The Journal of Mind and Behavior
Summer 1991, Volume 12, Number 3
Pages 347–366
ISSN 0271-0137

Toward a Model of Attention and Cognition, Using a Parallel Distributed Processing Approach Part 2: The Sweeping Model

Gregory Christ

University of Ottawa

This paper outlines a theoretical model of attention and cognition using artificial neural network properties, the Sweeping Model. This model includes a sensory memory, a self modifying neural network (capable of forming associations, and of pattern completion), a system applying synchronous sweeps of lateral inhibition across the neural network, and a subsystem to evaluate the body state of the system to control this sweep frequency. After further description, it will be argued that these components can result in behaviours such as learning and performing according to primary and secondary reinforcement rules.

As an extension of a previous review of relevant literature (Christ, 1991), this paper outlines a model of attention, the Sweeping Model. As is evident from reviewing the literature, attention encompasses or involves so many other processes (for example, perception, memory, learning) that it cannot be studied in isolation but must be presented as part of a larger model of cognition. Thus the Sweeping Model takes such an overview.

The use of neural networks of a parallel distributed processing (PDP), or connectionist, type is appropriate because such systems show certain properties that mirror those of the human brain, like content-addressable memory, graceful degradation, spontaneous generalization, and default assignment (see Rumelhart and McClelland, 1986; chapters 1–4).

The Sweeping Model takes a view of cognition much like that described by Martindale (1981), in which any cognitive activity, or experience, is a result of a pattern of neuronal activation in the brain (or any highly interconnected

This work was conducted first at the Cognitive Science Department at Sussex University, England, with the aid of a Commonwealth Scholarship, and was then continued at the School of Psychology at the University of Ottawa, with a Natural Sciences and Engineering Research Council Scholarship. Requests for reprints should be sent to Gregory Christ, School of Psychology, 125 University, University of Ottawa, Ottawa, Ontario, Canada, K1N 6N5.

neural network), with the most active parts representing the more "conscious" activity.

Figure 1 depicts an overall view of the Sweeping Model and some of its main features. Information enters from the environment to the sensory memory (through sensory receptors), with sensory memory for each modality located in distinct areas. Each modality may have a different duration of activity in sensory memory (for example, echoic appears to be longer than iconic; see Christ, 1991, p. 258). The activity in sensory memory is then periodically boosted into a neural network by relatively synchronous sweeps of lateral inhibition (corresponding to the arousal system in the brain). The frequency of these sweeps of lateral inhibition is influenced by an evaluation of positive or aversive body states, that uses information from within the neural network and possibly also information from the sensory memory. The responses of the system (for example, motor activity) should be seen as part of a pattern completion by the network, and would go either out into the environment in the form of an overt behavioural response, or perhaps go to other structures in the system. It should be noted that this process could operate with a neural network divided into a number of similar but relatively distinct modules. There could even exist several of these systems, each with its own lateral inhibition sweep frequency, as subsystems of a larger whole. It should also be noted that the different parts of the model need not be physically separate, as Figure 1 may imply. For example, a sensory memory would best be viewed as units near the edge of a neural network, and not as a separate structure. Below, the various parts of the model will be described in detail to suggest how they operate together as an attentional device.

Description of the Neural Network[1]

As shown in Figure 1, the neural network of this model appears to be the "black box," where all the magic of attentional processes occurs. However, all that is needed for this model is that the network can be an association-forming, pattern completion device. Here, the "black box" only requires the ability to modify the connections between units such that the network could represent associations and patterns from the environmental input, and could complete these learned representations if provided with a significant subset of any part of that same input (that is, content-addressable memory). Findings from work conducted with neural networks indicate that these properties are possible and even emerge somewhat naturally from such models (Hin-

[1]For grammatical simplicity, the Sweeping Model will be discussed in the present tense. This is not meant to imply that the model has already been implemented on a computer; it is still at a conceptual stage.

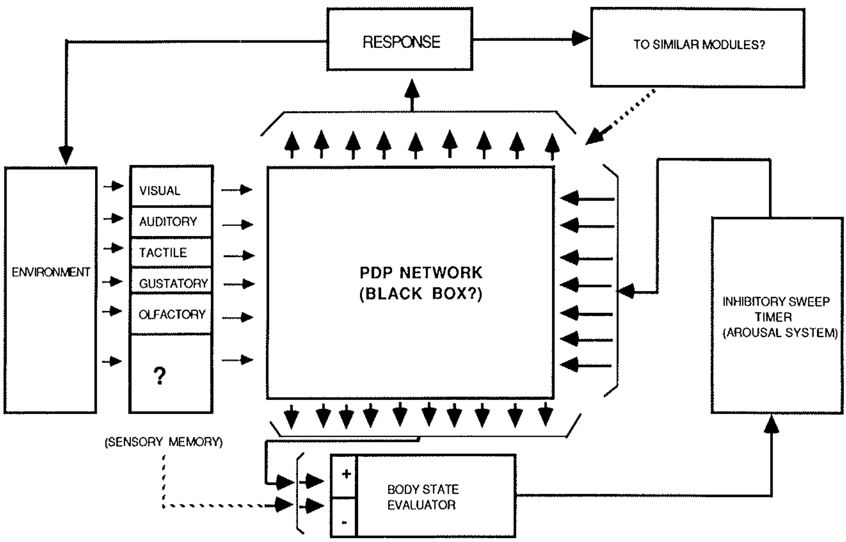

Figure 1. Schematic diagram of the Sweeping Model of attention and cognition. See text for detailed explanation.

ton and Anderson, 1981; Rumelhart and McClelland, 1986). Below, some suggestions for the units and learning rule for implementing this system are made, but it must be stressed that these are only suggestions. The important factors for this model are that the network be able to learn statistical regularities of environmental stimulation, representing these as patterns of activation within the network, and that activation should be able to spread amongst units, in a content-addressable way, to complete these patterns. These properties could be implemented in many ways.

Suggestions for the Units

The units in the neural network, to be both biologically plausible and flexible in behaviour, possess degress of activation, most likely coded as firing rates between a maximum and minimum value. The units have excitatory and inhibitory connections to many other units.[2] The connection strengths are depicted as fairly random at the start; that is, the network is uniform with no preset representations, and then is modified by learning. Units spatially closer together tend to have greater influence on each other (1) to act more like neurons, and (2) so that different areas of the network can more easily operate relatively independently, as in the case of divided attention. This

[2]Exactly how interconnected the units have to be or whether the network would be layered or not will not be examined in this paper.

spatial relation between units serves only as a general starting point that would be changed by learning.

Suggestions for the Learning Rule

The purpose of the learning rule is to bias the network toward completing or recreating the current pattern of activation given a significant subset of the current inputs. A basically Hebbian rule is followed: at the time of update, simultaneously active units have excitatory connections between them strengthened, and inhibitory connections weakened. It might also be added that between active to non-active units[3], inhibitory connections are strengthened and excitative connections are weakened; this aids pattern completion by actively suppressing activity in any units that were non-active during the learning of that pattern. In order to learn some things more quickly (for example, things involving intense stimulation, or very many active inputs), the magnitude of connection-strength changes should be proportional to the activity level in the units involved.

Sensory Memory

Sensory memory serves as a buffer in which incoming environmental stimulation can be summated to form an iconic or echoic image.[4] This image, or pattern of activity, is then allowed deeper into the neural network once for each inhibitory sweep. As previously mentioned, this summation of activity is likely to occur at the edge of the network, rather than in a physically separate structure. The visual system will be used here to illustrate the concept.

The iconic memory consists of a set of units connected to the receptors in the retina, with each iconic memory unit connected to one recpetor, or to a small number of them (a receptive field). Thus, there is a fairly direct mapping of patterns of activity on the retina to patterns of activity in the iconic memory. The receptors are binary, firing once each time the environment stimulates it beyond threshold, but being capable of firing very often (perhaps in the order of 100 or 200 times per second). The iconic units have variable levels of activity (likely coded as firing rates) that decay locally[5] and continuously; that is, each single impulse from a receptor raises the activity level of the iconic unit, and activity then immediately begins to decay gradually,

---

[3]Here, active means firing at or above some rate; and non-active means below that rate or not at all.

[4]There would be a different sensory memory buffer for each sense modality.

[5]Here, "locally" means that each iconic unit acts independently of other iconic units, free from lateral inhibition.

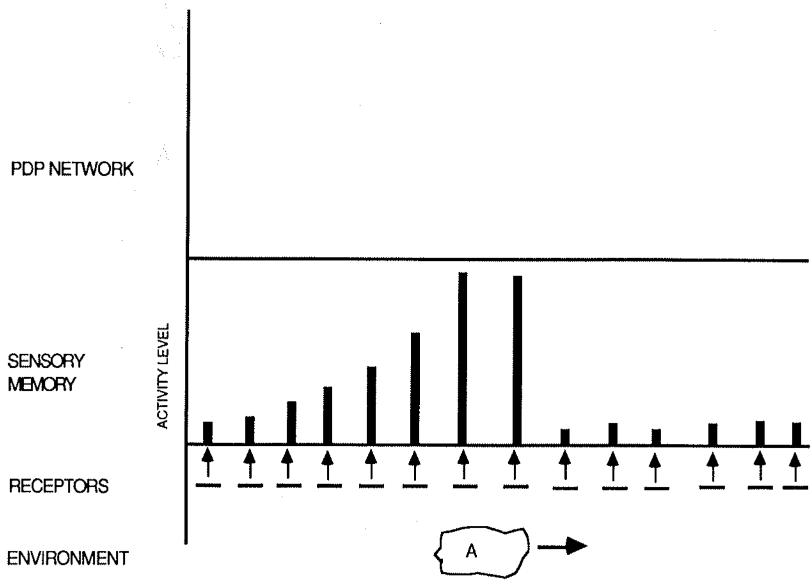

Figure 2. Levels of activity in the sensory memory zone, at the edge of the PDP network, when an object (A) in the environment is moving to the right, stimulating the sensory receptors. Note the trail of activation left by A.

according to some standard function. For vision, this decay function lasts about 250 msec before reaching resting rate. Also, as the activity in the iconic units decreases, any new receptor impulse raises the activity level again.[6] In this way, receptors firing at a high rate (corresponding to some intense and/or regular stimulation in the environment) cause iconic units to become very active, while random or transient stimulation leads to low levels of iconic activity. This buffer of activity is important for filtering out random noise – if receptor firing rates were allowed directly into the network, there would be no way to distinguish random noise from regular patterns, which may all have the same firing rate at any one instant. This would have an effect much like wave averaging the noise out of a signal. The iconic memory would be especially important for perception of motion.

Figure 2 illustrates how the sensory (here iconic) memory unit activity shows differences between random noise and regular stimulation, and also how a moving object leaves a trail of activation. The sensory receptors, which can fire very fast (once for each stimulation beyond threshold), lead to meaningful patterns of activity in the sensory memory units. As Object A moves to the

[6]Iconic activity is raised either by a standard amount per receptor impulse, or according to some function related to the activity level of the iconic unit at the time, so as not to exceed the firing rate of the iconic unit.

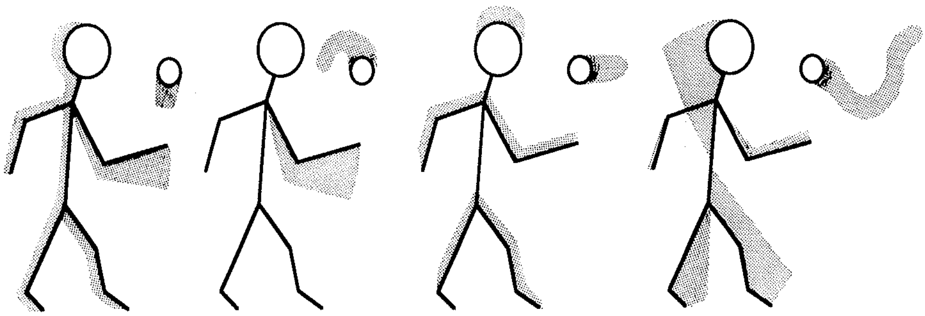

A B C D

Figure 3. Illustration of how different trails of activation in sensory memory can indicate different situations in the environment, even though identical instantaneous information (position of the stick figure) is present in the environment.

right along the row of receptors, those receptors directly over it at this instant are firing at maximum rate, leading to high sensory memory unit activity. The receptors to the left of Object A have just been firing at maximum rate, but are at this instant firing at a random baseline rate (determined by the background noise in the environment); thus their sensory memory units have high activity levels from the previous activity, but have had time to decay, and have decayed more the further away they are from the present position of Object A. All other receptors are firing at some lower rate due to the background stimulation. Thus the sensory memory indicates regularities by high levels of activity, and changes of regularities (such as motion of objects) by trails of activity.

Some Benefits of a Sensory Memory

The sensory memory presents a constantly but gradually changing image of the activity at the receptors, with regularities being most active, and leaving trails of activation when they move. This image then periodically enters the rest of the neural network. The trails of activation are of particular interest, because they at any one instant contain information about changes over time, such as direction of motion (the trail is where the object used to be), speed (longer trails denote faster moving objects), and even object identity.[7]

Figure 3 shows, given the same instantaneous state in the environment, how different trails of activation in sensory memory can indicate very different

[7]Object identity is established as follows: the trail of activity behind moving Object A at time T, leads back to the position A was in at time $T-1$, thus establishing that it is the same stimulus in the environment. Perhaps if Object A moved too quickly to leave such a connecting trail, it might be perceived either as two separate objects spontaneously appearing, or it might change position so far as to leave the visual field. There is very likely a relationship between width of visual field (or receptive field in general) and duration of sensory memory activity.

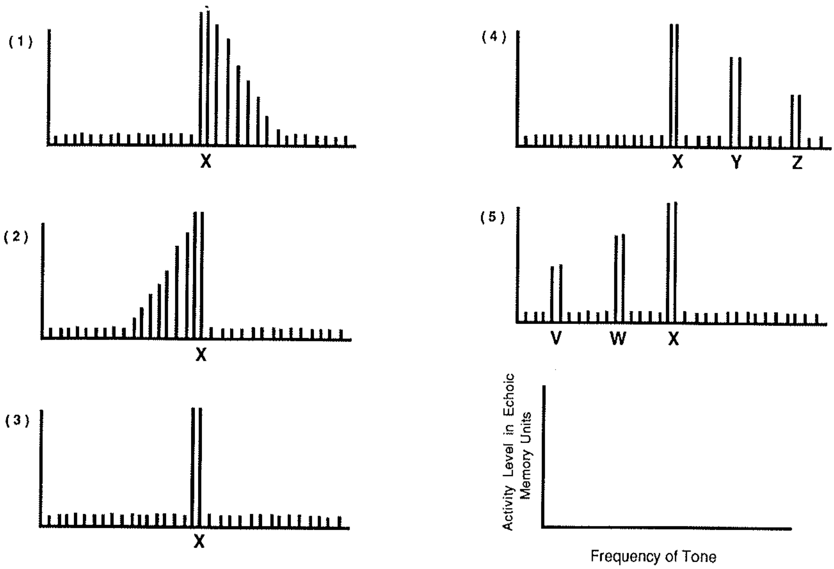

Figure 4. Trails of activation in the auditory sensory (echoic) memory for various inputs. The patterns represent the following situations: (1) one tone gradually decreasing in frequency to frequency X; (2) one tone rising to frequency X; (3) a single tone at frequency X; (4) scale of discrete tones descending to X; (5) scale of tones ascending to X.

situations at the same instant. (A) shows that the figure is moving to the right (trails of activation are to the left of the figure), fairly slowly (trails are short, therefore there was little change over time), and has tossed the ball (faster motion of arm), and the ball is still rising (trail below ball); (B) shows the figure is not moving left or right but has just tossed the ball which is already descending. The sensory memory can indicate curved trajectories without needing several frames to calculate vectors; (C) shows the figure falling and the ball moving left toward it; (D) shows the figure rotating clockwise about its center, and the ball following a curvy trajectory, moving very fast (long, less intense trail).

Figure 4 shows how trails of activation can also indicate "motion" of one sound as opposed to discrete, separate sounds in echoic memory at any one instant. (1) shows the pattern of activation for one tone falling in frequency to frequency X; (2) shows one tone rising to frequency X; and (3) shows a single tone of frequency X. Note that the environment was identical at the instants that (1), (2), and (3) were viewed in the echoic units; (4) shows three tones of identical intensity but different frequency, being presented in order of descending frequency, ending on frequency X. Note how (4) is different from the single decreasing tone in (1); (5) similarly presents three tones in order of ascending frequency, ending on frequency X, to be compared to (2).

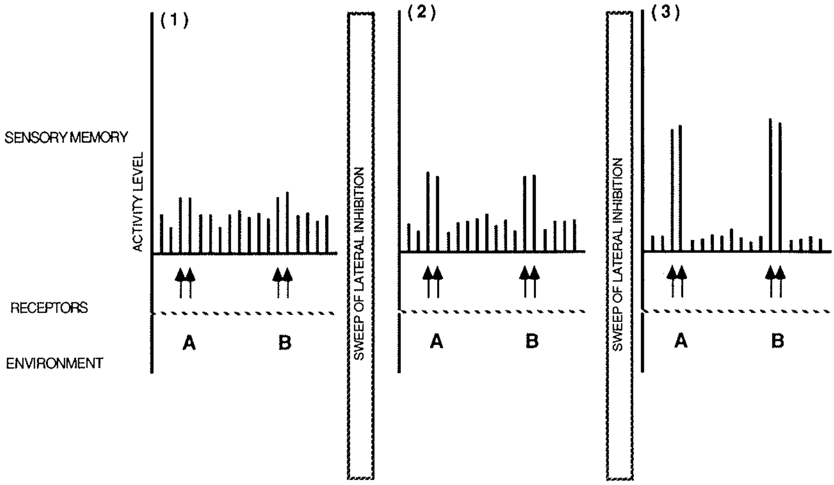

Figure 5. Illustration of how sweeps of lateral inhibition between units in sensory memory could act to diminish noise from an environmental input. Objects A and B are in the environment, with corresponding levels of activity in sensory memory units, at times (1), (2), and (3). Each time is separated by a sweep of lateral inhibition, which exaggerates differences and reduces random noise.

How Activation Spreads (The Filter Mechanism)

Activation would spread deeper into the neural network from the sensory memory of each modality, and then continue to influence the system, via a two-process mechanism.

(1) The network itself is essentially excitatory, with some inhibition between units, and asychronous. Any pattern of activity in the network will tend to spread, and relax toward some resultant state.

(2) Lateral inhibition between units is cyclically applied in a more synchronous sweep across the network (corresponding in the brain to nonspecific input from the arousal system; see Harter, 1967). This lateral inhibition operates separately from other inhibitory connections that are simply parts of learned representations.

Each inhibitory sweep causes activity to become concentrated or focused in the already most active parts of the network. As shown in Figure 5, lateral inhibition refers to how neighbouring units inhibit each other in proportion to how active each one is at the time.[8] This acts to accentuate differences in activity levels, hence greatly decreasing low level random noise.

[8]That is, the most active units can inhibit surrounding units much more than less active units. Also, the more a unit inhibits the surrounding ones, the less inhibition there is upon it, so it becomes even more active and can inhibit other units further.

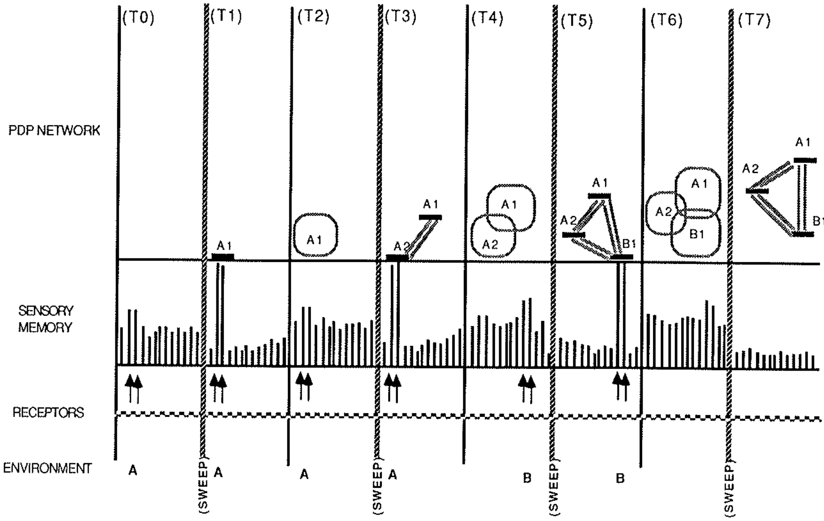

Figure 6. Illustration of how a sequence of sweeps of lateral inhibition leads to activity from environmental inputs entering, then propagating through the neural network. Object A is in the environment for time T0 to T3, and Object B is present for T4 and T5. Activation in the network is shown as black bars when concentrated following a sweep of lateral inhibition. Dotted boundaries show how far activation has spread before the next sweep. Note that this process also reduces noise from the environment.

In Figure 5, (1) shows activity levels from a noisy environment containing Objects A and B; (2) shows how after a sweep of lateral inhibition, the somewhat higher levels of activity in the sensory memory units related to A and B have inhibited their surroundings, with the surroundings now having less inhibitory effect; and (3) shows that after the next sweep of lateral inhibition, the differences are exaggerated even more. Note that any one unit can inhibit only a small surrounding area of units, so areas for both A and B can be quite active without interfering with each other much. This suppression of random noise leads to a clearer pattern of A and B being passed into the network, and is not possible without the sensory memory.

Another aspect important to this model is that the learning rule is applied to update connection strength once per sweep of lateral inhibition. And also one sweep corresponds to one psychological moment (PM).[9]

To describe how the two process mechanism works, Figure 6 depicts a succession of states of the system, including the environment, at times T0 to

[9]The "psychological moment" refers to the overall pattern of activation in the network that includes the latest bottom-up information from the environment. New information from the environment enters in a succession of chunks, or rapid attentional integrations, and is added to the current state of activation in the network to form a sequence of PM's. This concept is discussed in Christ (1991).

T7. (T0) shows that Object A is in the environment, giving noisy levels of activity in sensory memory; (T1) immediately follows a sweep of lateral inhibition, and shows that random noise has been reduced from the pattern for A in the sensory memory. In fact, A has become so active that the activity [A1] has spread to units actually in the neural network; (T2) shows that environmental noise is again creeping into the sensory memory pattern for A, and the activity A1 has spread as far as the dotted boundary just prior to the next sweep; in (T3), immediately after the sweep, more bottom-up activity [A2] is passed into the system, from the sensory memory, about Object A being in the environment. Lateral inhibition also caused the spreading activity of A1 to become concentrated in whatever parts of the network were most active at the time (here depicted as one point, A1). The dotted lines between A1 and A2 represent the connections between these simultaneously active areas of units; (T4) shows how A1 and A2 have spread and relaxed together toward a resultant state in the network. Also, Object A has disappeared and Object B has appeared in the environment, giving rise to changes in the sensory memory activity; (T5) immediately follows the next sweep, which has boosted activity about B from sensory memory into the network [B1], as well as concentrating A1 and A2 into new, and somewhat intermingled areas; (T6) shows how A1, A2, and B1 relax together and move into the system. There is now only a random environment; (T7) after this sweep, no new bottom-up activity enters (unless one part of the noise happened by chance to be much greater than surrounding noise), and activation A1, A2, and B1 are concentrated and blend together, to start spreading from this new configuration.

In this way, activation can propagate throughout the network, while being influenced by and integrated with incoming bottom-up activity. At this point, the system is a filter for accepting only the most active stimulation[10], and supressing random noise. This is desirable because intense stimuli tend to catch the focus of attention.

How the System Chooses Relevant Over Only Intense Stimuli

The system can choose relevant stimuli over merely the most intense stimuli due to top-down influences from the state of activation already in the network adding to the bottom-up sensory activation. Thus prior activation can select from the environmental inputs. Figure 7 illustrates this point.

Figure 7 shows at (T0) that Objects A, B, C, and D are in the environment, and each is causing an identical amount of bottom-up activity in the

[10]Stimuli could highly activate the sensory memory by being the most consistently present or the most intense stimuli in the environment.

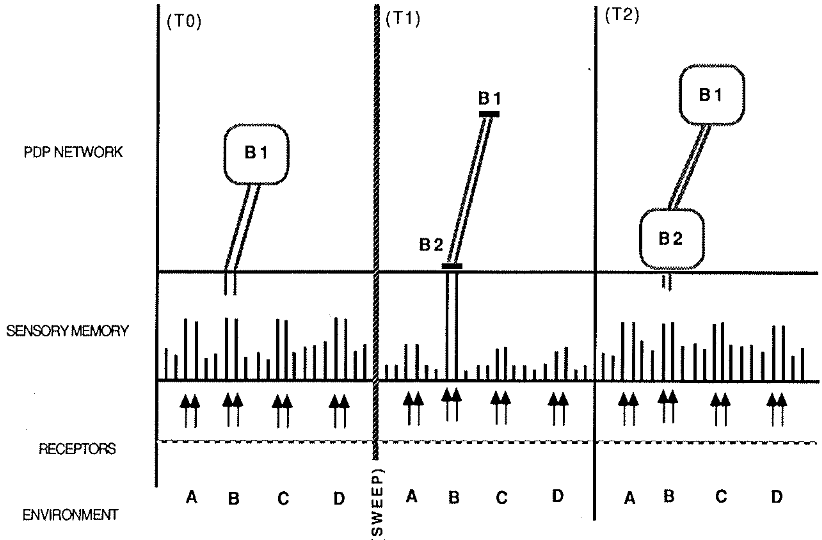

Figure 7. Illustration of how prior activation in the neural network can (top-down) add to and thus select from equally active bottom-up inputs. Activity B1, already in the network, adds to sensory memory activity from Object B in the environment. This boosts B activity above A, C, and D levels, allowing B activity to inhibit the others during the next sweep, and be the only information passed into the network.

sensory memory. B1 is the prior state of activation in the neural network, which has some relation to Object B (either by content, or perhaps to that area of the sensory memory, for example, a place on the retina), indicated by the dotted lines connecting B1 to the B region in the sensory memory. Some activation would spread top-down from B1 to increase the bottom-up B activity over A, C, and D activity in sensory memory; (T1) shows that after the next sweep of lateral inhibition, the slight increase in B activation is enhanced, and B becomes clearly stronger than A, C, or D. Thus, only B activity [B2] is able to enter the neural network and influence its activity; (T2) shows how bottom-up activity tends to exert more influence (that is, noise and new stimulation would make the activity levels between units much more even) until the next sweep exaggerates differences again. In this model, bottom-up influences tend to be stronger than top-down, and thus constrain the spread of activation. For example, a steady environmental stimulation tends to overpower or inhibit the weaker influence of top-down prior activation.

To further describe the spread of activaton resulting from this two process mechanism, activation spreads, after each sweep, along the strongest connections, in a content-addressable way. Thus, each subsequent state will be related in content to the previous one, resulting in a flow between states, with new environmental input blending into and influencing (relaxing with) the prior activation toward a resultant state. The most active parts of the network, which are all likely to be connected to each other, make up the current experience of the system (the current PM).

In more PDP terms, each sweep can be seen to (1) input new environmental information, and (2) raise the computational temperature of the system by concentrating activity into certain areas. Between sweeps, the system relaxes toward a resultant pattern of activation, but is interrupted by the next sweep (bringing new input as well) before reaching a final state. The result is the flow between related patterns of activation. The input vector to the system involves both the environmental input *and* the state of prior activation (including things like motor output). Here, relaxation is a pattern completion, and the system, being probabilistic, is likely to complete the most learned patterns for any input (that is, patterns that in the past have been updated the most often by the learning rule).

This proposed mechanism, so far, explains how some relevant stimuli are selected[11], and also, generally, how a flow of related patterns of activation (corresponding to the subjective experience of a flow of related PM's, thoughts, or mental states) in a neural network can be achieved. However, it still remains to be explained how this process, if allowed to run freely in an environment, comes to result in intelligent behaviours like those seen in animals or humans. I propose that this is largely due to the variable sweep frequency of the lateral inhibition, and just how this frequency is controlled by the environment through the prewired knowledge of positive and aversive body states.

Effects of Lateral Inhibition Sweep Frequency

(1) A faster sweep frequency results in more concentrated or focused activity in the network. Activity patterns have less time between sweeps to spread or relax out. As a result of so much lateral inhibition, only a few areas become very active (a narrow focus), rather than many widespread patterns being somewhat active (a wide focus); this corresponds well to the relationship of arousal to scope of attention in humans (see Martindale, 1981). In this way, sweep frequency can be seen as a limit to how widely activation can spread during any one PM, and hence what representations could be activated during that sweep frequency. Figure 8 illustrates how a slow sweep frequency allows activity to spread to simultaneously activate areas X and Y, while a faster sweep frequency does not.

---

[11]The system will select for intense stimuli (bottom-up), and stimuli related to prior activation (top-down).

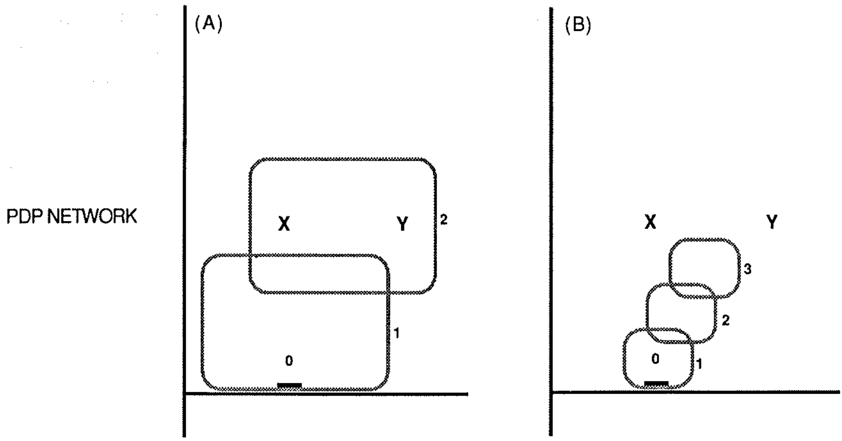

Figure 8. Illustration of how a longer sweep frequency for lateral inhibition, shown in (A), allows wider spreading of activation such that it can spread to simultaneously activate the separate areas X and Y in the neural network. This is not possible with the shorter sweep frequency in (B).

In Figure 8(A), O is an area of activation immediately following a sweep of lateral inhibition. O spreads as far as the dotted square 1 before the second sweep. Dotted square 2 shows how far activation was able to spread after the second sweep but before the third. Note that the spread was great enough to simultaneously activate areas X and Y. (B) shows the same situation but with a much faster sweep frequency, only allowing a small amount of spreading between sweeps. The spreading during three sweeps is shown. Note that X and Y are too far apart to ever be simultaneously reached by activity spreading at this sweep frequency (although they could be activated with specific bottom-up activity leading to X and Y by separate routes).

Thus, it appears that activation has to spread to approximately the same distances for reactivation (remembering) as were present when the representation was initally learned. So having the same learning and recall sweep frequencies would seem to enhance the chances of reactivating a representation. This agrees well with research on state dependent learning in humans and animals, which demonstrates that information is easier to recall if the state during learning is the same as that during the attempt to recall. State would be largely determined by arousal level, which here corresponds to the lateral inhibition sweep frequency.

(2) A faster sweep frequency also results in more frequent sampling of the environment. Upon each sweep, differences in activity levels of units in the sensory memory are exaggerated, allowing the most active parts to spread deeper into the neural network. This also serves to constrain the flow of activity in the network, due to the large amount of bottom-up activation influencing the network.

(3) A faster sweep frequency means more connection strength updates, hence more learning of environments associated with high sweep frequencies. Thus, any patterns of activity in the network (from bottom-up or top-down sources) during a high sweep frequency are learned more than patterns present during a low sweep frequency. Also, the parts of patterns most learned (that is, most updated) are the parts in which activity became concentrated or focused.

Control of the Sweep Frequency

Innate Knowledge of Positive and Aversive Stimulation (Body States)

Some prewired or innate knowledge has to be contained in the system in order to control the sweep frequency in response to certain body states related to the environment. This prewiring involves knowledge of certain positive or aversive body states, as determined by some very general sensory input parameters (for example, very intense stimulation). This kind of prewiring or innate knowledge is plausible – neonates can apparently distinguish pleasant (positive) from painful (aversive) from neutral body states. It is likely that the positive and negative evaluations are accomplished by relatively separate systems, as evidenced by the existence of different response patterns to positive and aversive reinforcement (Flaherty, 1985). The existence of drugs (for example antianxiety drugs like benzodiazepines) that appear to affect responses to only aversive stimuli, but not positive stimuli in humans (Gray, 1982), further supports the independence of the two evaluation systems. However, both positive and aversive body state evaluations directly influence the sweep frequency, by temporarily increasing it. The sweep rate then gradually returns to baseline, according to some decay function which would be different for the positive and aversive evaluators.

Shifts in Attention

In order to achieve shifts of attention, it is necessary to first defocus activity before reconcentrating it with the increased sweep rate. This occurs as a result of the positive or negative evaluation of some stimulus and involves a diffuse burst of excitation throughout the neural network[12], uniformly raising the activity level in many units to blur the current pattern. Then activity is refocused, possibly in different patterns, by the increased sweep rate. Thus, each new positive or negative evaluation involves a non-specific burst of ex-

[12]This burst of activation corresponds to ponto-geniculo-occipital [PGO] spikes in human beings that are associated with alerting responses to new environmental inputs (see Morrison and Reiner, 1985).

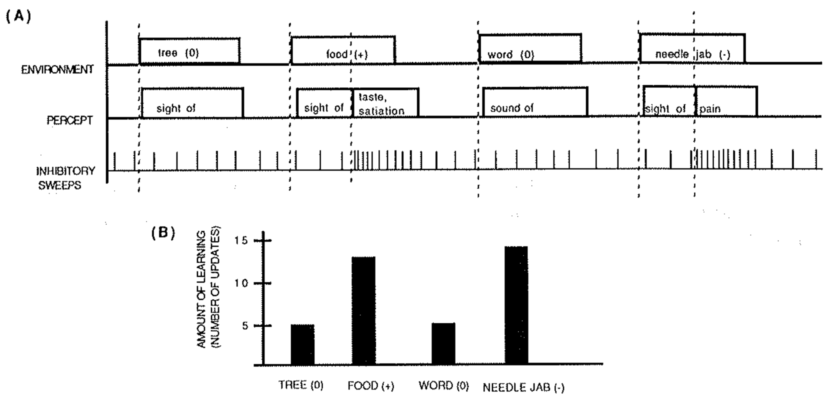

Figure 9. (A) Illustrates how, in the proposed Sweeping Model, environmental stimuli that influence body state would change the lateral inhibition sweep frequency, and thus the number of connection strength updates for the corresponding pattern of activation (here referred to as "percepts") in the neural network at the time. (B) shows how much learning (number of connection-strength updates) occurs for each stimulus.

citation from an internal source, followed by a temporarily increased sweep rate. This results in the observed orienting response (disruption of ongoing behaviour and shifting of focus of attention).

Consequences of the Sweeping Model

Primary Reinforcement

This overall mechanism results in behaviours following well-established reinforcement schedules observed in animals and humans. Any activity that happens to be in the network at the time of positive or aversive stimulation is subject to the increased sweep frequency, which then concentrates or focuses the activity; furthermore, the increased number of connection strength adjustments results in more learning of any patterns of activity (representations) that were frequently concurrent with such stimulation. So, with bottom-up activity, environmental contingencies would be learned. The types of patterns learned, at first, would be ones easily described as related to reward and punishment in terms of body states, and animals and children appear to learn things much in this way. Figure 9 illustrates how different amounts of learning of the environment (from bottom-up activation) result from this evaluation system.

Figure 9(A) shows three time lines indicating events in the Environment, and simultaneous active representations (subjective Percepts) in the system, along with the lateral inhibition Sweeps. In the environment, "+" indicates that the object gives rise to a positive body state, "−" to an aversive body state, and "0" to a neutral body state. Starting at the left of the diagram, a Tree is presented in the environment, which gives rise to a percept (pattern of activation) for a tree. The Sweep frequency was previously at baseline, and remains at that frequency since a tree has a neutral evaluation (0) and thus only a small amount of learning occurs (few connection strength changes). Next, Food is presented, giving rise first to just a visual percept of food, which is initially neutral the first time this food was ever seen. However, if eaten, the percepts of taste and satiation may arise and due to their positive (+) evaluation (they change the body state), they temporarily increase the sweep frequency, as shown by the increased number of sweeps at the time of taste and satiation. This results in an increased amount of learning (more connection strength updates) for this state of activation in the network. Note that the sweep frequency gradually returns to baseline. Next a Word is presented for the first time, which gives rise to an auditory percept or the visual percept of the printed word, both of which have a neutral (0) evaluation. Thus the sweep frequency stays at baseline and little learning occurs. Next is presented a Needle Jab, which starts out as the neutral percept at the sight of a needle. However, then follows the strongly aversive (−) percept of pain (a change in body state), that results in an increased sweep frequency. The sweep frequency remains increased until the pain-causing stimulus is removed and the body state returns to neutral, after which the sweep frequency returns toward baseline. There would be many connection strength updates, and hence more learning of this state. Figure 9(B) shows the amount of learning, in terms of the number of connection strength updates, for each of the four situations in the environment in Figure 9(A). Note that the neutral Tree and Word are learned much less than the positive Food or aversive Needle Jab.

The learning illustrated in Figure 9 involves primary reinforcers: things in the environment that themselves physically change the body state of the system. Any model of attention or cognition has to account for primary reinforcement schedules of learning. But this still leaves the system somwehat rigidly dependent upon the evironment for all responding. More cognitively based, environment independent behaviours – that are less easily explained by traditional behaviourism (such as language) – must also be taken into account.

Secondary or Conditioned Reinforcement (The Development of Symbols)

Due to the highly interconnected nature of the network, units located anywhere in the network have connections to the units in the positive and aversive evaluation systems. Thus, the normal learning rule could strengthen connections between certain representations in the network and the evalua-

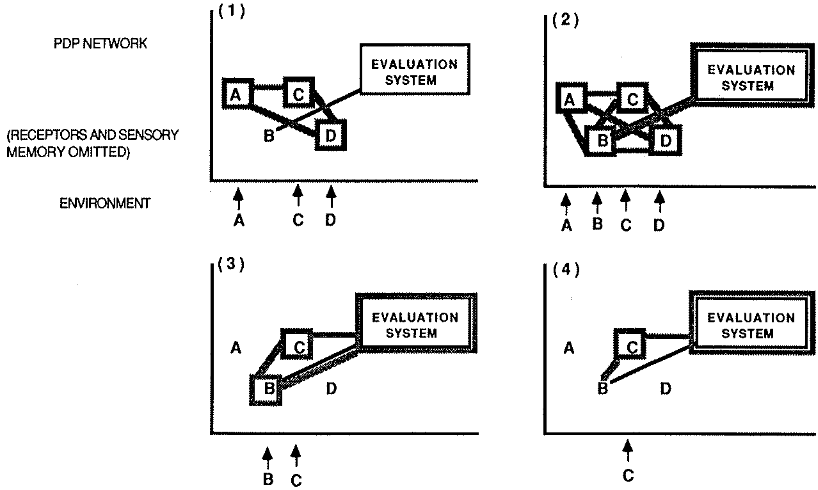

Figure 10. Illustration of how an initially neutral stimulus (C) can, through learning, come to act as a conditioned reinforcer or punisher (that is, directly influence lateral inhibition sweep frequency), and thus direct further learning.

tion system itself, provided that they were often simultaneously active. The result is that these once-neutral representations come to directly influence the evaluation system, and hence the sweep frequency; these representations thus become secondary reinforcers, and by being able to control the sweep frequency directly, without requiring a change in body state (as primary reinforcers do), they can direct further learning with a very flexible relationship to the environment.

Figure 10 shows how a neutral representation or pattern of activation [C] can become a secondary reinforcer and directly influence the Evaluation System, which changes the sweep frequency. The areas A, B, C, and D denote various representations in the network: A for the sight of a tree; B for the experience of pain; C for the sight of a needle; D for the sound of a word. These letters (A, B, C, D) are written across the bottom of each diagram to indicate when each is present in the environment to activate (bottom-up) the corresponding areas in the network. Note that B (the experience of pain) is a primary reinforcer,[13] reflecting an aversive change in body state, with

[13]Here the term "reinforcer" is used, for simplicity, to illustrate the concept that certain stimuli can influence subsequent behaviour, although "punisher" could be substituted. With the terminology used here (see Skinner, 1953; p. 73, p. 185), "positive" denotes the presentation of a stimulus, "negative" denotes the removal of a stimulus, a "reinforcement" means increasing the probability of a behaviour, and "punishment" means decreasing the probability of a behaviour. Thus any aversive stimulus could be a positive punisher (presenting a stimulus to decrease the frequency of a behaviour) or a negative reinforcer (removing a stimulis to increase the frequency of a behaviour), depending upon what behaviour is being observed.

prewired connections (shown in the diagram as a solid line) to the Evaluation System. Diagram (1) shows the environment activating areas A, C, and D (seeing a tree, a needle, and hearing a word in the environment). Note that they are all connected to each other, so the resultant pattern of activation forms one overall representation (ACD). Note also that the Evaluation System is not activated, so relatively few learning updates will occur; thus this representation (ACD) will not be learned to a great extent. Diagram (2) shows A, B, C, and D being activated by the environment (seeing a tree, a needle, hearing a word, and feeling pain). Again, all the activated areas are connected to all the others, making up an overall representation (ABCD). However, once B, the primary reinforcer, is activated, the activation immediately spreads along the prewired connections to activate the Evaluation System to temporarily increase the sweep frequency. This increased sweep frequency results in many more connection strength updates than in Diagram (1), and hence more learning of ABCD, than of ACD. Note also that in Diagram (2) connections between A, C, and D to the Evaluation System itself are also strengthened on each update. Diagram (3) shows areas B and C being activated by the environment (the sight of a needle, and experience of pain), these two being very likely to go together in a needle jab situation. Here, B activates, via the prewiring, the Evaluation System, so there will be many updates strengthening connections between B and C, and C to the Evaluation System. Note that in the situations for both Diagrams (2) and (3), the connections from C to the Evaluation System are strengthened many times, and more than from A or D to the Evaluation System. The result is shown in Diagram (4), in which only C is activated by the environment (the sight of a needle). The activation at C will spread along the strongest connections, and (if the connections had been strengthened enough) alone could spread to activate the Evaluation System and change the sweep frequency. This means that C would have become a secondary reinforcer, and then could influence the learning of things in the environment associated with needles, even in the absence of a change in body state (pain). Note again how this system acts like an averager, strengthening connections between frequently associated units and not between units rarely paired, which corresponds to learning contingencies in the environment, that is, bottom-up activity.

These secondary, or conditioned, reinforcers can then in a similar way lead to the development of tertiary reinforcers, and so on, until there is little or no correlation between behaviour and things in the environment that alter body state. In the above example, this makes C a symbol, perhaps predicting pain, but in any case having similar effects as pain upon the sweep frequency and learning. It is difficult if not impossible to trace conditioned reinforcers (symbols) back to their original body state relationships to the environment, and complex, relatively environment-independent behaviour patterns result.

Symbolic reinforcers might eventually become strong enough to override other symbolic reinforcers and even primary reinforcers. The ability to develop these strong symbolic reinforcers is a major difference between humans and lower animals, and adds flexibility and complexity to response patterns, avoiding mechanical behaviours tied directly to cues in the environment. Unfortunately, this flexibility disrupts any environment-to-behaviour mapping, which makes studying such systems complicated, and tends to obscure the mechanical nature of the learning and behaviour in this model. This may even make mechanical models seem less plausible as an approach than more global theories that appeal to "common sense" – that often end up with homunculi or empty, magical black boxes.

Conclusion

The Sweeping Model of attention and cognition, outlined above, puts forth a mechanical system that could account for many types of learning and behaviour. Although leading toward a computer simulation, this model should be viewed more as a speculative psychological theory of attention and cognition. Beyond simply accounting for behaviourist reinforcement rules and the more cognitive symbolic activities, this model has implications for perception, memory, and also sleep and dreaming, to be examined elsewhere.

References

Christ, G. (1991). Toward a model of attention and cognition using a parallel distributed processing approach. Part 1: Background. *Journal of Mind and Behavior, 12*, 247–262.

Flaherty, C.F. (1985). *Animal learning and cognition.* New York: Alfred A. Knopf.

Gray, J.A. (1982). Precis of: The neuropsychology of anxiety: An enquiry into the functions of the septo-hippocampal system. *Behavioral and Brain Sciences, 5*, 469–534.

Harter, M.R. (1967). Excitability cycles and cortical scanning: A review of two hypotheses of central intermittency in perception. *Psychological Bulletin*, 68(1), 47–58.

Hinton, G.E., and Anderson, J.A. (1981). *Parallel models of associative memory.* New Jersey: Lawrence Erlbaum Associates.

Martindale, C. (1981). *Cognition and consciousness.* Homewood, Illinois: Dorsey Press.

Morrison, A.R., and Reiner, P.B. (1985). A dissection of paradoxical sleep. In D.J. McGinty, R. Drucker-Colin, A. Morrison, and P. Parmeggiani (Eds.), *Brain mechanisms of sleep* (pp. 97–110). New York: Raven Press.

Rumelhart, D.E., and McClelland, J.L. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition.* Cambridge, Massachusetts: MIT Press.

Skinner, B.F. (1953). *Science and human behavior.* New York: Free Press.