

# Sequence Learning in a Single Trial

Ali A. Minai and William B. Levy

Department of Neurosurgery  
Box 420, Health Sciences Center  
University of Virginia  
Charlottesville, VA 22908

## Abstract

While recurrent neural networks can store pattern sequences through incremental learning, there could be a trade-off between network capacity and the speed of learning. The brain may solve this problem by using a two-stage system: a compact, low-capacity subsystem for rapid temporary storage of a few sequences, and a larger, high-capacity, slow learning subsystem for long-term storage of all sequences. In this study, we evaluate the ability of sparsely connected networks to learn pattern sequences in a single exposure using very high learning rates. The key factor is the amount of recurrent inhibition in the system. Our results indicate that post-synaptic gating in the learning rule enhances the rapid learning ability of networks. We also suggest how such rapid learning networks could transfer their memories to long-term storage.

## 1. Introduction

Learning sequential information is necessary for a cognitive system to develop a viable model of its temporally-varying environment and underlies such basic functions as recognition, prediction and induction. While recurrent networks of threshold neurons have been investigated primarily as associative memory devices, there is now a substantial literature on using asymmetrically connected networks for learning sequences of patterns (see, e.g., Sompolinsky and Kanter, 1986). In solving this problem, a cognitive system faces a three-way trade-off between the speed of learning, representational adequacy, and capacity. The stimuli to be learned are complex and numerous, but they must be learned *very* quickly — within one or two trials — in order to be useful. On the other hand, large capacity necessitates sparse coding of the stimuli, which makes them vulnerable to intrinsic noise-like activity of the recurrent network. To avoid learning such noise, these stimuli must be learned slowly over a large number of repetitions so that the noise averages out and the signal correlations are preserved. Also, if the network has a complex structure (as the cerebral cortex does), it might be difficult to store the required associations without going through some complicated optimization process — again precluding rapid learning.

There is evidence to suggest that the mammalian brain solves the above dilemma by using a "buffered architecture". It is well-known that the long-term storage of relational memories is in the cortex. However, studies of lesioned animals and humans have shown that the memories are initially stored in the hippocampus and are then transferred to the cortex over a period of days or weeks (e.g., Squire and Zola-Morgan, 1991) This suggests that the hippocampus uses its simple, compact architecture and a relatively dense coding for the rapid learning of stimuli. The decrease in capacity does not matter because the hippocampus acts only as a palimpsest, or temporary storage; new memories can overwrite old ones provided they do not come in too rapidly, perhaps using some attentional mechanism to prioritize the importance of new stimuli. Once the stimuli are stored in the hippocampus, they can be "replayed" repeatedly as stimuli to the cortex, which can learn them through some slower mechanism.

In this study, we consider the "hippocampal" part of the problem: rapid learning and replay. We idealize the hippocampus as a sparsely connected recurrent network of threshold neurons, using a previously investigated network model (Minai and Levy, 1992a,b, 1993). In a sense, we see the CA3 region of the hippocampus, with its recurrent connectivity, as the basic functional part of the system. The network learns through modifying its synapses using some modification rule with some *learning rate*. We consider the problem of learning a sequence of orthogonal patterns *in just one exposure*, using a very high learning rate. We then see whether allowing the trained network to relax from random states produces spontaneous regeneration of the learned sequence. If so, the output of the relaxing network can be used repetitively to train a noisier network. The main interest of the paper lies in exploring three different learning rules and trying to evaluate the strengths and weaknesses of each.

## 2. The Network Model

The network model we use is derived from our earlier work (Minai and Levy, 1992a,b, 1993). A network consists of  $n$  excitatory threshold neurons. Each neuron,  $i$ , receives exactly  $k$  excitatory, modifiable recurrent connections from randomly chosen neurons (including itself). The 0/1 variable  $c_{ij}$  indicates the absence/presence of a connection from neuron  $j$  to  $i$ , while  $w_{ij}$  is the strength of that synapse. Each neuron,  $i$ , also receives a single, unique external input,  $x_i$ , from an  $n$ -dimensional binary stimulus through a very strong, non-modifiable excitatory synapse of strength  $v$ . All neurons have identical firing threshold  $\theta$ , feedback inhibitory weight  $K$ , and feedforward inhibitory weight  $L$ . All neurons receive identical shunting inhibition, which is supplied by a single, globally-connected inhibitory interneuron. If  $z_j(t)$  is the binary output of neuron  $j$  at time  $t$  and  $h(x)$  denotes the Heaviside function, the activation of  $i$  at  $t$  is given by:

$$y_i(t) = \frac{vx_i(t) + \sum_j w_{ij} c_{ij} z_j(t-1)}{\sum_j w_{ij} c_{ij} z_j(t-1) + K \sum_j z_j(t-1) + L \sum_j x_j(t)}$$

$$z_i(t) = h(y_i(t) - \theta) \quad 0 < \theta < 1$$

## 3. Learning Rules

We consider three learning rules with learning rate  $\epsilon$ :

(1) **The Post\_Synaptic Rule:**  $w_{ij}(t) = w_{ij}(t-1) + \epsilon z_i(t) [z_j(t-1) - w_{ij}(t-1)]$ .

(2) **The Pre\_Synaptic Rule:**  $w_{ij}(t) = w_{ij}(t-1) + \epsilon z_j(t-1) [z_i(t) - w_{ij}(t-1)]$ .

(3) **The Symmetric Rule:**

$$w_{ij}(t) = w_{ij}(t-1) + \epsilon \left[ z_i(t) [z_j(t-1) - w_{ij}(t-1)] + z_j(t-1) [z_i(t) - w_{ij}(t-1)] - z_i(t) z_j(t-1) [1 - w_{ij}(t-1)] \right].$$

There is strong evidence that the post-synaptic rule obtains in some parts of the hippocampus (Levy and Steward, 1979; Levy, 1982). However, there is also evidence for a pre-synaptic rule in the CA1 region (e.g., Mulkey and Malenka, 1992), and the symmetric rule is analogous to the Hebb rule with  $\pm 1$  neurons (Hopfield, 1982). The main difference between the rules is in the situations where synapses are depressed.

## 4. Methods

Most recurrent network models use off-line learning (e.g., Hopfield, 1982) and begin with all weights set to 0, thus side-stepping the whole issue of noisy activity. To get a more meaningful situation, we set all recurrent weights to a small but significant initial value of 0.2. This means that, given a fixed sized stimulus, fixed recurrent inhibitory and excitatory weights, and a fixed fan in,  $k$ , the activity in the network depends on the recurrent inhibition,  $K$ . Thus, we varied  $K$  as the main parameter during our simulations to show that higher values of  $K$  allowed the use of larger learning rates — up to a point. We tried three values of  $\epsilon$ : 0.05, 0.5, and 0.8. For each learning rule and each  $\epsilon$ , we used five different  $K$ 's: 0.06, 0.08, 0.10, 0.12, and 0.14. In each case, we averaged over six randomly generated networks, all with  $n = 200$  and  $k = 60$ . In all cases, we took  $v = 2.0$ ,  $L = 0.1$  and  $\theta = 0.5$ . The sequence to be learned was 20 patterns long, where the  $\mu$ th pattern consisted of the ten bits  $10(\mu-1)+1$  through  $10\mu$  active and the rest inactive. Thus, the patterns were mutually orthogonal and each neuron was stimulated by exactly one pattern in the sequence. This setting was created to avoid the complicating effects of interference between stimuli. Of course, the successive states of the network during learning were not mutually orthogonal due to the firing caused by recurrent excitation, and that was the factor controlled by  $K$ .

Training was accomplished by stimulating the network with one pass through the entire sequence followed by an all-0 stimulus while modifying the weights using one of the three rules. The quality of learning was tested by the network's ability to reproduce the entire sequence given only the first pattern. If  $z_i(t)$  represents the output of neuron  $i$  at time  $t$ , and  $\zeta_i^\mu$  represents the  $i$ th bit of pattern  $\mu$ , then the overlap between the network state and pattern  $\mu$  at time  $t$  is defined as:

$$q_\mu(t) = \frac{\sum_i \zeta_i^\mu z_i(t)}{\sum_i \zeta_i^\mu} - \frac{\sum_i (1 - \zeta_i^\mu) z_i(t)}{\sum_i (1 - \zeta_i^\mu)}$$

This gives  $q_\mu(t) = 0$  if the network produces all 0's or all 1's, and an expected value of 0 if neurons are

independently activated with some fixed probability. It returns 1 for perfect match with pattern  $\mu$  and  $-1$  for perfect mismatch. If the first pattern is given at the first step, the *completion quality measure*,  $Q$ , is then defined to be:

$$Q = \frac{1}{20} \sum_{\mu=1}^{20} q_{\mu}(\mu)$$

The second part of the simulation was concerned with the spontaneous regeneration problem. We evaluated the state of each network while it relaxed 15 times for 40 iterations each from different randomly chosen initial states (a total of 600 steps). Learning was switched off in these simulations, and the recurrent inhibition was lowered by 0.02 to encourage spontaneous activity. Each bit in the initial state had a probability 0.1 of being active. The relaxation data was collected both before and after training and evaluated for the spontaneous emergence of the learned sequence. This was done by padding the data with 19 all-0 states at the beginning and the end and sliding the sequence as a 20 step window over it, calculating the  $Q$  value at each step as  $I(t)$ . Any fragment of the sequence thus showed up as a peak in the  $I(t)$  time series. Taking  $I_{av}$  as the time average of  $I(t)$ , the measure of *impression quality* was calculated as:

$$I = \frac{\sum h(I(t) - I_{av}) I(t)}{\text{number of relaxation steps} / 20}$$

Thus, if the relaxation produced the entire sequence every 20 steps,  $I$  would be about 1 (allowing for the initial and final padding). However,  $I$  could also exceed 1 if fragments of the sequence overlapped in time.

## 5. Results and Discussion

The results from the simulations are shown in Figure 1. As expected, they show that more recurrent inhibition allows the use of a higher learning rate by suppressing recurrent activity. However, the three rules showed distinct differences in performance. The most noticeable fact was the viability of the post-synaptic rule over a broad range of  $K$  values. The pre-synaptic rule did poorly in completion except for very high  $K$ . However, since high  $K$  impeded the spontaneous replay of the sequence, the overall performance of the presynaptic rule was not satisfactory. Presumably, this inferior performance could be improved if training occurred at high  $K$ , and then  $K$  was lowered by more than 0.02 during relaxation. The reason why the post-synaptic rule does better than the pre-synaptic at low inhibition is as follows. Suppose the association to be learned is from pattern  $A$  to pattern  $B$ , but when  $A$  is activated from the outside, a small part,  $C_1$ , of a third pattern,  $C$ , is also activated. When this state is followed by the activation of  $B$ , both rules strengthen the synapses from  $A$  to  $B$  and from  $C_1$  to  $B$ . However, the post-synaptic rule also *weakens* synapses from  $C - C_1$  to  $B$ , probably leaving the spurious  $C$  to  $B$  association weakened overall. The presynaptic rule, on the other hand, leaves the synapses from  $C - C_1$  untouched, thus enhancing the spurious association from  $C$  to  $B$ . Then, when  $C$  is activated from the outside, it elicits a portion of  $B$ , which interferes with the learning of the correct  $C$  to  $D$  association. At a high  $\epsilon$ , such false associations can be very strong and can very quickly swamp the true ones, precluding the learning of the sequence. The reason why the pre-synaptic rule does better at high  $K$  is that false activations such as  $C_1$  are suppressed to begin with. Of course, the post-synaptic rule also enhances some false associations, and to see why these matter less, consider the situation when the activation of  $A$  is followed by the activation of the correct pattern,  $B$ , and a part,  $C_2$ , of pattern  $C$ . Both rules strengthen the  $A$  to  $C_2$  synapses, the pre-synaptic rule weakens the  $A$  to  $C - C_2$  synapses while the postsynaptic rule leaves them unchanged, thus enhancing the spurious association from  $A$  to  $C$ . However, this hardly matters, since  $A$  will never be re-activated from the outside and its false association does not interfere with learning. Thus, one main conclusion from our simulations is that the postsynaptic rule is better for one-step learning in noisy situations. The symmetric rule's performance is virtually as good as that of the post-synaptic one. This rule might be even better when learning sequences of patterns with overlap, or when learning multiple sequences which share patterns.

Turning to the relaxation situation, it is very noticeable that, even though training with  $\epsilon = 0.5$  at high  $K$  does lead to reasonable completion performance, it produces virtually no impression on the network's spontaneous dynamics. Only training with  $\epsilon = 0.8$  succeeds in impressing the sequence strongly enough for spontaneous elicitation. However, very high  $K$  works against such spontaneous replay by squelching activity. Thus, the post-synaptic rule, which learns well at lower  $K$ , also has the best impression performance. Figure 2 shows the  $I(t)$  time series before and after learning for a network with training  $K = 0.1$  and  $\epsilon = 0.8$ . Both relaxations used  $K = 0.08$ .

## 6. Conclusion

The results of this brief comparative study indicate that a learning rule with post-synaptic gating can be used to rapidly embed pattern sequences in the dynamics of relatively sparse recurrent networks. Spontaneous activity

under conditions of reduced recurrent inhibition can then internally elicit fragments of these sequences, which could be used to train larger, noisier networks at much lower learning rates. Future stimulus sequences can then overwrite earlier ones, thus requiring little capacity from the small network.

**References:**

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities, *Proc. Nat. Acad. Sci., USA*, **79**: 2554-2558.

Levy, W.B. & Steward, O. (1979). Synapses as associative memory elements in the hippocampal formation, *Brain Research*, **175**: 233-245.

Levy, W.B. (1982). Associative encoding at synapses, *Proc. 4th Ann. Conf. Cognitive Sci. Soc.*, pp 135-136.

Minai, A.A. & Levy, W.B. (1992a). The dynamics of sparse random networks, *In review*.

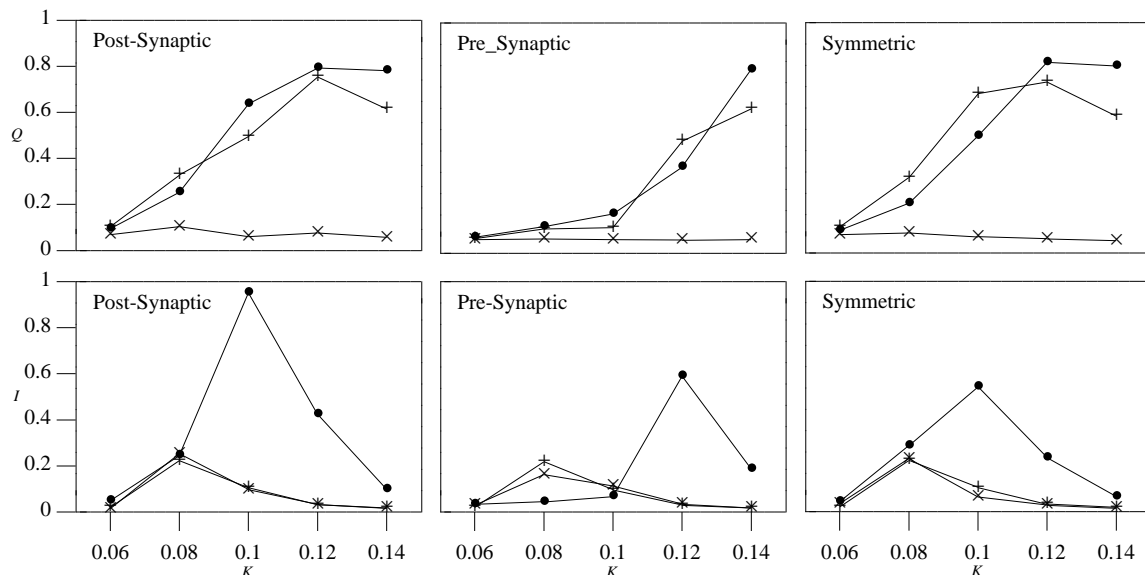
Minai, A.A. & Levy, W.B. (1992b). Setting the activity level in sparse random networks, *In review*.

Minai, A.A. & Levy, W.B. (1993). Predicting complex behavior in sparse asymmetric networks, *Advances in Neural Information Processing Systems 5*, San Mateo, CA: Morgan Kaufmann Publishers (to appear).


Mulkey, R.M. & Malenka, R.C. (1992). Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus, *Neuron*, **9**: 967-975.

Sompolinsky, H. & Kanter, I. (1986). Temporal association in asymmetric neural networks, *Physical Rev. Lett.*, **57**: 2861-2864.

Squire, L.R. & Zola-Morgan, S. (1991). The medial temporal lobe memory system, *Science*, **253**: 1380-1386.



**Figure 1:** Results for completion quality (top) and impression quality after training (bottom) for the three rules. The learning rates are: 0.05 (x); 0.5 (+); and 0.8 (●).



**Figure 2:** Time series of impression values in a network trained with  $\epsilon = 0.8$ ,  $K = 0.1$ . Graph (a) shows the situation before training and graph (b) after training.  $K = 0.08$  during both runs. Peaks indicate spontaneous emergence of sequence fragments.