R to generate a saccade toward the most salient stimulus. As a result the most salient stimulus will be moved to the center of the retina. The response system is thus responsible for exogenous saccades.

In principle, it would be possible to scan the visual scene with this response system if the target for each saccade was selected stochastically based on the relative salience of the different stimuli. However, in the model, a different solution is used. Each time a stimulus is attended, the fixed response system generates a fixed nominal value for the target. This expected value is compared to the actual value of the target. If the actual value is less than the expected value, the contextual system will learn to inhibit further saccades to the target. As a result, the orienting to the stimulus will habituate. In the short time scale, the suppression from the context implements inhibition-of-return. On a longer time-scale, the selection of targets for saccades comes under contextual control since the inhibition of the saccade will be specific to the current context.

The sensory buffer together with the response system will scan a scene in a sequential manner gradually losing interest in stimuli that are not of value to it. Since habituation is under contextual control, the attention to stimuli will reappear when the context is changed.

Stimulus Evaluation

The module S* is responsible for the evaluation of stimuli. (The star is used here to indicate valuation). Stimuli can gain value in two ways. First, there is a set of innately valued stimuli, called rewards. These have a fixed value independently of any learning. Second, there is a set of initially neutral stimuli that can gain a value by association with a reward. The process involved is classical conditioning where a stimulus that predicts reward will increase its value as described above.

The value of a stimulus is used as a bias in the sensory buffer in such a way that stimuli with higher value will obtain increased representation. As a result, a highly valued stimulus may gain access to the response system and generate a saccade toward it. This implies that S* can control saccades even though it does not directly control any responses. These saccades are exogenous since the stimulus that will be attended must be present in the visual field before it can be enhanced and selected in the sensory buffer.

A second role of the evaluation is to prevent the habituation in the response system to stimuli that are of value. If a stimulus no longer predicts a reward, the contextual system will learn to suppress its value. As a consequence, S* will no longer

enhance its storage in the sensory buffer. Additionally, the suppressed value will make habituation to the stimulus possible. The suppression of value and the subsequent habituation is the equivalent of extinction in ordinary classical conditioning.

The inhibition of stimulus value is controlled by the context and as a consequence, attention to an extinguished stimulus will reappear if the context is changed. After repeated extinction in different contexts, this will allow the evaluation mechanism to become highly context specific.

Response Evaluation

The final way to focus attention is controlled by the response learning system R^* . This subsystem learns the correlation between a saccade performed by the response system R and a subsequent reward. When a saccade makes a valued stimulus visible, the response learning system will associate the current state of the sensory buffer with that saccade. This is a form of instrumental conditioning. As a result, this saccade can be produced automatically at a later time. Since the saccade is triggered by something other than the target-stimulus, the attentional shift can in principle occur before the target appears. This is suggested to be the mechanism behind endogenous saccades.

Like stimulus evaluation, the evaluation of responses may change as a result of extinction. When a saccade fails to acquire the expected target, the context system will learn to inhibit the response and lower its value. The mechanism is analogous to that for stimulus evaluation and will place instrumental saccades under contextual control.

The extinction mechanism is necessary to cope with changing environments, but poses a problem for saccades that occur before the target has appeared. During the time when the target is not yet present the response should not be extinguished. To overcome this problem, the response learning system sends a prediction of the forthcoming stimulus value to the contextual system. This value will prevent extinction of the response during the inter-stimulus interval. When the target stimulus eventually appears, this value is retracted to prevent over-expectation (Kremer, 1978). Since the stimulus itself now enforces its value through the stimulus evaluation system, the expected value is no longer needed. To make extinction possible, the attention to the expected location of the target stimulus eventually times out.

Context System

The role of the context system in the control of habituation and extinction was described above, but we now need to consider how this contextual representation itself is constructed. I want to propose that a contextual representation can be constructed from a sequence of attentional fixations. Each time a stimulus enters into the sensory buffer and is selected as target for an attentional shift, it is also included in the context. The context is thus a sequence of fixations where each fixation includes both a spatial location and the stimulus information at that location.

The contextual system has two main functions in the model. The first is to inhibit unsuccessful learning in R , R^* and S^* as described above. The second function of the context is to generate expectations of stimuli that can be matched to the stimulus

in the current focus of attention. If expectations are not met, the contextual representation is reset. This will temporarily shut down the inhibition of the other learning systems and cause disinhibition or dishabituation.

Computer Simulations

The model described above has been realized in a computer implementation where the behavior of the model can be assessed. This section describes the implementation and a number of simulation experiments that illustrates the function of the model.

The scene used in the simulation consists of a two dimensional array with 21×21 locations, each of which can contain a stimulus of different color and luminance. Since visual recognition is not simulated, these stimuli are considered as separate and already segmented. Color substitutes for different stimulus properties and luminance represents the salience of a stimulus.

A stimulus can have three different roles in the simulations. First, it can act as a *target*. A target stimulus is rewarding and initiates learning in both S^* and R^* . Second, a stimulus can be a *cue*. A cue has the potential to inform the system about the location or timing of a target. Finally, a stimulus can be a *distractor*. Such a stimulus attracts attention and decreases the time left for looking at targets. Note that it is possible for a stimulus to function as both cue and distractor at the same time.

A retina covering 9×9 locations can move over the scene and sample the visual pattern at different locations. A fovea consisting of 5×5 cells is located in the middle of the retina and reacts to color as well as salience and produces four color sensitive feature channels. The rest of the retina only reacts to the salience of a stimulus and produces the salience channel. A relatively higher salience is assigned to transient stimuli compared to sustained.

To simplify the implementation, only two coordinate systems are used: one centered at the retina and one centered at the body. Since the retina is the only moveable component of the model, there is no difference between body and environment coordinates. The retinal coordinates are relative since they move with the fixation, and the body-centered coordinates are absolute.

The fixed response system reacts to both transient and sustained visual stimulation and generates a saccade toward the most salient stimulus. Using only the fixed response system R the model will saccade to the stimuli in the visual field with a probability proportional to their salience. Since R is color-blind, it is not able to distinguish the targets from distractors, but has the potential to move the color sensitive fovea to a target or a cue.

The contextual system inhibits saccades to locations that are not rewarding. The actual value of the stimulus currently in focus is constantly compared to the expected value of the stimulus. The actual value of a stimulus is set high if it is a target and zero otherwise. The expected value of a target selected by the fixed response system is set to a nominal low value. A discrepancy when the system fixates a non-target will lead to inhibitory learning in the connections from the context system to R , which will eventually cause habituation.

S* collects reward from the environment and biases the processing within the sensory buffer as described above. The learning processes involved are described in more detail in Balkenius and Morén (2000).

The response learning system R* was implemented as a reinforcement learning architecture. This part of the model is described in detail in Balkenius (1996) and Balkenius and Morén (1999).

Finally, the contextual system was implemented in a very coarse way. A context is assumed to consist of a template for the scene where the probability for the occurrence of a stimulus at each location is stored. This estimate is updated each time a location is fixated. When a large discrepancy is found between the stored probability and the actual target, the contextual template is reset and the learning starts over. This covers the basic intuition of a contextual representation but does not support the shift back and fourth between different stored contexts.

Simulation 1: Habituation

In the first simulation, a simple scene was presented to the system consisting of one target in the middle and four blinking distractors (figure 3, phase 1). The isolated target was shown for 5 simulated time steps (ticks), and then the target and four distractors were shown for 5 ticks. The behavior of the model is shown in figure 4. The percentage of the time spent at the target is plotted as a function of time. It takes approximately 200 ticks for the attention to habituate to the distractors. When the system had habituated, the blinking distractors were moved to four new locations and the system started to look at the distractors again (figure 3, phase 2). If the distractors are considered a part of the context, then this is a form of dishabituation. As can be seen in figure 4, the model eventually habituates a second time.

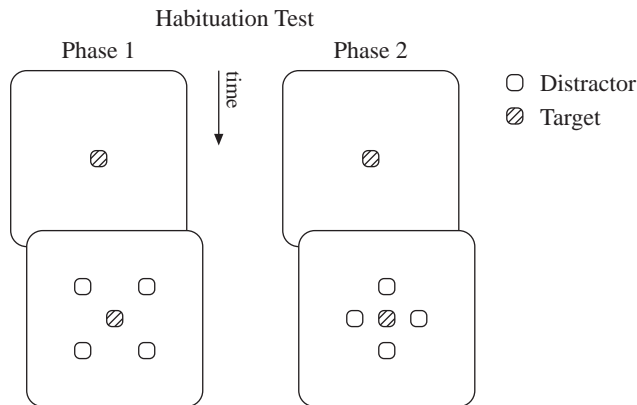


Figure 3: Two stimulus sequences used in the simulation of habituation.

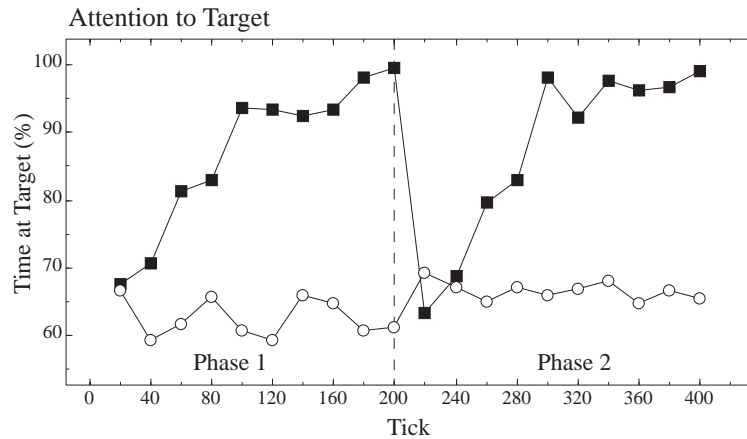


Figure 4: Attention to the target (%) as a function of time for the fixed response system R alone (circles) and for the complete model (solid squares). The increased attention to the target is a result of habituation to the distractors. Average of ten simulations.

The same simulation was also run with the response system R on its own, without the influence of the context system. In this case, the model could not habituate. This behavior is shown in the lower curve in figure 4. However, since the target was the only stimulus in the scene half of the time, more than 50% of the time was still spent at the target.

Simulation 2: Instrumental Saccades

The next simulations were devised to test if the model could use visual cues to control instrumental saccades. Two different types of trials were presented to the system in semi-random order. In the first type of trial, a red cue presented at the center of the scene was followed by the presentation of the target to the left (figure 5, phase 1). In the second type of trial, a blue cue identified that the target would appear to the right (figure 5, phase 2).

Figure 6 shows the time spent at the cue and the target location as a function of time. As can be seen, the model quickly learns to use the cue to shift attention to the expected location of the target. On the first trial, the shift of attention to the target is controlled by an exogenous saccade since it occurs after the presentation of the target. However, on the next trial the saccade is already endogenous and is executed directly when the cue is presented. It is shown that the model eventually spends 25% of the time looking at the target. To spend this time on the target, the model must already be fixating the target location when the target appears. This performance level indicates that the model has learned to perform an anticipatory saccade.

The graph also shows how much time the model waits for the cue at the center of the scene. This time gradually increases from around 40% to close to 70%. The extra 5% of the time not accounted for in the graph is the time spent at the target locations before the target appears.

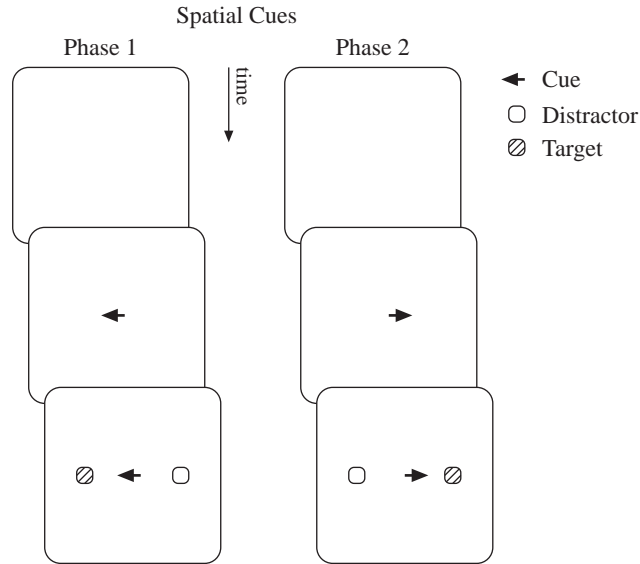


Figure 5: Two instrumental trials. The cue (actually a red or blue dot) signals the direction to the target (a yellow dot). The distractor is used to rule out that an exogenous saccade could find the target.

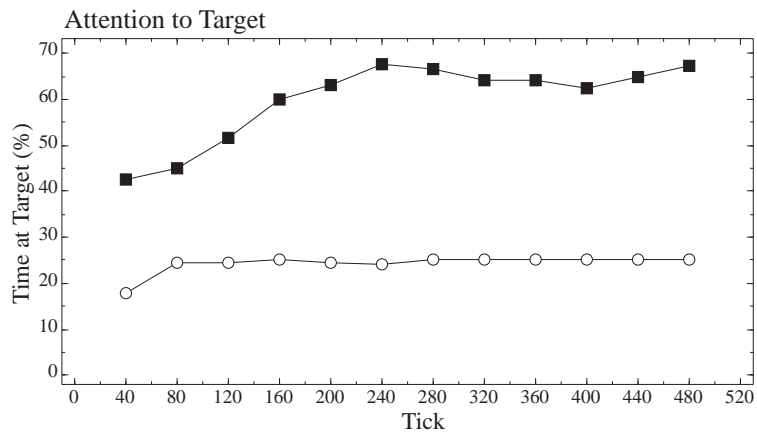


Figure 6: The time spent at a target (circle) and at the central cue location (solid square). The model quickly learns to saccade directly to the target when the cue is detected. It also gradually learns to wait at the center for the cue. Each trial consists of 40 ticks and the cues are visible for 20 ticks and the targets for 10 ticks. Average of ten simulations.

A different form of the simulation was also run where the cue location was selected at random in the scene and signaled the relative location of the target. In this case, the model learned to use the relative coordinate system rather than the absolute. This is a form of discrimination learning. The overall score was lower than in the previous simulation since there were many cases when the cue appear at a location outside the retina and what thus not seen. As could be expected in a discrimination procedure, learning was also slower since the two coordinate systems did not agree.

Simulation 3: Latent Inhibition

The final simulation was run to see if the model was susceptible to latent inhibition. First a cue stimulus was presented on its own for ten trials in a context of ten other stimuli. This eventually lead to habituation. Next, the cue was made a predictor of the target location as in simulation 2. The learning rate in this simulation was compared with that of simulation 2 where the cue was a predictor from the start. As could be expected, learning was substantially slower when the cue had first been shown without the target. This is a form of latent inhibition.

The ten context stimuli where necessary to prevent that the introduction of the target caused dishabituation which would prevent latent inhibition. With this more complex context, a single novel stimulus was not enough to reset the context.

Discussion

The goal of the above model is to show that a view of attention as action lets the various learning systems of the brain make sense as parts of an attentional system. To do this we need to identify the part of the computational with different brain areas. Figure 7 shows an overview of the brain systems involved in the processes addressed by the model. The next few sections discuss the relation between the sub-systems of the computational model and various brain systems.

Exogenous Saccades and Habituation

In the brain, the lowest level of attentional control is governed by the superior colliculus (SC) which is the equivalent of the fixed response system R in the model. In contrast to the cortical visual system, which receives its input through the lateral geniculate nucleus, the superior colliculus receives its input directly from the retina (Stein & Meredith, 1993). This attentional system reacts directly to salient stimuli without any high level of visual analysis. The system reacts best to transient stimulation in the temporal hemifield but receives input from the whole retina (Rafal & Robertson, 1995). The receptive fields in superior colliculus are retinotopic but it is likely that the representation of space is anchored to the body rather than the retina. This is necessary for the production of meaningful orienting movements regardless of eye position. Such body-centered coordinates are assumed to result from a biasing of the receptive fields with eye and body position in the form of planar gain-fields (Andersson & Zipser, 1990).

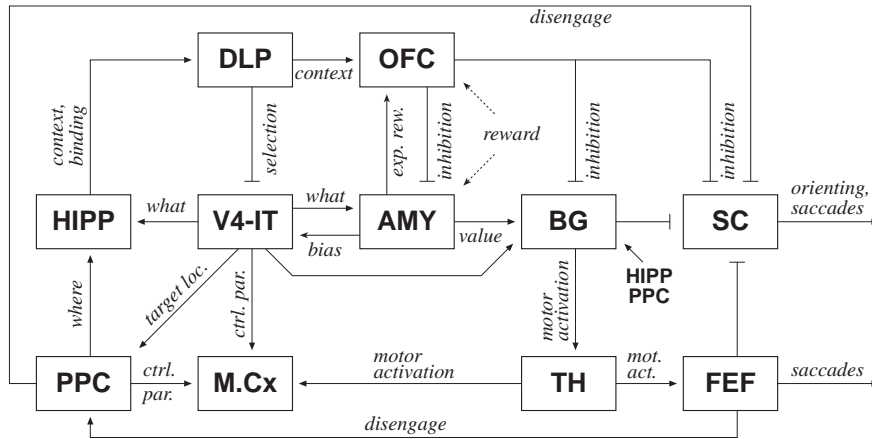


Figure 7: Overview of the components of the model and their main interactions. PPC: posterior parietal cortex, M.Cx: primary and premotor cortex including supplementary motor cortex, TH: thalamus, FEF: frontal eye field, HIPP: hippocampal system including entorhinal cortex, V4-IT: The ventral visual stream, AMY: amygdala, BG: basal ganglia and related structures including substantia nigra, SC: superior colliculus, DLP: dorsolateral prefrontal cortex, OFC: orbital prefrontal cortex. (Some of the connections are indirect).

The superior colliculus probably also includes a mechanism implementing inhibition of return (Johnson, 1995, Shimojo, Tanaka, Watanabe, 1996). This mechanism prevents that focus of attention remains at the same location too long by inhibiting attention to the currently attended location. As a consequence, a saccade to some other target will be generated. Inhibition of return functions in an environmental or at least body centered frame of reference (Shimojo, Tanaka, Watanabe, 1996). It is possible that also the posterior parietal cortex (PPC) is involved in this process (Rafal & Robertson, 1995). Such a mechanism was not included in the model however.

The superior colliculus implements a low-level mechanism for visual scanning of the environment. Since the pattern processing abilities of the superior colliculus are limited, the generated saccades are guided by simple low-level cues such as the contrast or intensity of the visual stimulus (Stein & Meredith, 1993). On its own, the superior colliculus system is hyper-reactive, it will continue to perform saccades to everything in the environment without end (Fuster, 1997). This is essentially what happened in the model in the second part of simulation 1 when the response system R was simulated on its own.

To reduce the number of saccades, it becomes necessary to inhibit saccades to parts of the scene that are irrelevant to the animal. It has been suggested that the orbital prefrontal cortex (OFC) is involved in this function (Fuster, 1997). The OFC could act in a way similar to the contextual system of the model. The OFC receives input coding both the current context and whether a reward is received or not. If a reward does not follow a saccade to a particular stimulus, the OFC can gradually acquire inhibitory control of the saccade. As a consequence, the saccade or orienting

reaction will habituate. This is what happened when contextual inhibition was included in simulation 1.

Stimulus Evaluation in the Brain

When an animal fixates a stimulus, it will be processed in the ventral visual stream (Mishkin, Ungerleider & Macko, 1983). This system is involved with the identification of a stimulus that starts in striate cortex and continues through the inferior temporal area (IT). The highly analyzed visual stimulus in IT is then associated with the reward in the amygdala (Rolls, 1995). The amygdala is suggested to be the equivalent of the sensory evaluation system S* in the model, while the ventral path is analogous to the feature half of the sensory buffer.

When the stimulus has been associated with reward, it will cause the amygdala to increase the general arousal of the organism (Holland & Gallagher, 1999). A second function of the amygdala, which has been suggested to be implemented through the back-projections to sensory cortex, is to selectively enhance processing of stimuli that have previously been paired with reward (Rolls 1995, LeDoux 1998). This enhancement will bias the attentional competition within the visual system toward important stimuli. This biased competition exploits the competition for storage in the visual system (Moran & Desimone, 1985, Desimone, 1996). By backprojections throughout the ventral visual system, the representation of the selected stimulus is selectively enhanced all the way back at least to V2 (Luck, Chelazzi, Hillyard & Desimone, 1997); a process that is probably aided by the pulvinar (Rafal & Posner, 1987). As a result of this competition, a saccade is made toward the winning target. Such saccades are controlled by the connections between V4-IT to the PPC with the support of the pulvinar. In the model, this corresponds to the connection between the spatial representations in the sensory buffer and the fixed response system R.

Working together, the superior colliculus, the orbital prefrontal cortex and the amygdala will make attention to irrelevant stimuli wane as a result of habituation, while attention to rewarded stimuli is enhanced.

In an unchanging world, this is all that would be needed of an attentional system. In reality, the environment often changes and some mechanism is needed that can retune the attentional system to new situations. The OFC is assumed to compare actual rewards received with the expected rewards generated by the amygdala (Rolls, 1995). When a mismatch occurs it will trigger learning that will establish an inhibitory influence on the amygdala that will cancel out the initial learning in a way analogous to habituation (Rolls, 1995). In the model, this was accomplished by the connection between the contextual system and S*.

Context and Expectation Matching

We now need to consider how the contextual representation enters the system in the first place. In the model, it is assumed that the context is a compressed representation of a sequence of location-object bindings. When the animal looks around, it will perceive a number of stimuli at different locations. The collection of all these glances will constitute the current context. The binding of a stimulus to a location has been suggested to take place in the hippocampus (Mishkin, Ungerleider & Macko, 1983), which receives input both from the visual system for stimulus identi-

fication (V4-IT) and the system for localization in posterior parietal cortex (PPC). A set of such bindings would constitute a good starting point for a contextual representation based on sequences of temporally organized stimulus events. This view of hippocampal function is well in accord with its role in configurational learning (Wickelgren, 1979, Rudy & Sutherland, 1989), spatial mapping (O'Keefe & Nadel, 1978), sequence recognition (Solomon, 1979), and working memory (Rawlins, 1985, Olton, 1986).

The dorsolateral prefrontal cortex is closely connected to the hippocampus (Nauta, 1964, Goldman-Rakic, Selemon & Schwartz, 1984, Amaral, 1987). It is possible that these connections carry the contextual representation to the prefrontal system where it can control inhibition and short-term memory. Another possibility would be that the hippocampus sends the contextual representation directly to the OFC (Fuster, 1997). In either case, an intact hippocampus would be necessary for normal habituation and extinction (O'Keefe & Nadel, 1978).

It has also been suggested that an expectation based mechanism which compares expected and actual states of the world exists in the hippocampal system, (Gray, 1982). If the mismatch between expected and actual stimuli occurs, the contextual representation from hippocampus will be reset. As a consequence, inhibition of the amygdala as well as superior colliculus will be removed causing disinhibition and dishabituation respectively. The intuition behind this mechanism is that unexpected events should be taken as evidence for a new context or situation. The behavior of the model agrees well with the responses to unexpected events in animals (Gray, 1975). As mentioned above, this mechanism was implemented in a very coarse way in the contextual system of the model.

The orbital system and the hippocampal system play two complementary roles in attention. While the OFC system inhibits motor actions, that is, orienting responses (Fuster, 1997), the hippocampal system works with sensory representations comparing actual and expected sensory situations (Gray, 1982).

Endogenous Saccades

We now turn to the role of the basal ganglia in the control of voluntary saccades, that is, attentional shifts that are not directly under stimulus control. The basal ganglia are assumed to collect information about the current sensory situation and state of the motor system from all over neocortex (Schultz et al., 1995). This information is correlated with value information received from the amygdala (Rolls & Williams, 1987). The suggested role of the basal ganglia is to learn by instrumental conditioning to trigger attentional shifts that agree with current goals. This is the function served by R^* in the model.

One type of instrumental control of attention is essentially a process of selection. The visual scene suggests a number of possible targets and the basal ganglia are responsible for selecting the target. This function could in principle be controlled by an inhibitory influence on the superior colliculus from the basal ganglia. Selection could function by inhibiting all but the appropriate saccade (Goldberg, Eggers & Gouras, 1991).

A second type of instrumental shift of attention occurs when the target of the attentional shift is not in view. In this case, prefrontal control is necessary to perform

anticipatory saccades. Especially the frontal eye fields (FEF) are crucial when monkeys make voluntary saccades (Sommer & Tehovnik, 1997). The frontal eye field is also able to disengage attention from the previous location through its connection with PPC, which in turn inhibits superior colliculus (Rafal & Robertson, 1995).

Such attentional shifts are assumed to be controlled by the pathway from the basal ganglia through the thalamus and the frontal eye fields which in turn inhibits the superior colliculus (Segraves & Goldberg, 1987, Rafal & Robertson, 1995) as well as activates the brain stem saccade generator (Segraves, 1992). The direct connection to the brain stem is necessary when the target is outside the visual field since in this case the superior colliculus will not be externally activated. Another conceivable way in which the basal ganglia could influence ocular movements is through the supplementary eye field (Schall, 1997).

The response learning system R^* of the computational model thus corresponds to a number of structures in the brain. While it makes instrumental saccades possible, it does not distinguish between disengaging from one location, shifting attention, and engaging attention at the new location (Posner, 1988). This is one of the main extensions that will be made to the model in the future.

It has been suggested that the prefrontal cortex plays an important role in extinguishing instrumental saccades through its connections with the basal ganglia (Nauta, 1964, Fuster, 1997). From a learning theoretical perspective, the generation of attentional shifts poses some special problems for the instrumental learning mechanisms. When an anticipatory saccade is performed, the target will, by definition, not yet be present. Without special precautions, this would trigger the inhibitory learning system in prefrontal cortex, which would extinguish the instrumental saccade while waiting for the target. One solution to this dilemma would be if the basal ganglia could generate an expectation of the target that can substitute for it until it appears (Hikosaka, Sakamoto & Usui, 1989, Apicela, Scarnati, Ljungberg & Schultz, 1992, Schultz, Apicela, Scarnati & Ljungberg, 1992). This expectation could be used to prevent extinction. When the target appears however, the expectation must be immediately shut off to prevent overexpectation which again would trigger inhibitory learning (Kremer, 1978). This corresponds to the expected value generated by R^* while the system is waiting for the target to appear as described above.

Although it is assumed that the learning process in the basal ganglia is governed by value generated by the amygdala, it is not the case that a visual target must be a primary rewarding object. In most cases, the basal ganglia could generate its own instrumental value through higher order conditioning (Houk et al., 1995). This is an additional role that could be played by the generated expectations.

The inhibitory inputs from prefrontal cortex could also be used for discrimination learning of the appropriate situation for an attentional shift as illustrated in simulation 2. Apart from direct sensory information from visual cortex, the basal ganglia also receive spatial information from PPC. This input includes a large number of coordinate descriptions of the location of the stimulus that is currently in focus of attention. The basal ganglia are thus in an eminent position to learn to discriminate between the different stimulus components and coordinate systems that should be used to perform an attentional shift.

Motor Control

According to the view of selection-for-action (Allport, 1990), the current focus of attention is used to control motor actions. This could be accomplished through the connections from parietal, and indirectly from temporal, to motor cortex where it sets the parameters for motor actions (Goodale & Milner, 1992, Murata, Fadiga, Fogassi et al., 1997, Fagg & Arbib, 1998). These actions may themselves be triggered by the basal ganglia in a way similar to attentional shifts.

The control of attention outlined above works with two sensory channels. The first is controlled by the focus of attention and selects one stimulus or stimulus configuration at a time. The other channel represents the context in which the focal attention should be interpreted. In many situations this is all that is required to control action. However, if several objects need to control action or have to be manipulated, the single attentional channel is not enough. It is suggested that this limitation is overcome by the use of deictic variables or visual indexes stored in a working memory in DLP (Ballard *et al.* 1997, Pylyshyn, 1994). These variables can support a change of attentional focus between a number of stimuli in the environment through their influence on visual processing in IT and through its interaction with the rest prefrontal cortex. The inclusion of such a mechanism is an important future development of the model.

Implications for Robots

Although the model presented above is primarily intended as a cognitive model, it also suggests a new way to implement artificial attention in a robot. The framework presented fits well with reinforcement learning architectures for robot control (Sutton & Barto, 1998), but introduces the novel idea that stimulus selection, as well as response selection, could be controlled by learning. This will make it possible for an autonomous robot to learn where to look to gain the most information about its environment. A robot implementation of the model, which uses input from two video cameras, is currently under way. Some of the first steps in this direction are described in Balkenius & Hulth (1999).

Limitations of the Model

The computational model is able to qualitatively reproduce many attentional and learning theoretical phenomena. It is not the intention of the model at present to produce quantitatively correct predictions. It is not clear whether a model at this level can do that. Instead, the goal of the simulations has been to show that the overall architecture is computationally sound.

As a learning model it introduces a new component in that stimuli are spatially located in relation to each other. This is an aspect of the environment that has been overlooked in previous accounts for conditioning. The framework presented above offers the possibility to study how placement of stimuli influences learning and attention. However, there are several components that have been left out that will need to be addressed in the future.

First, the context representation uses a local coding. The model is thus not able to recognize the similarity between different situations. Another problem is that the

contextual reset is all or none. A reasonable assumption is that different levels of mismatch should be able to reset the context to different degree.

Second, the role of working memory has not been investigated although the main components for such a memory are implemented in the computer program. It would be interesting to extend the model with a contextually controlled deictic memory store. An intriguing possibility is that the basal ganglia could control the storage and recall of working memory in the same way that it is assumed to control attention and motor actions in the model.

Third, the cerebellar system has not yet been included in the model. How will the addition of such a system influence the dynamics of the learning processes? According to the view described earlier, the cerebellum learns to make attentional shifts automatic after repeated training. This function is supported by the connections from the cerebellum to motor cortex (and FEF). This could possibly be the mechanism behind automatic processing, as opposed to the controlled processing described above (Shiffrin & Schneider, 1977).

Fourth, another simplification is that attention was identified with orienting. This excluded the possibility of covert attention that does not cause eye movements. This is an aspect of attention that will be included in the model in the future by extending the role of R*. A possible mechanism would be to let R* learn to either activate or inhibit representations in the sensory buffer.

Fifth, the bilateral structure of the different brain areas was not modeled. This excludes the possibility to model unilateral neglect.

References

- Allport, A. (1987). Selection-for-action: Some behavioral and neurophysiological considerations of attention and action. In H. Heuer & A. F. Sanders (Eds.), *Perspectives on perception and action*. Hillsdale, NJ: Erlbaum.
- Allport, A. (1990). Visual attention. In M. Posner (Ed.), *Foundations of Cognitive Science* (pp. 631-682). Cambridge, MA: MIT Press.
- Amaral, D. G. (1987). Memory: Anatomical organization of candidate brain regions. In F. Plum (Ed.) *Handbook of physiology: The nervous system V, Part 1* (pp. 211-294). Bethesda: American Physiological Society.
- Amaral, D. G., Price, J. L., Pitkänen, A., & Carmichael, S. T. (1992). Anatomical organization of the primate amygdaloid complex. In J. P. Aggleton (Ed.), *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 1-66). New York: Wiley.
- Amsel, A. (1962). Frustrative nonreward in partial reinforcement and discrimination learning: some recent history and a theoretical extension. *Psychological Review*, 69, 306-328.
- Apicela, P., Scarnati, E., Ljungberg, T., & Schultz, W. (1992). Neuronal activity in monkey striatum related to the expectation of predictable environmental events. *Journal of Neurophysiology*, 68, 945-960.
- Balkenius, C. (1995). Natural Intelligence in Artificial Creatures. *Lund University Cognitive Studies*, 37.
- Balkenius, C. (1996). Generalization in instrumental learning. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack & S. W. Wilson (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: The MIT Press/Bradford Books.

- Balkenius, C., & Hulth, N. (1999). Selection-for-action: a scheme for active perception, In *Proceedings of EUROBOT '99*, ETH, Zurich.
- Balkenius, C., & Morén, J. (1999). Dynamics of a classical conditioning model, *Autonomous Robots* 7 (1), 41-56.
- Balkenius, C., & Morén, J. (2000). Emotional learning: A computational model of the amygdala. *Cybernetics and Systems* (in press).
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997) Deictic codes for the embodiment of cognition, *Behavioral and Brain Sciences*, 20 (4), 723-767
- Bisiach, E., Perani, D., Vallar, G. & Berti, A. (1986). Unilateral neglect: Personal and extra-personal. *Neuropsychologia*, 24, 759-767.
- Bouton, M. E., & Nelson, J. B. (1998). Mechanisms of feature-positive and feature negative discrimination learning in an appetitive conditioning paradigm. In N. A. Schmajuk & P. C. Holland (Eds.), *Occasion Setting: Associative Learning and Cognition in Animals* (pp. 69-112). Washington, DC: American Psychological Association.
- Catiello, U. (1999). Mechanism of selection for the control of hand action. *Trends in Cognitive Science*, 3 (7), 264-271.
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Science USA*, 93, 13494-13449.
- Fagg, A. H., & Arbib, M. A. (1998) Modeling parietal-premotor interactions in primate control of grasping. *Neural Networks*, 11, 1277-1303.
- Fuster, J. M. (1993). Frontal lobes. *Current Opinion in Neurobiology*, 3, 160-165.
- Fuster, J. M. (1997). *The prefrontal cortex: Anatomy, physiology and neurophysiology of the frontal lobe*. Philadelphia: Lippincott-Raven.
- Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neuroscience*, 15, 20-25.
- Goldberg, M. E., Eggers, H. M., & Gouras, P. (1991). The ocular motor system. In E. R. Kandel, Schwartz, J. H., & Jessel, T. M. (Eds.), *Principles of Neural Science*. Amsterdam: Elsevier.
- Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum & V. B. Mountcastle (Eds.), *Handbook of Physiology: The nervous system V, Part 1* (pp. 373-417). Bethesda, Md: American Physiological Society.
- Goldman-Rakic, P. S., Selemon, L. D., & Schwartz, M. L. (1984). Dual pathways connecting the dorsolateral prefrontal cortex with the hippocampal formation and parahippocampal cortex in the rhesus monkey. *Neuroscience*, 12, 719-743.
- Gray, J. A. (1975). *Elements of a two-process theory of learning*. London: Academic Press.
- Gray, J. A. (1982). *The neuropsychology of anxiety: An enquiry into the function of the septo-hippocampal system*. New York: Oxford University Press.
- Gray, J. A. (1995). A Model of the Limbic System and Basal Ganglia: Application to anxiety and schizophrenia. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1165-1176). Cambridge, MA: MIT Press.
- Hikosaka, O., Sakamoto, M., & Usui, S. (1989). Functional properties of monkey caudate neurons. III. Activities related to expectation of target and reward. *Journal of Neurophysiology*, 61, 814-832.
- Holland, P. C., & Gallagher, M. (1999). Amygdala circuitry in attentional and representational processes. *Trends in Cognitive Science*, 3 (2), 65-73.
- Houk, J. C., Adams, J. L., & Barto, A. G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In J. C. Houk, J. L. Davis, & D. G. Beiser, *Models of information processing in the basal ganglia* (pp. 249-270). Cambridge, MA: MIT Press.
- James, W. (1890). *The principles of psychology*. New York: Holt.

- Johnson, M. H. (1995). The development of visual attention: a cognitive neuroscience perspective. In Gazzaniga, M. S. (Ed.), *The cognitive neurosciences* (pp. 735–747). Cambridge, MA: MIT Press.
- Kimble, D. P., Bagshaw, M. H., & Pribram, K. H. (1965). The GSR of monkeys during orienting and habituation after selective partial ablations of the cingulate and frontal cortex. *Neuropsychologia*, 3, 121-128.
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). *Science*, 273, 1399-1402.
- Kolb, B., & Whishaw, I. Q. (1990). *Fundamentals of human neuropsychology*. New York: W. H. Freeman.
- Kosslyn, S. M. (1994). *Image and Brain*. Cambridge, MA: MIT Press.
- Kremer, E. F. (1978) The Rescorla-Wagner model: Losses in associative in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavioral Processes*, 11, 15-34.
- LeDoux, J. E. (1995) In search of an emotional system in the brain: Leaping from fear to emotion and consciousness. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 1049-1061). Cambridge, MA: MIT Press.
- LeDoux, J. E. (1998). *The Emotional Brain: The mysterious Interpinnings of Emotional Life*. New York: Touchstone.
- LeDoux, J. E., Romanski, L. M., & Xagoraris, A. E. (1989). Indelibility of subcortical emotional memories. *Journal of Cognitive Neuroscience*, 1, 238-243.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & 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.
- Mackintosh, N. J. (1975). A theory of attention: variations in associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford: Oxford University Press.
- Marr, D. (1982). *Vision*. New York: Freeman.
- Maurice, N., Deniau, J.-M., Glowinski, J., & Thierry, A.-M. (1998). Relationships between the prefrontal cortex and the basal ganglia in the rat: Physiology of the corticostriatal circuits. *Journal of Neuroscience*, 18 (22), 9539-9546.
- Moran, J. & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229, 782-784.
- Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neuroscience*, 6, 414-417.
- Murata, A., Fadiga, L., Fogassi, L., Gallese, V., Raos, V., & Rizzolatti, G. (1997). Object representation in the ventral premotor cortex (F5) of the monkey. *Journal of Neurophysiology*, 78, 2226-2230.
- Nauta, W. J. H. (1964). Some efferent connection of the prefrontal cortex in the monkey. In J. M. Warren & K. Akert (Eds.), *The frontal granular cortex and behavior* (pp. 397-407). New York: McGraw-Hill.
- Nauta, W. J. H., & Domesick, V. B. (1984). Afferent and efferent relationships of the basal ganglia. In *Functions of the Basal Ganglia* (pp. 3-29). London: Pitman.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- Olton, D. S. (1986). Hippocampal function and memory for temporal context. In R. L. Isaacson & K. H. Pribram (Eds.), *The hippocampus, Vol 4* (pp. 281-298). New York: Plenum Press.
- Ono, T., & Nishijo, H. (1992). Neurophysiological basis of the Klüver-Bucy syndrome: responses of monkey amygdaloid neurons to biologically significant objects. In J. P. Aggleton (Ed.), *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 167-190). New York: Wiley.

- Pavlov, I. P. (1927). *Conditioned reflexes*. Oxford: Oxford University Press.
- Pearce, J. M., & Hall, G. (1980). A model of Pavlovian learning: variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review*, *87*, 532-552.
- Posner, M. I. (1988). Structures and functions of selective attention. In T. Boll & B. K. Bryant (Eds.), *Clinical neuropsychology and brain function*. Washington, DC: American Psychological Association.
- Pylyshyn, Z. W. (1994). Some primitive mechanisms of spatial attention. *Cognition*, *50*, 363-384.
- Trick, L. 1990. Subitizing and counting. Paper presented
- Rafal, R. D., & Posner, M. I. (1987). Deficits in human visual spatial attention following thalamic lesions. *Proceedings of the National Academy of Science USA*, *84*, 7349-7353.
- Rafal, R., & Robertsson, L. (1995). In M. S. Gazzaniga (Ed.) *The cognitive neurosciences* (pp. 625-648). Cambridge, MA: MIT Press.
- Rawlins, J. N. P. (1985). Associations across time: The hippocampus as a temporary memory store. *The Behavioral and Brain Sciences*, *8*, 479-496.
- Rolls, E. T. (1986). A theory of emotion, and its application to understanding the neural basis of emotion. In Oomura, Y. (Ed.) *Emotions: neural and chemical control* (pp. 325-344). Tokyo: Japan Scientific Societies Press.
- Rolls, E. T. (1992). Neurophysiology and function of the primate amygdala. In J. P. Aggleton (Ed.), *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (pp. 143-166). New York: Wiley.
- Rolls, E. T. (1995). A theory of emotion and consciousness, and its application to understanding the neural basis of emotion. In M. S. Gazzaniga (Ed.) *The cognitive neurosciences* (pp. 1091-1106). Cambridge, MA: MIT Press.
- Rolls, E. T., & Williams, G. V. (1987). Sensory and movement related neuronal activity in different regions of the primate striatum. In J. S. Schneider & T. I. Kinsky (Eds.), *Basal ganglia and behavior: Sensory aspects and motor functioning* (pp. 37-59). Berne: Hans Huber.
- Rudy, J. W., & Sutherland, R. J. (1989). The hippocampal formation is necessary for rats to learn and remember configurational discriminations. *Behavioral Brain Research*, *34*, 97-109.
- Schall, J. D. (1997). Visuomotor areas of the frontal lobe. In K. Rockland, A. Peters, & J. Kaas, *Extrastriate cortex of primates, volume 12 of cerebral cortex* (pp. 527-638). New York: Plenum.
- Schiffman, R. M. & Schneider, W. (1977). Controlled and automatic information processing: II Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*, 127-189.
- Schmajuk, N. A., Lam, Y. W., & Gray, J. A. (1996) Latent inhibition: A neural network approach. *Journal of Experimental Psychology: Animal Behavioral Processes*, *22*, 321-349.
- Schultz, W., Apicela, P., Scarnati, E., & Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *Journal of Neurophysiology*, *12*, 4595-4610.
- Schultz, W., Romo, R., Ljungberg, T., Mirenowicz, J., Jollerman, J. R., & Dickinson, A. (1995) Reward-related signals carried by dopamine neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233-248). Cambridge, MA: MIT Press.
- Segraves, M. A. (1992). Activity of the monkey frontal eye field neurons projecting to oculomotor regions of the pons. *Journal of Neurophysiology*, *68*, 1967-1985.
- Segraves, M. A., & Goldberg, M. E. (1987). Functional properties of the corticotectal neurons in the monkey's frontal eye fields. *Journal of Neurophysiology*, *58*, 1387-1419.

- Shimamura, A. P. (1995). Memory and frontal lobe function. In Gazzaniga, M. S. (Ed.) *The cognitive neurosciences* (pp. 803-813). Cambridge, MA: MIT Press.
- Shimojo, S., Tanaka, Y., & Watanabe, K. (1996). Stimulus driven facilitation and inhibition of visual information processing in environmental and retinotopic representations of space. *Cognitive Brain Research*, 5, 11-21.
- Smythies, J. R. (1966). *Brain Mechanisms and Behavior*. New York: Academic Press.
- Sokolov, E. N. (1960). Neuronal models and the orienting reflex. In M. A. B. Brazier (Ed.), *The central nervous system and behaviour* (pp. 187-276). Madison, NJ: Madison Printing.
- Solomon, P. R. (1979). Temporal versus spatial information processing theories of the hippocampal function. *Psychological Bulletin*, 86, 1272-1279.
- Sommer, M. A., & Tehovnik, E. J. (1997). Reversible inactivation of macaque frontal eye field. *Experimental Brain Research*, 116, 229-249.
- Stein, B. E., & Meredith, M. A. (1993). *The merging of the senses*. Cambridge, MA: MIT Press.
- Sutton, R. S. & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Tanaka, D. (1973). Effects of selective prefrontal decortication on escape behavior in the monkey. *Brain Research*, 53, 161-173.
- Thach, W. T. (1998). What is the role of the cerebellum in motor learning and cognition? *Trends in Cognitive Science*, 2, 331-337.
- Thompson, R. F. (1986). The neurobiology of learning and memory. *Science*, 233, 941-947.
- Thompson, R. F. (1988). The neural bases of basic associative learning of discrete behavioral responses. *Trends in Neuroscience*, 11, 152-155.
- Thompson, R. F., & Spencer, W. A. (1966). Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychological Review*, 73, 16-43.
- Tipper, S. P., Howard, L. A., & Houghton, G. (1998) Action-based mechanism of attention. *Philos. Trans. R. Soc. London Ser. B* 353, 1385-1393.
- Treisman, A. M. (1988). Features and objects: The fourteenth Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, 40A, 201-237.
- Treisman, A. M. & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97-136.
- Wagner, A. R. (1981). SOP: a model of automatic memory processing in animal behavior. In N. Spear & R. R. Miller (Eds.), *Information processing in animals* (pp. 5-47). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Wang, J., Aigner, T., & Mishkin, M. (1990). Effects of neostriatal lesions on visual habit formation of rhesus monkeys. *Society for Neuroscience*, 16, 617.
- Wang, D. (1995). Habituation. In M. Arbib (Ed.), *Handbook of brain theory and neural networks* (pp. 441-444). Cambridge, MA: MIT Press.
- Wickelgren, W. A. (1979). Chunking and consolidation: A theoretical synthesis of semantic networks, configuring in conditioning, S-R versus cognitive learning, normal forgetting, the amnesic syndrome, and the hippocampal arousal system. *Psychological Review*, 86, 44-60.
- Yeo, C and Hesslow, G. (1998). Cerebellum and conditioned reflexes. *Trends in Cognitive Science*, 2 (9), 322-331.