Effectiveness of Neural Network Learning Rules Generated by a Biophysical Model of Synaptic Plasticity

David J. Jilk, Daniel M. Cer, and Randall C. O'Reilly
Department of Psychology
University of Colorado, Boulder
345 UCB, Boulder, CO 80309-0345
djilk@jilk.com, cerd@cs.colorado.edu, oreilly@psych.colorado.edu

Abstract

We describe our initial attempts to reconcile powerful neural network learning rules derived from computational principles with learning rules derived "bottom-up" from biophysical mechanisms. Using a biophysical model of synaptic plasticity (Shouval, Bear, and Cooper, 2002), we generated numerical synaptic learning rules and compared them to the performance of a Hebbian learning rule in a previously studied neural network model of self-organized learning. In general, the biophysically derived learning rules did not perform as well as the analytic rule, but their performance could be improved by adjusting various aspects of the biophysical model. These results show that some progress has been made in integrating our understanding of biological and artificial neural networks.

Most artificial neural networks rely on analytic learning rules that are selected for their ability to successfully solve problems (e.g., Hopfield 1982, Oja 1982, Kohonen 1984), or that represent biologically plausible but highly idealized functions of the activity of the input and output neurons over a set of stimuli (e.g., Bienenstock et al. 1982). Conversely, many biophysical models of synaptic plasticity are developed primarily for the purpose of comparison with known experimental biological results (e.g., Shouval, Bear, and Cooper 2002, Froemke and Dan 2002, Sjöström et al. 2001).

In this research we attempt to bridge that gap by generating numerical "lookup table" learning rules for artificial neural networks using simulations of a biophysical model of synaptic plasticity. We applied those learning rules to a previously studied neural network model to ascertain their learning effectiveness as compared with idealized functions.

Methods

We based our simulations on a biophysical model developed by Shouval, Bear, and Cooper (2002), henceforth referred to as the SBC model. Although there are numerous other models that attempt to account for plasticity and incorporate recent findings about spike-timing dependent plasticity (STDP), they are either entirely analytic (Abarbanel et al. 2002), account only for STDP or joint activity scenarios (Froemke and Dan 2002, Sjöström et al. 2001), or model potentiation processes more closely associated with medium-term plasticity (Castellani et al. 2001). The SBC model includes an analytic function to calculate changes in potentiation based on calcium concentration; however, this function is derived from a prior biophysical model by Shouval et al. (2002b). In one experiment, we also combined the calcium dynamics of the SBC model with a different biophysical model of potentation (Lisman 1989, Zhabotinsky and Lisman 2000).

The SBC model is based on three key assumptions: (1) that different calcium levels trigger different forms of synaptic plasticity, (2) that NMDARs are the primary source of calcium, (3) that back-propagating action potentials, which contribute to STDP, have a slow after-depolarizing tail. Calcium current through NMDARs is calculated through a combination of glutamate activation which decays over time and a magnesium block that is sensitive to membrane potential. Actual calcium concentration is a result of the NMDAR current and a decay that is concentration-dependent. Changes in membrane potential result from both back-propagating action potentials (BPAPs) with a 100mV peak depolarization and excitatory post-synaptic potentials (EPSPs) with a 1.0mV peak

depolarization. In Shouval, Bear, and Cooper (2002), the SBC model is applied to model three key experiments: a pairing protocol, a presynaptic frequency protocol, and a spike-timing protocol.

Our simulations use the formulas directly from the SBC model as presented in Shouval, Bear, and Cooper (2002) and its supporting information, with one minor modification to account for a typographical error (Shouval 2002, private communication). The simulation engine calculates spike timing, membrane potential, NMDAR calcium current, calcium concentration, and synaptic potentiation in discrete time steps. A time step representing 0.1 mS was found to be adequate to exactly reproduce the results of Shouval, Bear, and Cooper (2002).

This simulator was then applied to generating a lookup table of weight changes for combinations of pre- and post-synaptic activation levels. Neuron activations were in the range [0,1], with an activation of 1.0 assigned to a 25 Hz average firing rate. Fractional activation was defined as the product of this maximum firing rate and the activation fraction. The lookup table was calculated in activation increments of 0.05. Spike timing was calculated as a Poisson distribution around the average rate. Synaptic weights were also in the range [0,1], and the weight change tables were calculated assuming a beginning weight of 0.5.

In some cases, we performed further adjustments on the learning rules. Some learning rules were biased toward potentiation or depotentiation, and we experimented with constant offset adjustments and multiplicative scaling ("re-centering") to remove this bias. We also created semi-synthetic learning rules by adjusting each value to be closer to an analytic rule (CPCA, see below) using a weighted-average. The results are reported as average success rate as a function of the percentage adjustment to CPCA.

Network model simulations were performed using the Leabra modeling framework in the PDP++ simulation environment (O'Reilly 1996, O'Reilly and Munakata 2000, McClelland et. al. 1986). The Leabra framework attempts to balance biologically plausible methods with computational power and feasibility. Individual neuron activation behavior is modeled through equations of ion current flow through the membrane of a point-neuron. Among other innovations, Leabra models can combine unsupervised associative learning with task-oriented error driven learning through a combination of biologically plausible feedback loops, a *k*-winners-take-all inhibition method, and a two-phase expectation/outcome learning process.

In a normal Leabra simulation, Hebbian weight changes are calculated with an analytic learning rule that has been used in a variety of neural network models (including competitive learning networks and Kohonen networks); we refer to this as the CPCA (conditional principal components analysis) rule (O'Reilly and Munakata 2000). The update function for CPCA is:

$$\Delta w_{CPCA} = y(x - w) \tag{1}$$

where w is the current synaptic weight, x is the rate-code activation of the pre-synaptic neuron, and y is the activation of the post-synaptic neuron. In our simulations, synaptic weight changes were calculated by interpolating within the simulation-generated lookup tables using a variant of Shepard's method (Gordon and Wixom 1978) and then applying a soft weight bounding function, as follows:

$$\Delta w = \begin{cases} (1 - w)\Delta w_{CPCA}, & \Delta w_{CPCA} > 0 \\ w\Delta w_{CPCA}, & otherwise \end{cases} \tag{2}$$

A sigmoidal contrast-enhancement function commonly used in Leabra models was also applied in some cases:

$$\hat{w} = \frac{1}{1 + \left(\frac{\theta w}{1 - w}\right)^{-\gamma}} \tag{3}$$

where $\hat{w}$ is the effective weight, θ is the offset of the sigmoid, and γ is the gain. This formula sharpens the contrast in the weights among the inputs to a given receiving unit: weights from strongly correlated inputs are enhanced, while those from weakly correlated inputs are reduced. Contrast enhancement is important for producing selectivity of feature detectors in self-organizing learning.

In this study, we applied the new learning rules to a previously studied model of self-organizing learning (O'Reilly and Munakata 2000, pp. 137-142). Although simplified, the learning processes in this model are analogous to early perceptual learning in V1 and V2 neurons. The model contains a 5x5 array of input units with full feed-forward connectivity to a hidden layer of 20 units. It uses a "soft" *k*-winners-take-all inhibition process operating on the hidden layer to restrict total activation. The training set for the model contains all 45 possible combinations of two vertical or horizontal "lines" (Fig 1a). A training session for the model runs 30 passes through the entire training set. Fig. 1b shows the resultant weights from the input layer to the hidden layer after training with the CPCA rule. Learning success in the model is gauged by a "probe" method (present a single line as the input) to determine the number of distinct lines (out of a possible 10) uniquely identified by the model. Each learning rule was run in a batch of 25 train/probe sessions each having random initial weights, and the success rate was averaged over these sessions.

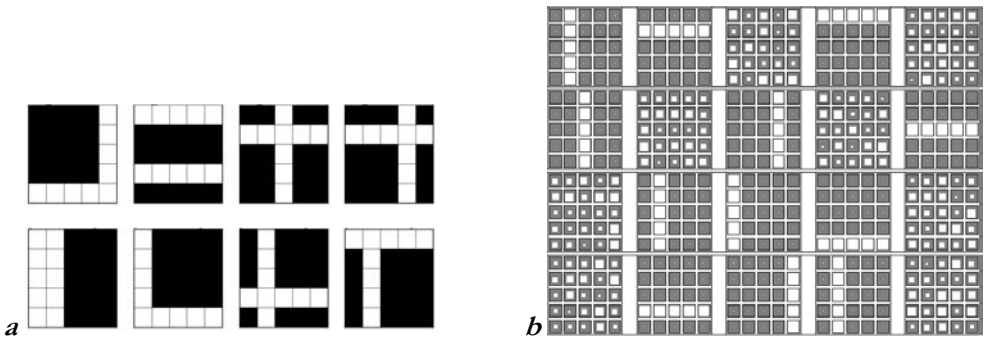

Figure 1: (a) A representative sample of the training inputs. The full set contains all 45 possible combinations of two "lines." (b) The weights from the input layer to the hidden layer after training. Each of the smaller grids in the 5 × 4 array represents the weights to one of the hidden units; the cells in the small grid represent the weights from each of the input units to that hidden unit. Note that each individual "line" is represented by at least one hidden unit.

Results

We generated a learning rule as described above using the SBC model as specified in Shouval, Bear, and Cooper (2002). As mentioned above, this model reproduced the experimental results in Shouval, Bear, and Cooper (2002) exactly, confirming that our simulator was performing as expected. A surface graph of the resulting lookup table is shown in Fig 2b.

Without re-centering, this rule performs poorly, averaging only 1/10 unique lines. Only by adjusting with a mix of at least 50% of the CPCA rule does the performance significantly improve. After re-centering, and using contrast enhancement (γ =6), the rule performs nearly perfectly, averaging 9.88 lines, but without contrast enhancement it performs poorly. Results for the various combinations are shown in Fig 3a.

In examining the details of the learning rule resulting from the SBC model, we found two areas where it deviated qualitatively from the CPCA rule. First, significant calcium concentrations were generated by pre-not-post activation, resulting in depotentiation. This is a result of the fact that the magnesium block of the NMDAR is not 100% effective at resting potential (Jahr and Stevens 1990). Second, no calcium transients resulted from post-not-pre activation, due to the fact that NMDARs are the only source of calcium in the SBC model and are inactive without glutamate present.

Although there is much debate about the impact of voltage-dependent calcium channels on synaptic plasticity, there is certainly some evidence that both L-type channels (Raymond and Redman 2002) and T-type channels (Song et al. 2002) contribute under certain circumstances. Consequently, we added a voltage-dependent calcium current that was proportional to the membrane potential (g = 0.003 μM/mS*mV) as it exceeded a reversal potential (-30 mV).

The effectiveness of the magnesium block of the NMDAR is strongly dependent on both extracellular [Mg^{2+}] and on membrane potential. Using the parameters in the SBC model, this block is only 94% effective at resting potential, resulting in significant influx of calcium in the face of pre-synaptic activity. To reduce this effect, we reduced the resting potential from -65.0 mV to -70.0 mV, and increased the extracellular [Mg^{2+}] from 1.0 mM to 5.0

mM. Although physiological levels of [Mg^{2+}] are normally assumed to be in the range of 1.0 to 1.2 mM, experiments typically consider ranges up to 10 mM (Jahr and Stevens 1990, Mayer et al. 1984, Nowak et al. 1984), and the actual concentration in the synaptic vicinity appears not to have been measured.

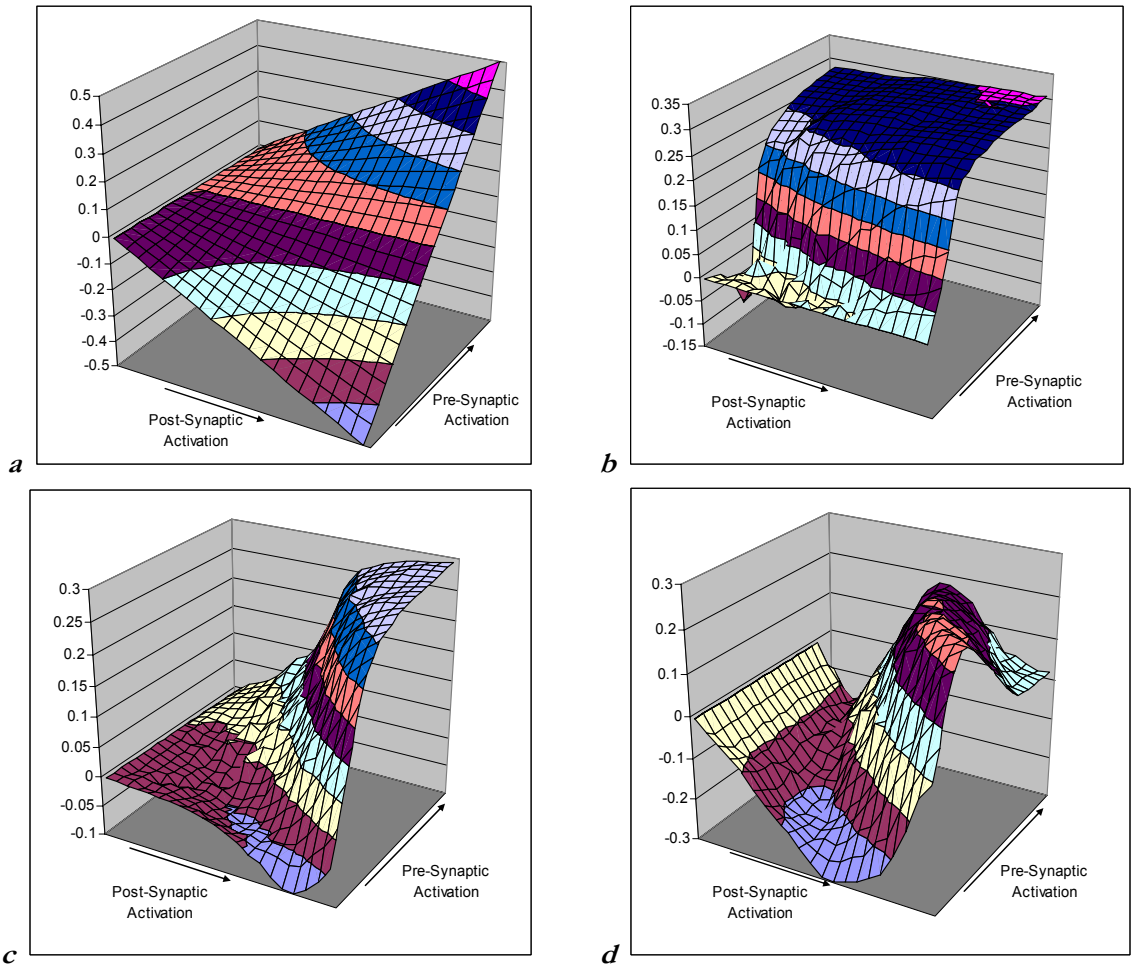

Figure 2: Learning rules viewed as surface functions. (a) The CPCA learning rule viewed numerically. Note the relative balance between potentiation and depotentiation. (b) The rule generated from the SBC model, prior to re-centering. Note the strong potentiation bias and the dramatic qualitative difference with CPCA when the post-synaptic neuron is inactive. (c) The rule generated from a modified SBC model, including voltage-dependent calcium currents and a stronger Mg block on NMDARs. Note the qualitative similarity to the CPCA rule. (d) A rule generated from the modified SBC model using an alternative potentiation function.

These changes resulted in a learning rule that is qualitatively very similar to the CPCA rule (Fig. 2c). Without re-centering, its learning performance is better than the original rule, but with re-centering it is worse. Complete results are shown in Fig. 3b. We also tested the experimental simulations of pairing, frequency, and spike timing from Shouval, Bear, and Cooper (2002) with these new parameters. We found qualitative but not quantitative agreement in the pairing protocol, and qualitative disagreement in the other experiments.

The potentiation function used in the SBC model is inconsistent with some studies of LTP/LTD. First, it contains no gap between calcium concentrations that cause depotentiation and those that cause potentiation, despite recent evidence that there is such a gap (Cho et al. 2001, Cormier et al. 2001). Second, it causes potentiation at all calcium concentrations above the initial threshold, despite evidence that there is a second threshold above which no potentiation occurs (Conti and Lisman 2002). Consequently, we also tested learning rules based on a potentiation function using the principles of Lisman (1989) and the experimental data of Cormier et al. (2001).

The learning rule generated in this case is shown in Fig. 2d, and results of this approach are shown in Fig 3c. This rule was very successful overall, learning 9/10 lines without any adjustment, although it only achieved perfect

performance with a 30% adjustment. When tested on the experimental protocols in Shouval, Bear, and Cooper (2002), this model did not show qualitative agreement with any of the results.

Finally, we tested each of the three learning rules using contrast enhancement and a multiplicative scaling approach to re-centering rather than an additive offset. As is apparent from the results in Fig. 3d, this approach brings the performance of the modified SBC model, and with small adjustment the original SBC model, to the level of performance of the CPCA rule.

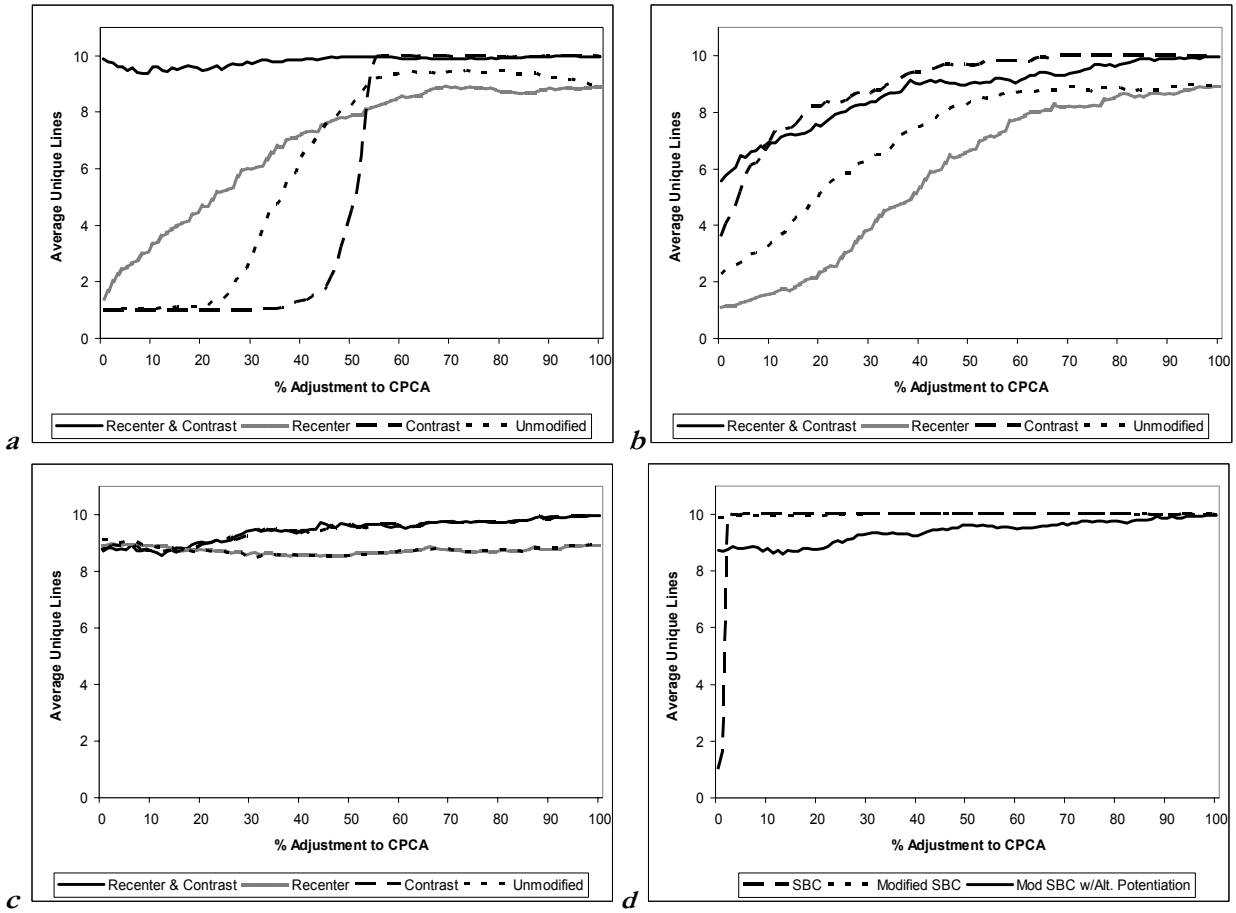

Figure 3: Learning performance as a function of weighted average adjustment to the CPCA rule. Except in (d), re-centering is a simple offset. (a) Results from the learning rule as generated directly from the SBC model. (b) Results from the modified SBC model. (c) Results from the modified SBC model with an alternative potentiation function. (d) Results from all three learning rules with a "multiplicative" recentering.

Discussion

An integrated theory of learning and memory would hopefully combine a detailed understanding of the biochemistry underlying synaptic plasticity with the progress made in artificial neural network modeling over the past two decades. A significant chasm still exists between the analytic learning rules used in artificial networks and the known biochemistry. We attempted to bridge that gap by assessing the performance of a learning rule derived from simulations of a comprehensive biophysical model of synaptic plasticity that is known to match many experimental results.

Our results show that the biophysical model must be adjusted in qualitatively significant ways to achieve reasonable performance on a simple associative learning task. A re-centering operation improves the performance nearly to perfection, but it also implies a change in an important qualitative aspect of the biophysical model, namely that the post-synaptic neuron does not depotentiate if the pre-synaptic neuron is inactive. Further, this qualitative change appears to be important to its learning success: when re-centering is performed multiplicatively (thereby retaining this qualitative aspect), the rule learns poorly unless it is slightly adjusted.

Addition of voltage-dependent calcium currents and a strengthening of the Mg block on the NMDARs, both potentially reasonable modifications, result in a learning rule that is qualitatively very similar to the analytic CPCA rule. Use of an alternative potentiation function generates a rule that learns effectively. However, with these changes the biophysical model no longer matches experimental results. An important question for future study is whether there is a combination of biophysical parameters that both generates a successful learning rule and qualitatively matches experiment.

In future work, we plan to extend this approach to the error-driven learning mechanism used in the Leabra framework. In error-driven learning, Leabra computes a difference between two temporally adjacent phases of activation: expectation and subsequent outcome. Preliminary results using a rule generated from the SBC model have shown that a brief expectation phase followed by a longer-lasting outcome phase can produce qualitatively-appropriate learning rules. However, these rules are much more difficult to analyze than the single-phase rules considered above.

References

Bienenstock E.L., Cooper L.N., Munro P.W. (1982), *Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex*, J Neurosci 2: 32-48.

Castellani G.C., Quinlan E.M., Cooper L.N., Shouval H.Z. (2001), *A biophysical model of bidirectional synaptic plasticity: dependence on AMPA and NMDA receptors*, Proc Natl Acad Sci USA 98: 12772-12777.

Cho K., Aggleton J.P., Brown M.W., Bashir Z.I. (2001), *An experimental test of the role of postsynaptic calcium levels in determining synaptic strength using perirhinal cortex of rat,* J Physiol 532: 459-466.

Conti R., Lisman J. (2002), *A Large Sustained Ca^{2+} Elevation Occurs in Unstimulated Spines During the LTP Pairing Protocol But Does Not Change Synaptic Strength*, Hippocampus 12:667–679.

Cormier R.J., Greenwood A.C., Connor J.A. (2001), *Bidirectional synaptic plasticity correlated with the magnitude of dendritic calcium transients above a threshold*, J Neurophysiol 85: 399-406.

Froemke R.C., Dan Y. (2002), *Spike-timing-dependent synaptic modification induced by natural spike trains*, Nature 416:433-438.

Gordon W.J., Wixom J.A. (1978), *Shepard's method of metric interpolation to bivariate and multivariate interpolation*, Mathematics of Computation 32: 253-264.

Hopfield, J.J. (1982), *Neural networks and physical systems with emergent collective computational abilities,* Proceedings of the National Academy of Sciences 79: 2554-2558.

Jahr C.E., Stevens C.F. (1990), *Voltage Dependence of NMDA-Activated Macroscopic Conductances Predicted by Single-Channel Kinetics,* J Neurosci 10: 3178-3182.

Kohonen, T. (1984), *Self-organization and associative memory,* Berlin: Springer Verlag.

Lisman J.E. (1989), *A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory*, Proc. Natl. Acad. Sci. USA 86: 9574-9578.

Mayer M.L., Westbrook G.L., Guthrie P.B. (1984), *Voltage-dependent block by Mg^{2+} of NMDA responses in spinal cord neurones*, Nature 309: 261-263.

McClelland J.L., Rumelhart D.E., and PDP Research Group (Eds.) (1986), *Parallel distributed processing, Volume 2: Psychological and biological models*, Cambridge, MA: MIT Press.

Nowak L., Bregestovski P., Ascher P., Herbet A., Prochiantz A. (1984), *Magnesium gates glutamate-activated channels in mouse central neurones*, Nature 307: 462-465.

Oja, E. (1982), *A simplified neuron model as a principal component analyzer*, Journal of Mathematical Biology 15: 267-273.

O'Reilly R.C. (1996), *The Leabra model of neural interactions and learning in the neocortex*, Ph.D. thesis, Carnegie Mellon University.

O'Reilly R.C., Munakata Y. (2000), *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*, Cambridge, MA: MIT Press.

Raymond C., Redman S. (2002), *Different calcium sources are narrowly tuned to the induction of different forms of LTP*, J Neurophysiol 88: 249–255.

Shouval H.Z., Bear M.F., Cooper L.N. (2002), *A unified model of NMDA receptor-dependent bidirectional synaptic plasticity*, Proc Natl Acad Sciences USA 99: 10831-10836.

Shouval H.Z., Castellani G.C., Blais B.S., Yeung L.C., Cooper L.N. (2002b) , *Converging evidence for a simplified biophysical model of synaptic plasticity*, Biological Cybernetics 87: 383-391.

Sjöström P.J., Turrigiano G.G., Nelson, S.B. (2001), *Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity*, Neuron 32: 1149-1164.

Song D., Wang Z., Berger T. (2002), *Contribution of T-Type VDCC to TEA-Induced Long-Term Synaptic Modification in Hippocampal CA1 and Dentate Gyrus*, Hippocampus 12: 689-697.