Synfire Chains and Catastrophic Interference

Jacques P. Sougné and Robert M. French {J.Sougne, rfrench}@ulg.ac.be
Department of Psychology, University of Liège
4000 Liège, Belgium

Abstract

The brain must be capable of achieving extraordinarily precise sub-millisecond timing with imprecise neural hardware. We discuss how this might be possible using synfire chains (Abeles, 1991) and present a synfire chain learning algorithm for a sparsely-distributed network of spiking neurons (Sougné, 1999). Surprisingly, we show that this learning is not subject to catastrophic interference, a problem that plagues many standard connectionist networks. We show that the forgetting of synfire chains in this type of network closely resembles the classic forgetting pattern described by Barnes & Underwood (1959).

Introduction

A professional pitcher can send a baseball hurtling towards a batter at speeds approaching 100 miles an hour. In a mere 2 milliseconds the ball moves three inches, more than the width of a baseball bat. Given the long chain of neurons that must fire sequentially with incredible precision in order for the bat to connect with the ball, how could anyone ever hit a baseball? Consider gymnastics. The landing of a skilled gymnast dismounting from the high bar with a triple somersault is completely determined by the millisecond-precise instant he releases the bar. How is it possible to achieve the extraordinary timing accuracy necessary to consistently hit this beautiful landing correctly?

The problem is to find a way to make imprecise neurons act in an extremely precise manner. Nature has clearly found a way to circumvent the imprecision of individual neuron firings. The solution seems to rely on the presence of large *populations* of interacting neurons. In this paper we will discuss a mechanism for achieving precise timing, *synfire chains* (Abeles, 1991), that has received considerable empirical support. We will consider how the brain might learn these synfire chains and will present a neurobiologically plausibly computer simulation of synfire chain learning (see Sougné, 2001). Most importantly, we will show that the problem of catastrophic interference, a problem that plagues many types of neural networks (see French, 1999, for a review), does not seem to be a problem for synfire chains implemented in a sparsely-distributed network of spiking neurons. We simulate the classic forgetting experiment of Barnes & Underwood (1959) on a network designed to learn synfire chains and show that forgetting of information encoded in these simulated synfire chains very closely resembles the forgetting patterns of humans, as demonstrated by Barnes & Underwood. This leads to the prediction that real neural synfire chains will be forgotten gradually, rather than catastrophically.

Synfire Chains

Empirical data demonstrate the existence of very precise temporal behavior in neuron firings. For example, researchers have recorded spike timing of different cortical cells in monkeys (Abeles, 1991, Prut & al, 1998) and have observed the following stimulus-dependent pattern: when an initial neuron, A, fired, a second neuron, B, would fire 151ms later, followed by a third neuron, C, that would fire 289ms later *with a precision across trials of 1 ms*! Intervals of this duration require dozens of transmission delays from neuron A to neuron *C*. One of the major hypotheses about how this phenomenon could occur involves so-called “synfire chains.” Abeles (1991). The other hypothesis is based on an increase in a population rate, which builds excitation in a downstream population, which, in turn, increases its firing rate, etc. (see Shadlen & Newsome, 1994). According to Abeles’ hypothesis, since cortical synapses are relatively weak, many inputs to cells must arrive at the same time for them to fire. Consequently, each step in the synfire chain requires a significant pool of neurons whose simultaneous firing raises the potential of the next pool of neurons to allow them to fire. Recent experiments (Prut & al, 1998) indicate that these precise temporal firing sequences correlate more to behavior than to rate modulation and do not seem to be a byproduct of rate modulation. This would seem to buttress the synfire chain hypothesis.

Previous work on synfire chain learning has focused on how they can develop from a chaotic net with an unsupervised Hebbian learning rule (Bienenstock, 1995; Hertz & Prügel-Bennet, 1996). These studies involved an external stimulus which makes a large pool of neurons fire simultaneously at a particular instant. Subsequently, a sequence of successive large pools of simultaneous neuron firings occurs, produced by the random connection weights of the network. A given neuron will only fire if a large enough number of its presynaptic neurons provoke an increase of postsynaptic potential *simultaneously*. Connections are modified by a Hebbian learning rule. After learning, when the previously learned stimulus is presented again, the same chain fire, thereby constituting a synfire chain. These studies show that these chains are stable, noise tolerant and that one network can store many different chains. Formal analysis showed that there is a relation between network size and the length of learnable synfire chains (Bienenstock, 1995), and that the recall speed should be faster than the training speed of the sequence (Sterratt, 1999).

In previous modeling work (Sougné, 2001), it has been shown how a synfire chain can develop, thereby linking two pools of neuron firings caused by two sequential external stimuli. After learning, the first external stimulus will reactivate the stored synfire chain. It was also shown that synfire chain learning depends on the size of the network, the presence of long term depression (LTD) and the sparseness of connections. It turns out, surprisingly, that these synfire chains are not subject to catastrophic interference.

Catastrophic Interference

Gradual forgetting is one of the fundamental facts of cognition, which means that plausible models of human cognition must exhibit progressive forgetting of old information as new information is acquired. Only rarely does new learning in natural cognitive systems *completely* (or “catastrophically”) interfere with previously learned information (see, for example, French & Ferrara, 1999). However, it turns out that for a very large class of commonly used connectionist models — those with a single set of shared (or partially shared) multiplicative weights (and most notably, standard feedforward backpropagation networks) — learning new information can quite easily completely destroy all traces of previously learned information (McCloskey & Cohen, 1989; Ratcliff, 1990). In fact, the very features that give these connectionist models of memory their much-touted abilities to generalize, to function in the presence of degraded input, etc., are the root cause of catastrophic forgetting (See French, 1999, for a review of research on catastrophic interference).

Catastrophic interference is a radical manifestation of a more general problem for connectionist models of memory — in fact, for *any* model of memory —, the so-called “stability-plasticity” problem (Grossberg, 1982). The problem is how to design a system that is simultaneously sensitive to, but not radically disrupted by, new input.

A number of ways have been proposed to avoid the problem of catastrophic interference in connectionist networks. In the connectionist network that is our brain, McClelland, McNaughton & O’Reilly (1995) proposed that the dual memory system consisting of our hippocampus and neocortex evolved, at least in part, in order to overcome the problem of catastrophic interference (see also French, 1997).

In what follows, however, we hope to show that, rather unexpectedly, there is no catastrophic interference in the implemented network of old information from newly learned precise firing sequence. We will begin by considering a now classic experiment on human forgetting by Barnes & Underwood (1959). We will then show that when this experiment is simulated for synfire chains in a sparsely distributed network of spiking neurons, the forgetting curves observed during new learning are largely the same as those observed in Barnes & Underwood.

Forgetting Caused by New Learning

Barnes & Underwood (1959) conducted a series of experiments that measured the extent of retroactive interference in human learning. We will consider two of their experiments in this paper. In the first, subjects first learned a set of paired associates (A-B) consisting of a nonsense syllable and an adjective (e.g. *dax* paired with *regal*, etc.) and then were asked to learn a new set of paired associates (A-C) consisting of the same nonsense syllables associated with a new set of adjectives (e.g. *dax* paired with *dirty*, etc.). (This was called the A-B/A-C paradigm.) The forgetting curve for the A-B associate pairs produced by interference from the learning of the new A-C pairs was relatively gradual (Figure 1a).

In a second experiment, participants first learned a list of paired associates A-B, as above. Then they were asked to learn a series of paired associate where the second word was very semantically close to the original I word in the A-B pairs. They called this paradigm the A-B/A-B′ paradigm. See Figure 1b for the results of this second experiment.

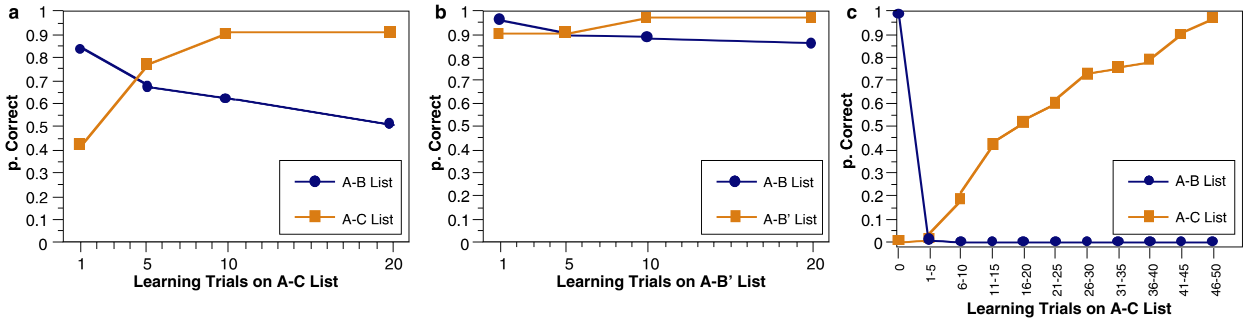

Figure 1, a: Gradual forgetting of previously learned information in Barnes & Underwood’s (1959) A-B/A-C. b: Results of the same experiment when the C list closely resembles the original B list (A-B/A-B′ paradigm). c: McCloskey & Cohen’s (1989) results showing the network’s catastrophic forgetting A-B/A-C paradigm.

When connectionist networks began to become widely used as models of human memory, McCloskey & Cohen (1989) used the Barnes & Underwood A-B/A-C paradigm to test forgetting in these networks. It came as a considerable surprise to most researchers in the field that, at least under certain circumstances, McCloskey & Cohen were able to show that forgetting in a standard backpropagation network was anything but gradual. In one set of experiments, for example, a standard backpropagation network thoroughly learned a set of "one's addition facts" (i.e., the 17 sums 1+1 through 9+1 and 1+2 through 1+9). Then the network learned the 17 "two's addition facts" (i.e., 2+1 through 2+9 and 1+2 through 9+2). Recall performance on the originally learned one's facts plummeted as soon as the network began learning the new two's addition facts. Within 1-5 two's learning trials, the number of correct responses on the one's facts had dropped from 100% to 20%. In five more learning trials, the one's knowledge was at 1%, by 15 trials, no correct answers from the previous one's problems could be produced by the network. The network had "catastrophically" forgotten its one's sums. (See Figure 1c).

Networks of Spiking Neurons

In a network of spiking neurons (Maass & Bishop, 1999), nodes can be in two different states: they can fire (on), or they can be at rest (off). A node fires at a precise moment and transmits activation to other connected nodes with some time course. When a node activation or potential $V_i^{(t)}$ reaches a threshold, it emits a spike. After firing, the potential is reset to some resting value V_r. Inputs increase the node potential, but some part of the node potential is lost at each time step. Spiking neuron models, and in particular, INFERNET, the network discussed here, use a quite realistic post synaptic potential (PSP) function.

INFERNET is not a fully connected network; its structure is organized by clusters of nodes which constitute subnets. Each subnet is fully connected. From each node of a subnet there is a connection to every other node within that subnet. Some subnet nodes have connections to external subnet nodes. This not only reduces the computational demands of the program, but also better corresponds to the actual organization of the brain.

Two variables affect each connection: weight and delay. Each weight corresponds to the synaptic strength between a presynaptic and postsynaptic cell. The weight between a presynaptic node *j* and a postsynaptic node *i* is designated by w_{ij}. Noise is added to this value and the resulting noisy connection is denoted by $\hat{w}_{ij}$. The delay *d* of a connection determines when the effect of the presynaptic node firing will be maximum on the postsynaptic node. There is also a noise factor on the delay. This delay corresponds to the axonal, synaptic and dendritic delays of real neurons.

A signal, whether excitatory or inhibitory, will be affected by a leakage factor. When the signal has reached its maximum, at each following step of 1 ms, the signal will be divided by 2. Delays and leakage factors define the Post Synaptic Potential function $\varepsilon_{ij}(x)$:

$$\varepsilon_{ij}(x) = \frac{1}{2^x} H(x) \qquad (1)$$

where: $H(x) = \begin{Bmatrix} 1 \text{ if } x \geq 0 \\ 0 \text{ if } x < 0 \end{Bmatrix}$

and x is the difference between the time t, the time of the presynaptic node firing, and the noisy delay on the connection: $x = t - t_j^{(f)} - d$.

When a node potential V_i reaches a threshold θ, it emits a spike. Thereafter, the potential is reset to its resting value. After emitting a spike, a node enters a refractory period. This corresponds to the membrane resistance of real neurons which increases after a spike. In INFERNET, the refractory state of node *i* depends only on the last spike of the node i: $t_i^{(f)}$. A value dependent on the refractory state is subtracted from the node state value V_i. This value is denoted by $\eta_i(u)$, where u is the difference between the current time t and the time of the last spike of node i: $u = t - t_i^{(f)}$, and where a and b are constants

$$\eta_i(u) = \left[1 + e^{-\left(u^a - b\right)}\right] \vartheta(u)\theta \qquad (2)$$

where: $\vartheta(u) = \begin{Bmatrix} +\infty, & \text{if } u < 1 \\ 1, & \text{otherwise} \end{Bmatrix}$

Variables affecting the potential of a node have now been defined. Equation (3) express how $V_i^{(t)}$ is calculated at each time step.

$$V_i^{(t)} = \left[\sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in F_j} \hat{w}_{ij}\, \varepsilon_{ij}(x)\right] - \eta_i(u) \qquad (3)$$

Node i fires when its potential $V_i(t)$ reaches the threshold Θ. This potential is affected by connection weights $\hat{w}_{ij}$ coming from each presynaptic node j. The set of presynaptic connections to node i is given by $\Gamma_i = \{ j \mid j \text{ is presynaptic to } i \}$. F_j is the set of all firing times of presynaptic nodes j: $t_j^{(f)}$. Noisy connection weights linking j node to i node are $\hat{w}_{ij}$.

Learning

Long term potentiation (LTP) and depression (LTD) are the basic mechanisms of long-lasting modifications of synaptic efficiency. Hebb (1949) postulated that when presynaptic activity coincides with postsynaptic activity, the connection between both neurons is strengthened. According to recent experiments, the modification of synaptic efficiency depends on precise timing of afferent signals (neurotransmitters binding to receptors) and the postsynaptic neuron spike. LTP seems to require that postsynaptic action potential be simultaneous or subsequent to postsynaptic currents (Markram et al., 1997; Zhang et al., 1998). In short, when the signal from the presynaptic neuron firing arrives before, or during the spike of postsynaptic neuron, the synapse is strengthened (LTP). When the signal from the presynaptic neuron arrives after the spike of postsynaptic neuron, the synapse is depressed (LTD).

In the present model, the plasticity of a synapse w_{ij} is a function of three parameters: the firing time of the presynaptic neuron: $t_j^{(f)}$, the transmission delay between this firing and its effect on the postsynaptic neuron (d_{ij}), and the firing time of postsynaptic neuron $t_i^{(f)}$. Learning in INFERNET consists of modifying the weights of connections between nodes w_{ij} by a value ˇ w_{ij} (weights are all short integers from -32767 to 32767, which explains why the weight change values run from −1024 to 1024). The Hebbian learning function used is shown in Figure 2. This function follows empirical studies (Markram et al., 1997; Zhang et al., 1998). Similar functions have been used in various simulations by others (Levy & Horn, 1999; Munro & Hernandez, 1999).

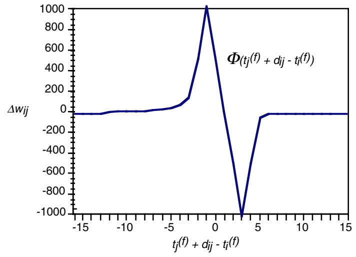

Figure 2: INFERNET's Hebbian learning function: when the signal from the presynaptic neuron arrives before or during the spike of the postsynaptic neuron, the synapse is strengthened (LTP); when the signal arrives after this spike, the synapse is depressed (LTD).

The learning algorithm attempts to reproduce the temporal relation between two successive inputs. This is particularly difficult because two successive inputs can be separated by several tenths of a second and a single connection cannot alone be responsible for such long delays. A long chain of successive pools of node firings is therefore required. This problem is illustrated in Figure 3. The problem is linking nodes *a* and *a′* that fire at time 0 with node *g* firing at time 49. In the learning phase, only nodes *a* and *a′* and *g*, 49 ms later, are externally stimulated. The system has to find a chain of node-firings that makes the target node *g* fire at time 49 when the probe nodes *a* and *a′* fire at time 0. The levels shown in Figure 3 are defined by the pools of firing nodes that separate the nodes firing in response to the input probe and the nodes responding to the target. Note that when simultaneous input from enough afferent neurons does not occur, the node will not fire. There is therefore a phenomenon of selection of only those nodes that have fired due to simultaneous inputs. This requires a large fan out of connections at all levels between the probe nodes and the target nodes.

The refractory state indicates when a particular node fired. We also know the delay of signal propagation from a presynaptic node to a postsynaptic node. From these two values we can, therefore, detect which synapse can contribute to a node firing at the right moment. In Figure 3, one can detect which nodes contribute to the firing of node *g* □ i.e., *e* and *f*, whose signals arrive at *g* virtually simultaneously. If *f* fired 9 ms before *g*, and *e* fired 11 ms. before *g*, their respective connections will be strengthened. Similarly, one can also determine which nodes contributed to the firing of *f* (i.e., *d* and *d′*), if *d* fired 7 ms before *f*, their connection will be strengthened (to a somewhat lesser extent because we are farther down the chain). This chaining rule acts as if a signal was going backwards from the target node to the probe nodes, losing a bit of its strength at each step. In order to reduce combinatorial explosion, only the *n* best contributing nodes are selected for the next level in this chaining rule. Connections between nodes will be modified according to equation (4):

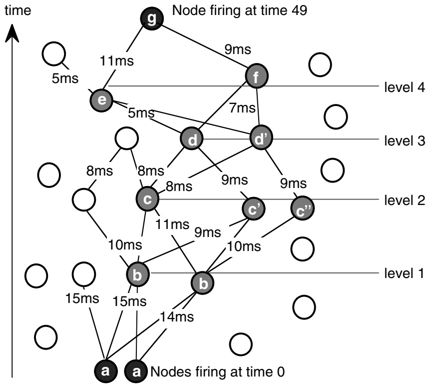

Figure 3: The chaining rule problem: Learning a path of neural firings that makes node *g* fire exactly 49 ms after nodes *a* and *a′*.

$$\Delta w_{ij} = \Phi\left(t_j^{(f)} + d_{ij} - t_i^{(f)}\right) - \lambda \tag{4}$$

$$\textit{where}\ \ \lambda = \begin{Bmatrix} -\textit{level if}\ \Phi\left(t_j^{(f)} + d_{ij} - t_i^{(f)}\right) \textit{is negative} \\ \textit{level otherwise} \end{Bmatrix}$$

This rule is based on the history of node firing and has neurobiological justification. For example, Markram, et al. (1998) show that the state of a synapse is indicative of its past activity. Moreover, empirical studies (Engert & Bonhoeffer, 1997) show that LTP also propagates from the originating synapse to neighboring synapses, lending further plausibility to the present chaining rule. In addition, each connection has a small decay factor (of −10 by epoch).

The learning algorithm is triggered only when external input is presented. We can imagine that external input provides a strong signal that triggers the chaining rule. Note that Hebbian learning does not seem to be dependent on this kind of signal and affects probably all synapses downstream from an action potential. Here, it is the target input that is the signal to launch the chaining rule. The objective is to link the probe nodes' firing to the target nodes' firing and to avoid reinforcing other irrelevant firings.

Simulation 1

In general, motor forgetting occurs more slowly than cognitive forgetting (Globerson, Nahumi, Ellis, 1998). By testing the present synfire chain algorithm for cognitive forgetting,, we reasoned that if catastrophic interference disappeared for this paradigm, the same algorithm would eliminate it for precise motor learning.

The following simulation is based on the original AB-AC paradigm used in Barnes & Underwood (1959). As in the original experiment, we created a list of "non-words" (A) and two associated lists of "words" (B and C). Each B and C word was coded over 6 nodes (out of 800 possible nodes) and each "non-word" A was coded over 16 nodes. Although the selection of nodes was made randomly, we ensured that there was very little overlap among the items in the A-list and among the items in the B-list. All lists consisted of eight items.

Each temporal firing sequence consisted of 16 nodes firing at time 0, corresponding to a presentation of a non-word from list A in the Barnes & Underwood experiments. Six nodes fired at time 60, triggered by the associated word from the B wordlist.

Once the network had learned to associate the items in list A with those in list B, the network then had to associate the node-firings associated with the items in the A-list with those of the C-list. As in Barnes & Underwood, we kept the similarity very low between the corresponding words in the B and C lists. This meant that very few nodes overlapped in the encoding of the corresponding words from the two lists. As the network learned the new set of associations, we tracked how fast new learning was taking place (i.e., how close the output of the network was to the desired word in List C) and, at the same, time, how far the output was from the originally learned word in List B.

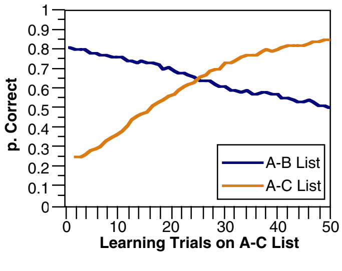

Figure 4: INFERNET performing AB-AC learning.

The results, based on 20 runs of the program, are shown in Figure 4. The Y axis indicates the proportion of correct node firings, i.e. the number of B and C node-firings within a –2ms window divided by the total number of nodes in B and C words. It is clear that, unlike the catastrophic interference observed in standard backpropagation networks (see Figure 1c), this sparsely-distributed network of spiking neurons can learn the second set of associations without catastrophically forgetting the previously learned list. These results are strikingly similar to those of Barnes & Underwood (Figure 1a).

Simulation 2

All of the model parameters for this simulation were identical to those of the preceding simulation, with the exception that wordlist C was replaced by a wordlist B′. All words in the B′ List were very similar to the corresponding words in the B list. Of the 6 nodes used by the representation of each B′ word, 4 of them were shared by the corresponding word in the B List.

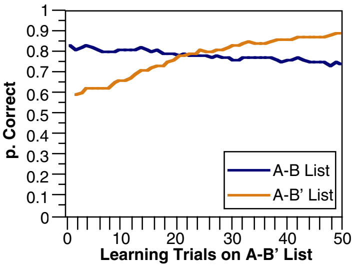

Figure 5: INFERNET performing AB-AB′ learning.

The results, based on 20 runs of the program, are shown in Figure 5. Again, this simulation closely reproduced the experimental results of Barnes & Underwood's second experiment (see Figure 1b). The results indicate the second associations are learned more quickly and the first associations are almost not forgotten, as for humans.

Conclusions

Human learning involves relating two signals separated in time, or linking a signal, an action and a subsequent effect. On occasion, the precise timing of these signals is of critical importance. A millisecond inaccuracy can mean that the spear thrown by the hunter will miss its target, that the gymnast will miss her landing, etc. Events may often be separated in time, but nonetheless, humans can link them, if necessary, with extraordinary accuracy, thereby allowing them to correctly perform a particular action at precisely the right moment. We have explored one major hypotheses concerning how the brain might achieve this — namely, synfire chains.

Clearly people are not born with encodings of this timing information. Hunters *learn* to throw projectiles accurately, gymnasts *learn* to land correctly. Precise temporal firing sequences must be learnable and permit the linking of two events with extreme precision.

A learning algorithm based on a Hebbian learning rule has been presented in this paper. We have briefly explored the ability of a sparsely-distributed network of spiking neurons, INFERNET, to learn synfire chains and, most importantly, we studied forgetting in this network of these chains. Unlike many current connectionist networks, we found that the forgetting of synfire chains is not subject to catastrophic interference, but rather, closely resembles the gradual forgetting curves exhibited in Barnes & Underwood□s (1959) paper on human forgetting. This is due to the sparseness of the number of paths (compared to the very large number of possible paths) from probe to target created by the learning algorithm. We hope to have demonstrated the importance of synfire chains for human cognition and to have shown an implementation in a network of spiking neurons. Finally, and crucially, our simulations indicate that synfire chains, so necessary for precision actions in the real world, may not be affected by catastrophic forgetting.

Acknowledgements

This research was supported by the Belgian PAI Grant p4/19 and the E.C. grant HPRN-CT-1999-00065.

References

Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex.* New-York: Cambridge University Press.

Barnes, J., & Underwood, B. (1959). □Fate□ of first-learned associations in transfer theory. *Journal of Experimental Psychology, 58*, 97-105.

Bienenstock, E. (1995). A model of neocortex. *Network: Computation in Neural Systems, 6,* 179-224.

Engert, F., & Bonhoeffer, T. (1997). Synapse specificity of long-term potentiation breaks own at short distances. *Nature, 388,* 279-284.

French, R. M. (1997). Pseudo-recurrent connectionist networks: An approach to the □sensitivity-stability□ dilemma. *Connection Science, 9,* 353-379.

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences, 3,* 128-135.

French, R. M., & Ferrara, A. (1999). Modeling time perception in rats: Evidence for catastrophic interference in animal learning. In *Proceedings of the 21st Annual Conference of the Cognitive Science Conference*. NJ:LEA, 173-178.

Globerson, S., Nahumi, A., & Ellis, S. (1998). Rate of forgetting for motor and cognitive tasks. *International Journal of Cognitive Ergonomics, 2*, 181-191.

Grossberg, S. (1982). *Studies of Mind and Brain: Neural Principles of Learning, Perception, Development, Cognition, and Motor Control.* Boston: Reidel.

Hebb, D. O. (1949). *The Organization of Behavior.* New York: Wiley.

Hertz, J.A., & Pr□gel-Bennet, A. (1996). Learning synfire chains by self organazation. *Network: Computation in Neural Systems, 7,* 357-363.

Levy, N., & Horn, D. (1999). Distributed synchrony in a Hebbian cell assembly of spiking neurons. In *Advances in Neural Information Processing Systems 11,* Cambridge Ma, MIT Press.

Maass, W., & Bishop, C. M. (1999). *Pulsed Neural Networks.* Cambridge, Ma, MIT Press.

Markram, H., Gupta, A., Uziel, A., Wang, Y., & Tsodyks, M. (1998). Information processing with frequency-dependent synaptic connections. *Neurobiology of Learning and Memory, 70,* 101-112.

Markram, H., L□bke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of Synaptic Efficacy by coincidence of Postsynaptic Aps and EPSPs. *Science, 275,* 213-215.

McClelland, J., McNaughton, B., & O'Reilly, R. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102,* 419-457.

McCloskey, M., & Cohen, N. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (ed.) *The Psychology of Learning and Motivation: Vol. 24*, pp. 109-164, NY: Academic Press

Munro, P., & Hernandez, G. (1999). LTD facilitates learning in a noisy environment. In *Advances in Neural Information Processing Systems 11,* Cambridge Ma, MIT Press.

Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovin, H., & Abeles, M. (1998). Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *Journal of Neurophysiology, 79,* 2857-2874.

Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psych. Review, 97*, 285-308.

Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology, 4,* 569-579.

Sougn□, J. P. (1999). *INFERNET: A neurocomputational model of binding and inference.* Doctoral dissertation, Universit□de Li□ge.

Sougn□, J. P. (2001). A learning algorithm for synfire chains. In R. M. French, & J. P. Sougn□ (Eds.) *Connectionist Models of Learning, Development and Evolution.* pp. 23-32 London: Springer Verlag.

Sterratt, D. C. (1999). Is biological temporal learning rule compatible with learning synfire chains? *Proceedings of the Ninth International Conference on Artificial Neural Networks.*

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W., & Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature, 395,* 37-44.