

UW CSE Technical Report 02-07-03

Temporal Sequence Learning With Dynamic Synapses

Aaron P. Shon
Dept. of CSE
114 Sieg Hall, Box 352350
University of Washington
Seattle, WA 98195
aaron@cs.washington.edu

Rajesh P. N. Rao
Dept. of CSE
114 Sieg Hall, Box 352350
University of Washington
Seattle, WA 98195
rao@cs.washington.edu

July 17, 2002

Abstract

Recent results indicate that neocortical synapses exhibit both short-term plasticity and long-term spike-timing dependent plasticity. It has been suggested that changes in short-term plasticity are mediated by a redistribution of synaptic efficacy. This paper investigates how learning rules based on redistribution of synaptic efficacy can allow individual neurons and small networks of neurons to extract temporal information from incoming spike trains. Our results suggest that spike-timing dependent rules for redistribution of synaptic efficacy can provide a powerful and flexible mechanism for temporal sequence prediction and delay learning.

1 Introduction

Understanding how the activity of neocortical neurons encodes temporal information is a crucial open question for computational neuroscience. Recent *in vitro* experiments [1, 9, 8] suggest some plausible mechanisms. First, neocortical synapses have been shown to be dynamic: postsynaptic responses are not simply a function of the presynaptic firing rate multiplied by a synaptic “weight” but rather, reflect the short-term history of input spike trains. Second, paired firing of pre- and postsynaptic neurons tends to redistribute total synaptic efficacy, so that a synapse responds much more strongly to the first presynaptic event than to subsequent events in a spike train [9]. Third, long-term synaptic plasticity appears to change via a temporally asymmetric spike-timing dependent learning rule: a synaptic connection is strengthened if a presynaptic spike occurs slightly before a postsynaptic spike, and is weakened in the opposite case [8, 3].

These findings raise three critical questions that motivate our study: (1) What is the role of dynamic synapses in cortical information processing? (2) What is the relationship between spike-timing dependent plasticity (STDP) and the redistribution of

synaptic efficacy observed in dynamic synapses? and (3) How does STDP in conjunction with dynamic synapses allow cortical neurons to learn spatiotemporal input patterns? Dynamic synapses have previously been identified as a possible mechanism for gain control of cortical activity [1]. They have been used as memory buffers for "remembering" a history of presynaptic stimulation [7], and have been shown to be useful in modeling irregular bursts of activity in the cortex [14]. STDP has been suggested as a mechanism for learning temporal sequences [2, 12, 10], although without taking dynamic synapses into account.

In this paper, we explore two learning rules for STDP-like modification of dynamic synapses. The first rule adapts dynamic synapses in a Hebbian manner, reproducing the basic experimental results on redistribution of synaptic efficacy [9]. We show that such a rule allows neurons to predict spatiotemporal input patterns. A complementary rule, which adapts dynamic synapses in an anti-Hebbian manner, is shown to be useful for learning delays and making neurons selective for specific temporal patterns. This provides a neural basis for some previously proposed algorithms for delay learning [5] and temporal clustering [11] (see also [6]). While short-term plasticity and STDP have both been studied in isolation, our simulation results demonstrate that their combination can provide a powerful mechanism for temporal sequence learning.

2 Methods

2.1 Synapse Model

Our simulations used the following dynamics for modeling short-term synaptic plasticity:

$$d_t^{syn} = f^{syn} d_{t-1}^{syn} \text{ if a presynaptic spike occurred at time } t \quad (1)$$

$$d_{t-1}^{syn} + \frac{\alpha}{\tau_d} (d_0 - d_{t-1}^{syn}) \text{ if no presynaptic spike occurred at time } t \quad (2)$$

where d_t^{syn} represents a fractional amount of excitatory neurotransmitter available to synapse syn at time t , with d_0 a constant value of 1, α an integration rate (0.7 msec in the simulations), and τ_d the depression time constant (50 in the simulations). Parameter f^{syn} represents the fraction of neurotransmitter remaining at synapse syn after a presynaptic spike. Higher values of f correspond to weaker short-term depression. More sophisticated models of synaptic depression could also be used (e.g. [1]), but the simpler model described above sufficed for the purposes of this study.

2.2 Learning Rules

The novel component of our model is the interaction between short-term depression and STDP. We assume that STDP modifies the depression parameter f^{syn} , in addition to optionally modifying the peak synaptic conductance g_{peak}^{syn} . For each postsynaptic spike, the parameters are changed according to:

$$\Delta f^{syn} = -\gamma \mathbf{w}^T \mathbf{s}^{syn} \quad (3)$$

$$\Delta g_{peak}^{syn} = \beta \mathbf{w}^T \mathbf{s}^{syn} \quad (4)$$

where \mathbf{w} is the vector representing the temporally-asymmetric learning window (Fig. 1 (a), based on [8, 3]), \mathbf{s}^{syn} is the vector of presynaptic spikes to synapse syn centered at the time of the postsynaptic spike, and β and γ are gain factors. Note that updates to g_{peak} and f have complementary signs, so that an increase in synaptic depression is compensated for by an increase in peak conductance and vice versa. It is precisely this compensation that causes redistribution of synaptic efficacy in the model (see Fig. 1 (b)).

We investigated two complementary forms of the above learning rule:

- **Hebbian rule:** The gain parameters γ and β were set to positive quantities (we used $\beta = 0.052$ and $\gamma = 0.26$ in the simulations). We call this rule Hebbian because presynaptic spikes that occur before postsynaptic spikes cause a decrease in f (i.e. faster peak) and an increase in g_{peak} (higher synaptic efficacy) and vice versa for the reverse order.
- **Anti-Hebbian rule:** γ and β were set to nonpositive quantities (we used $\beta = 0$ and $\gamma = -0.05$ in the simulations shown in Fig. 3, and $\beta = 0$ and $\gamma = -0.005$ in the simulations shown in Fig. 4). In this case, presynaptic spikes that occur before postsynaptic spikes cause an increase in f (i.e. slower peak) and a decrease in g_{peak} (lower synaptic efficacy) and vice versa for the reverse order. Note that since $\beta = 0$ here, our results using the anti-Hebbian rule do not involve changing the peak synaptic conductance g_{peak} .

2.3 Neuron Model

We used standard leaky integrate-and-fire neurons in the simulations with resting potential $V_{rest} = -60$ mV and threshold $V_{threshold} = -40$ mV. The membrane resistance and capacitance were: $R = 40\text{M}\Omega$ and $C = 0.5$ nF. The refractory period τ_{refrac} was 3.5 msec. Potentials for excitatory and inhibitory synapses were: $E_{syn} = 0$ mV and $E_{inh} = -80$ mV. Peak synaptic conductances for the network whose output is shown in Fig. 4 were fixed at $g_{peak} = 0.17$ nS for excitatory synapses and $g_{inh} = 0.07$ nS for inhibitory synapses.

3 Results

3.1 Hebbian Redistribution of Synaptic Efficacy

Our first set of simulation results demonstrate how the combination of Hebbian STDP with short-term plasticity can reproduce the experimental results on redistribution of synaptic efficacy [9]. Recall that in the Hebbian case, pairing a presynaptic spike with a postsynaptic spike that occurs a few milliseconds later causes a decrease in f^{syn} in our model (Equation 3). Fig. 1 (b) shows the effect of decreasing f^{syn} on the postsynaptic responses to a fixed input spike train. Decreasing f^{syn} causes the synapse to respond much more strongly to earlier presynaptic events than to subsequent events in a spike train, moving the peak response to closer to the onset of the input spike train. This is accompanied by a gradual decrease in the overall magnitude of the responses, which

is compensated for during learning by increases in g_{peak} (not shown, see Equation 4). This compensation stabilizes the learning process and captures the synaptic behavior seen in experiments on redistribution of synaptic efficacy (cf. Fig. 1 in [9]).

In the next set of simulations, we investigated whether Hebbian redistribution of synaptic efficacy is conducive to temporal sequence learning. We considered the simple case of a single integrate-and-fire neuron with two synapses. At the onset of each trial, one of the synapses received a train of input spikes at a fixed rate r_1 . After a delay τ_{delay} (Fig. 1 (c)), the other synapse received a train of input spikes at a rate r_2 . Each spike train lasted for a specific duration (35 msec in the simulations, with $r_1 = r_2 = 150$ Hz, $\tau_{delay} = 15$ msec, peak synaptic conductance before training $g_{peak} = 0.07$ nS, and maximal peak conductance $g_{max} = 0.2$ nS). We trained the neuron for 100 trials (each lasting 210 msec), sufficient to allow synaptic parameters to converge.

Fig. 1 (c) shows that the neuron has learned to redistribute its synaptic efficacies to recognize and fire in response to the temporal order of the input spike trains. Furthermore, the neuron’s output is predictive, in the sense that it learns to fire soon after the onset of the second spike train before receiving the spike train in its entirety.

Traditional learning rules that modify g_{peak} alone are unable to capture this behavior. Fig. 2 contrasts the adjustment of both f and g_{peak} with adjustment of g_{peak} alone. Again, a neuron capable of redistributing its synaptic efficacies learns to recognize the given input sequence (Fig. 2 (a)) and does not respond to a different ordering of the inputs (Fig. 2 (b)). On the other hand, a neuron that modifies g_{peak} but not its f values (which remain fixed at $f = 1$) repeatedly spikes on both the training sequence and on a different ordering of the sequence after learning (Fig. 2(c),(d)).

3.2 Anti-Hebbian Redistribution of Synaptic Efficacy

The previous section illustrated how a Hebbian rule for redistribution of synaptic efficacy can enable a neuron to learn to predict an input sequence of spike trains. The same rule however is not well-suited to learning the timing delays between various inputs. Consider once again the simple case of two input spike trains separated by a temporal delay (Fig. 1 (c)). Fig. 3 (a) (top panel) illustrates the corresponding EPSPs generated in a neuron with randomly chosen values for f . Ideally, for the neuron to be selective for this sequence, we would like to maximize the overlap between these two sets of EPSPs, thereby increasing the probability that the neuron will fire in response to this temporal sequence. Consider the effect of the Hebbian rule when an output spike occurs somewhere in the middle of these two input events (say at 50 ms). The Hebbian rule will decrease f for the first input synapse and increase f for the second synapse, thereby *minimizing* the overlap between the EPSPs as shown in Fig. 3(a) (bottom panel) (note the new values for f).

On the other hand, the anti-Hebbian rule is well-suited to learning input delays. As shown in Fig. 3(b), the rule adapts the depression parameter f in the correct direction so as to maximize the overlap between the EPSPs due to the temporally separated inputs (g_{peak} was held constant). This allows the neuron to become selective for this input sequence, generating an output spike after the onset of the second input train. The anti-Hebbian rule can also control the timing of output spikes by converging to stable values of f different from the extremum values 0 or 1. This is shown in Figs. 3 (c) and

(d) for a neuron with a single input synapse whose peak conductance is high enough to generate an output spike. In this case, the learning rule converges to a value for f that balances the contributions of presynaptic spikes before and after the output spike(s). Such stability is a prerequisite for maintaining selectivity for a given input sequence once training is completed.

3.3 Feedforward network with recurrent inhibition

In the final set of experiments, we investigated whether the anti-Hebbian learning rule for f (with g_{peak} held constant) could allow a network of neurons with mutually inhibitory recurrent connections to become selective for specific input sequences (Fig. 4 (a)). These simulations were intended to extend the result in Fig. 3 (b) to the case of multiple input sequences. The depression parameters f were randomly initialized to allow symmetry breaking and the mutual inhibition implemented a form of competition among neurons so that different neurons could code for different input sequences. The set of input sequences comprised 6 random permutations of input spike trains labeled A, B, and C; each sequence was 35 msec in length. Each of the 6 sequences was presented to the network 50 times, in round-robin fashion, for a total of 300 training iterations.

Fig. 4 (b) shows the broad initial selectivities of the 10 neurons in the network to each of the 6 different temporal sequences. As shown in Fig. 4 (c), the anti-Hebbian rule tailors the synaptic f values in such a way that some neurons code for no patterns while others become highly selective for a small subset of the possible input sequences. The space of input sequences is thus partitioned among the coding neurons as a result of learning. Compared to the firing rates before learning, firing rates after learning are much less noisy (in the sense that the large number of weakly-responding neurons before learning are suppressed from firing after learning). Additionally, the responses of those neurons that learn to code for each sequence are higher after learning than before learning. None of the neurons in the network shown here responds strongly to sequences that are pairwise opposites; e.g. the neuron that responds strongly (more than 20 Hz) to the sequence "ABC" will not respond strongly to the sequence "CBA". The ability to respond to one ordering of patterns in a sequence but not the opposite ordering may be useful for recognizing temporal subsequences in sensory data, e.g. motion in visual information or onset of different frequencies in hearing information.

4 Conclusion

Short-term synaptic plasticity and STDP have emerged as two important properties of neocortical synapses. The interaction between them has remained unclear. Using simulations, we showed that Hebbian STDP can reproduce the changes in short-term plasticity known as redistribution of synaptic efficacy observed in neurophysiological experiments. Our results suggest that such a rule allows prediction of temporal sequences. A complementary rule based on anti-Hebbian STDP is well-suited for learning the delays between various inputs. Redistribution of synaptic efficacy allowed a neuron to become selective for specific input patterns by introducing an asymmetry in

synaptic excitation over time. This asymmetry lead to temporal selectivity: the neuron fired only if the input spike trains generate a set of appropriately aligned EPSPs that pushed the membrane potential above spiking threshold.

Our results suggest a strong computational role for redistribution of synaptic efficacy in temporal sequence learning. As a specific example, for moving stimuli, our model predicts the development of direction selectivity (Fig. 2 (a),(b), Fig. 4), an important property of neurons in the visual cortex (see also [4]). The temporal range of pattern selectivity in the present model is clearly limited by the width of the STDP learning window. This range could potentially be increased by using recurrent excitation to provide contextual information [2, 12, 13]. Our current efforts are therefore focused on exploring the effects of redistribution of synaptic efficacies in recurrent spiking networks.

Acknowledgments

This research is being supported by a National Defense Science and Engineering Graduate Fellowship to APS, a Sloan Research Fellowship to RPNR, and NSF grant no. 130705.

References

- [1] L. Abbott, J. Varela, K. Sen, and S. Nelson. Synaptic depression and cortical gain control. *Science*, 275:220–224, 1997.
- [2] L. F. Abbott and K. I. Blum. Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex*, 6:406–416, 1996.
- [3] G. Bi and M. Poo. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, 18(24):10464–10472, 1998.
- [4] F. Chance, S. Nelson, and L. Abbott. Synaptic depression and the temporal response characteristics of v1 cells. *J. Neurosci.*, 18:4785–4799, 1998.
- [5] C. Eurich, K. Pawelzik, U. Ernst, A. Thiel, J. Cowan, and J. Milton. Delay adaptation in the nervous system. *Neurocomputing*, 32:741–748, 2000.
- [6] J. Hopfield. Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376:33–36, 1995.
- [7] W. Maass and H. Markram. Synapses as dynamic memory buffers. *Neural Networks*, 2001. In press.
- [8] H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275:213–215, 1997.
- [9] H. Markram and M. Tsodyks. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382:807–810, 1996.

- [10] M. R. Mehta and M. Wilson. From hippocampus to V1: Effect of LTP on spatiotemporal dynamics of receptive fields. In J. Bower, editor, *Computational Neuroscience, Trends in Research 1999*. Amsterdam: Elsevier Press, 2000.
- [11] T. Natschläger and B. Ruf. Spatial and temporal pattern analysis via spiking neurons. *Network: Comp. Neural Sys.*, 9(3):319–332, 1998.
- [12] R. P. N. Rao and T. J. Sejnowski. Predictive sequence learning in recurrent neocortical circuits. In *Advances in Neural Information Processing Systems 12*, pages 164–170. Cambridge, MA: MIT Press, 2000.
- [13] D. Tank and J. Hopfield. Neural computation by concentrating information in time. In *Proc. Natl. Acad. Sci. USA*, volume 84, pages 1896–1900, 1987.
- [14] M. Tsodyks, K. Pawelzik, and H. Markram. Neural networks with dynamic synapses. *Neural Computation*, 10(4):821–835, 1998.

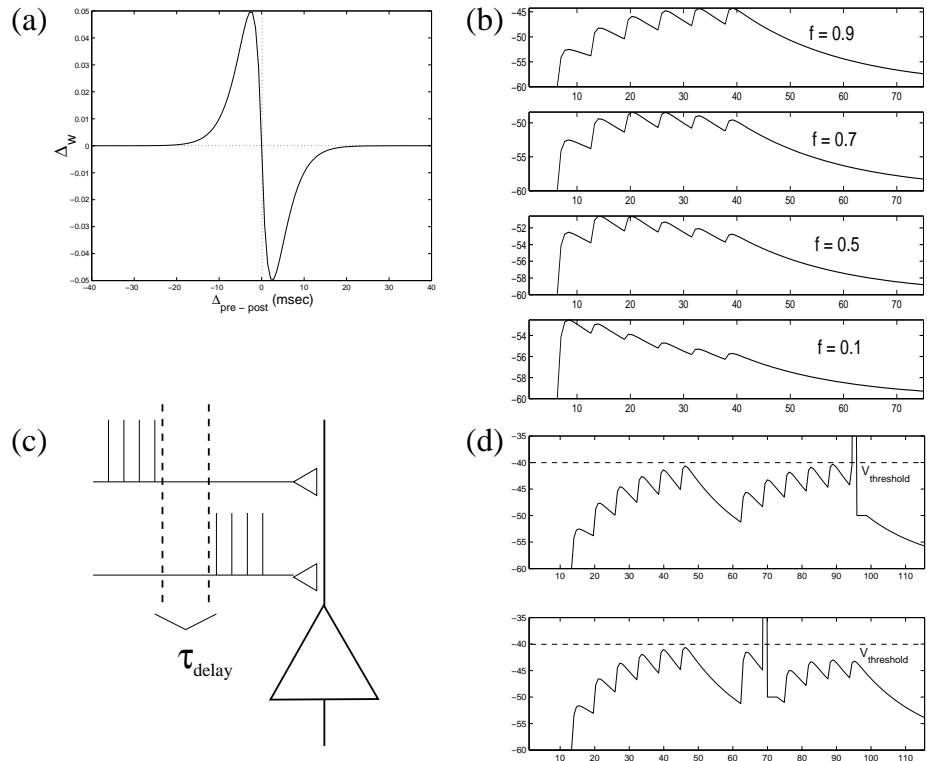


Figure 1: Learning to Predict using Redistribution of Synaptic Efficacy. (a) Temporally asymmetric learning window depicting the magnitude of change made to synaptic parameters f (and optionally g_{peak}) as a function of the relative timing of pre- and postsynaptic spikes (see Equations 3 and 4). (b) The four plots depict postsynaptic responses in the model to a fixed input spike train as a function of the depression parameter f ($= f^{syn}$ in the text). Decreasing f causes the synapse to respond with a faster peak for earlier presynaptic events. Note also the gradual decrease in the overall magnitude of the responses, which can be compensated for during learning by the complementary Hebbian rule for g_{peak} . (c) Depiction of the learning paradigm: synapses are adapted in response to input spike trains separated from each other by fixed delays. (d) Hebbian rule for modifying f leads to predictive firing: (Top panel) Response of the neuron before training. (Bottom panel) Response after training. The output spike now occurs at the onset of the second input.

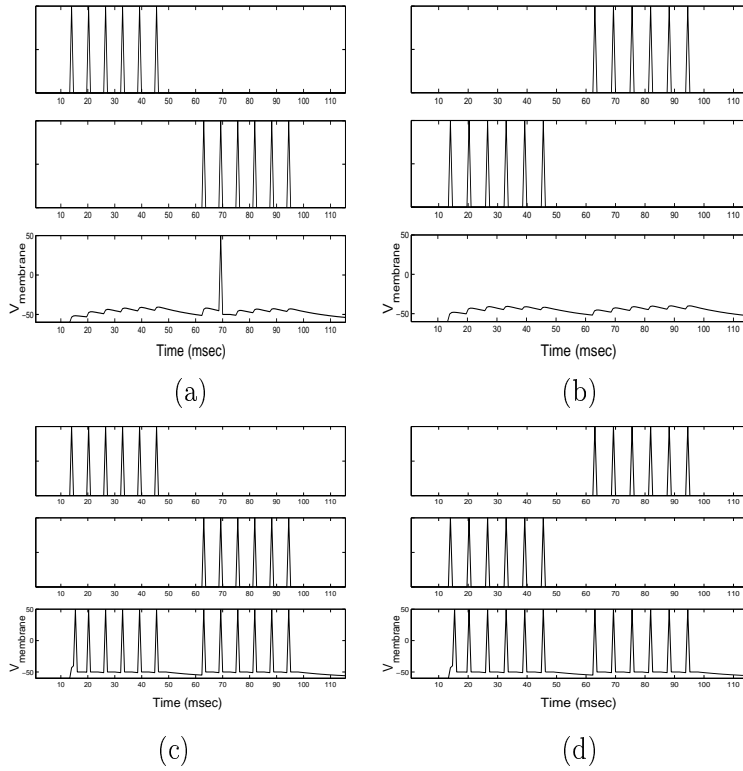


Figure 2: The Importance of Adapting Short-Term Synaptic Dynamics. (a) A model neuron that uses STDP to adjust both the depression parameter f and peak conductance g_{peak} spikes when presented with a temporal sequence used in training. (b) The same neuron does not fire when presented with a sequence that reverses the order of inputs in the training sequence. Note that the firing rate remains the same and only the temporal order is changed. (c,d) A model neuron that only adjusts peak conductance g_{peak} (as in traditional models) responds vigorously and indiscriminately to both the training sequence as well as the reverse-order sequence.

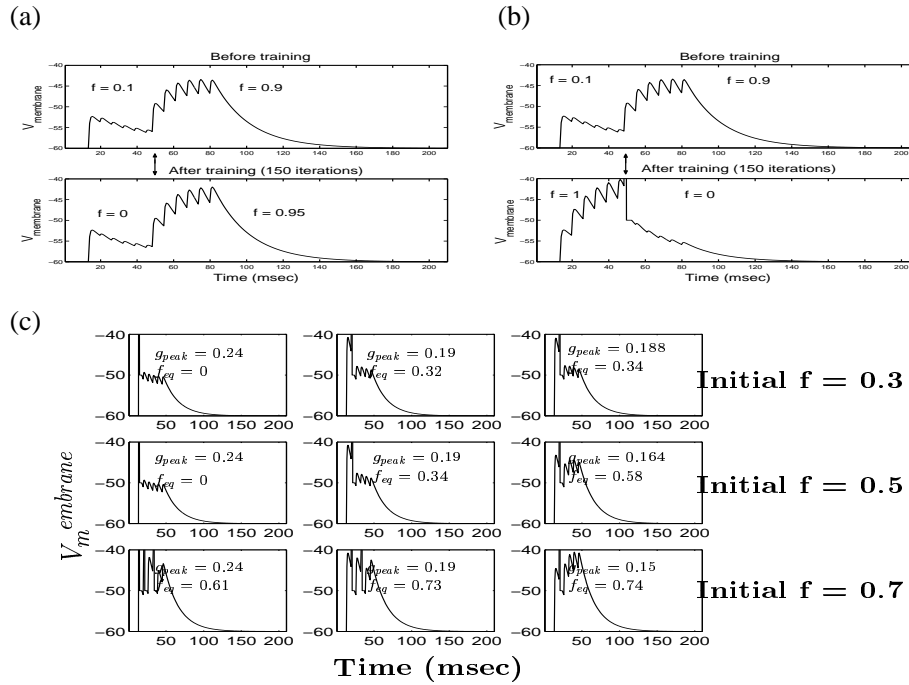


Figure 3: Anti-Hebbian Redistribution of Synaptic Efficacies. (a) A single neuron receives 2 input spike trains as in Fig. 1(c). The top panel shows the response of the neuron to the two spike trains. An output spike was elicited during training by injecting current at time $t = 50ms$ (double arrows). As shown in the bottom panel, the Hebbian rule for modifying f moves the peaks for the two trains in the opposite direction, preventing the neuron from learning the input delay. (b) The anti-Hebbian rule moves the peaks in the correct direction, allowing the neuron to learn to spike on its own at the appropriate time (without injecting an external current) after 150 training iterations. (c) Interactions between the timing of the presynaptic spike train, the value of g_{peak} , and the initial value of f lead to different equilibrium values f_{eq} following 150 training iterations. Each iteration involves stimulating a single neuron by a train of 6 spikes applied to a single synapse.

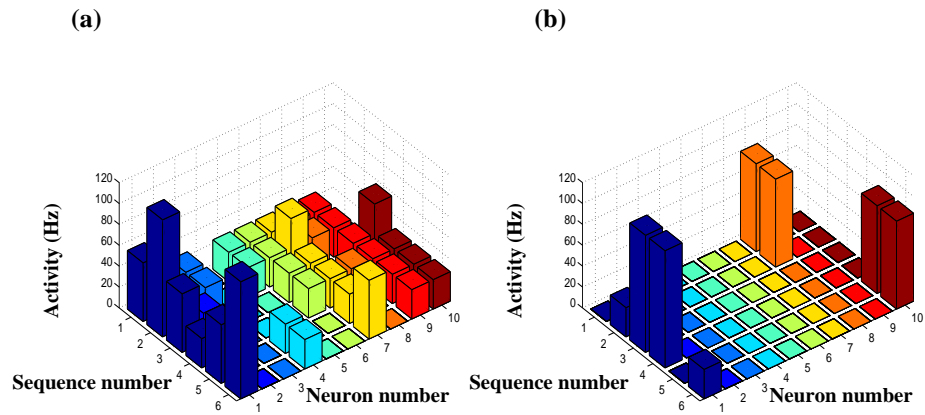


Figure 4: **Temporal Sequence Learning using Anti-Hebbian Redistribution of Synaptic Efficacies.** (a) shows the initial selectivities (measured as spike counts) of 10 neurons in a mutually inhibitory network for 6 input sequences (permutations of spike trains A, B, and C). (b) shows the selectivities developed by the neurons after learning using the anti-Hebbian rule for f . Note that the neurons that are active have become selective for a small subset of the original training set of sequences as a result of learning.