

Relating STDP to BCM

Eugene M. Izhikevich

Eugene.Izhikevich@nsi.edu, www.nsi.edu/users/izhikevich

Niraj S. Desai

desai@nsi.edu

The Neurosciences Institute, San Diego, CA, 92121, U.S.A.

We demonstrate that the BCM learning rule follows directly from STDP when pre- and postsynaptic neurons fire uncorrelated or weakly correlated Poisson spike trains, and only nearest-neighbor spike interactions are taken into account.

1 Introduction ---

Over the past several years, there has been increasing interest in a novel form of Hebbian synaptic plasticity called *spike-timing-dependent plasticity* (STDP; Markram, Lubke, Frotscher, & Sakmann, 1997; Bi & Poo, 1998; Debanne, Gahwiler, & Thomson, 1998; Feldman, 2000; Sjostrom, Turrigiano, & Nelson, 2001; Froemke & Dan, 2002), in which the temporal order of presynaptic and postsynaptic spikes determines whether a synapse is potentiated or depressed (see Figure 1A). While experiments to date have given us a fairly clear idea of how STDP affects synaptic weights when only isolated pairs of presynaptic and postsynaptic spikes are present, it is not clear how STDP should be applied to natural spike trains, which involve many spikes and many possible pairings of spikes (Froemke & Dan, 2002). Specifically, what is not clear is how plasticity at a given synapse builds up over time. This problem is interesting for its own sake, and also because it promises to shed light on the relationship between STDP and the best-studied forms of Hebbian plasticity, long-term potentiation, and depression (LTP and LTD), which explicitly make use of long spike trains (Bear & Malenka, 1994). Presumably all of these forms of plasticity arise from the same underlying biophysical mechanisms, and it should be possible to consider them within a single framework. Here we examine different implementations of STDP, compare them with a standard LTP/LTD implementation called the BCM (*Bienenstock-Cooper-Munro*) synapse (Bienenstock, Cooper, & Munro, 1982; Bear, Cooper, & Ebner, 1987), and in so doing arrive at certain constraints on how STDP should be implemented when considering natural spike trains.

In the BCM formulation, one considers instantaneous firing rates rather than individual spikes. Synaptic input that drives postsynaptic firing to

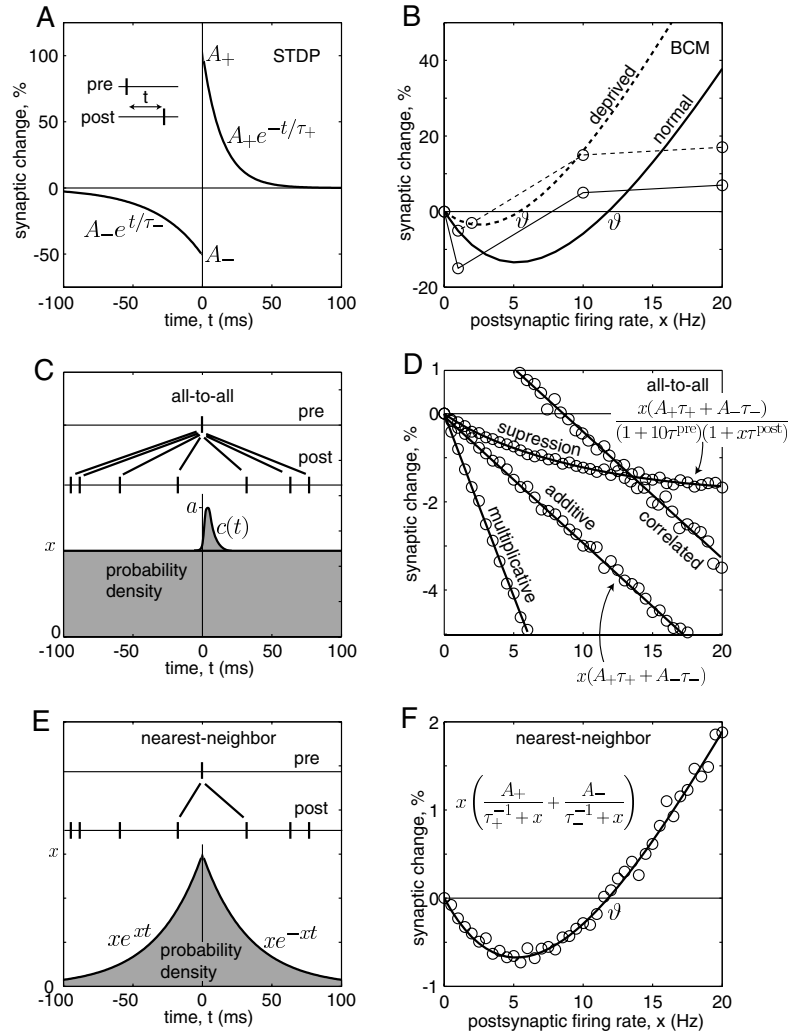


Figure 1: The STDP curve. The parameters $A_+ = 103\%$, $\tau_+ = 0.014$ sec, $A_- = -51\%$, $\tau_- = 0.034$ sec are taken from Froemke and Dan (2002). (B) Function controlling synaptic plasticity at the Cooper synapse receiving 20 Hz presynaptic stimulation. Data points (circles) are from visual cortex experiments by Kirkwood et al. (1996). Parameters from Figure 1A in equation 3.1 result in the normal curve. Increasing τ_+ by 10% results in the “deprived” curve. (C) All-to-all implementation of STDP: the net synaptic change is a combination of small changes induced by all possible pre- and postsynaptic pairs. (D) The result of application of STDP rule to Poisson spike trains: presynaptic: 10 Hz and 100,000 spikes; postsynaptic: x Hz and matching number of spikes. The analytical curves are derived in supplemental materials. (E) The nearest-neighbor implementation of STDP. For each presynaptic spike, only one preceding and one succeeding postsynaptic spike are considered. (F) The resulting BCM function. Parameters are as in Figure 1A.

high levels results in an increase in synaptic strength, whereas input that produces only low levels of postsynaptic firing results in a decrease (see Figure 1B). The threshold firing rate, the crossover point between potentiation and depression, is itself a slow function of postsynaptic activity, moving so as to make potentiation more likely when average activity is low and less likely when it is high. This sliding of the threshold serves to stabilize network activity and promote competition between synapses (Bienenstock et al., 1982). Considerable experimental evidence for this kind of plasticity has been obtained in neocortex and hippocampus at some of the same synapses at which evidence for STDP has also been obtained (Bi & Poo, 1998; Froemke & Dan, 2002; Kirkwood, Dudek, Gold, Aizenman, & Bear, 1993). Even so, it is not obvious how BCM plasticity and STDP are related or even if they are compatible. Consider, as an illustration, the extreme case in which all spikes of the postsynaptic neuron occur after those of the presynaptic one. This will always result in potentiation by an STDP rule, but it could result in either depression or potentiation by the BCM rule, depending on the exact value of the postsynaptic firing rate. To clarify this issue, we compare the two kinds of plasticity more closely in a more biologically realistic regime—that of uncorrelated or weakly correlated presynaptic and postsynaptic neurons that fire in a nearly Poisson manner, as do cortical neurons *in vivo*.

2 Classical Implementations of STDP

The form of the STDP curve, with its well-matched potentiation and depression portions, might cause one to guess that STDP applied to this kind of natural spike train should lead to a BCM-like curve similar to the one in Figure 1C. On the one hand, the peak of the potentiation half of the curve is larger than the minimum of the depression half, which suggests that high firing rates (small interspike intervals) should result in a net potentiation. On the other hand, the tail of the depression half is longer than the tail of the potentiation one, which suggests that low firing rates (large interspike intervals) should result in a net depression. And at some intermediate firing rate, a crossover should occur. But in fact, as we demonstrate in Figure 1D, this is not necessarily true.

A straightforward application of the STDP rule does not lead to potentiation in any case, when one uses independent Poisson spike trains with a mean postsynaptic firing rate x . In the standard additive implementation of STDP (Song, Miller, & Abbott, 2000), for each presynaptic spike, one takes into account all preceding and all succeeding postsynaptic spikes and then sums the contributions of the various pairings. If postsynaptic firing is a random process that is relatively independent of presynaptic firing, then the postsynaptic firing density is relatively flat, as in Figure 1C (no peak). Consequently, all postsynaptic spikes that precede the presynaptic one essentially sample the depression curve, and all postsynaptic spikes that follow the presynaptic one sample the potentiation curve. (Mathemat-

ical details of this procedure are given in the appendix.) The net effect is, then, the difference between the area under the positive portion of the STDP curve and the area under the negative portion, multiplied by the postsynaptic firing rate x . As a function of x , this is merely a straight line, in which potentiation never results, no matter how high the firing rate.

One can relax the assumption of completely uncorrelated pre- and postsynaptic spike trains and consider the case when postsynaptic firing density has a small peak right after the presynaptic firing (function $c(t)$ in Figure 1C). The magnitude of the peak a may be constant, or it may scale with the postsynaptic firing rate x . Both cases, considered in the appendix, result in a linear synaptic change similar to that produced by additive STDP without correlations. The case of constant a shifts the straight line up (see Figure 1D), while the case of a proportional to x changes the slope of the line. In any case, one does not see the transition from LTD to LTP, as is required by the BCM rule.

Also inadequate from this standpoint is a modification of the STDP rule recently proposed by Froemke and Dan (2002), in which the efficacy of each spike is suppressed by the preceding spikes of the same neuron, so that activity-induced synaptic modification depends not only on the relative spike timing between neurons, but also on the spiking pattern within each neuron. Specifically, these authors propose that associated with each presynaptic and postsynaptic neuron is an efficacy variable that drops to zero following a spike and that relaxes exponentially back to one (its maximal value) with a characteristic time τ^{pre} and τ^{post} on the order of tens of milliseconds. However, this proposal for treating multiple interactions cannot account for the potentiation of synapses at high firing rates. As we show in the appendix, the modification is basically equivalent to the standard STDP implementation, the only difference being that the presynaptic and postsynaptic spikes are now in chronic state of suppression. The precise numbers are changed, but the qualitative picture is not.

In the appendix, we explicitly consider eight implementations of STDP. In particular, we find that treating interactions between multiple pairs of spikes multiplicatively (van Rossum, Bi, & Turrigiano, 2000) rather than additively—that is, imagining that each pairing changes the conductance by a percentage of its existing value rather than by a fixed amount—makes no real difference. Additive and multiplicative models are in some cases equivalent because the latter often reduces to the former when one considers the logarithms of weights rather than the weights themselves.

3 Nearest-Neighbor Implementation of STDP

What does make a difference—what does make STDP compatible with BCM (and with classical LTP/LTD more generally)—is restricting which pairings contribute to plasticity (Sjostrom et al., 2001; van Rossum et al., 2000). Rather than considering all presynaptic and postsynaptic pairings equally,

there are good reasons to consider only nearest-neighbor pairs. One reason is that postsynaptic spikes backpropagate into the dendritic tree and reset the membrane voltage in dendritic spines. Consequently, the most recent postsynaptic spike overrides the effect of all the earlier spikes, so that the membrane voltage is really only a function of time since the latest postsynaptic spike. Similarly, the first succeeding postsynaptic spike may override the effect of subsequent spikes due to calcium saturation or glutamate receptor desensitization. Making this assumption, one finds that when the postsynaptic spike train is a Poisson process with firing rate x , the postsynaptic probability density (the probability of observing a spike with a certain delay t) becomes exponential in time, xe^{-xt} , as in Figure 1E (here we consider uncorrelated trains again). High (low) firing rates x result in predominantly small (large) intervals and hence in potentiation (depression). The expected magnitude of synaptic modification per one presynaptic spike has the form

$$\begin{aligned}
 C(x) &= \overbrace{\int_0^\infty A_+ e^{-t/\tau_+} x e^{-xt} dt}^{\text{average potentiation}} + \overbrace{\int_{-\infty}^0 A_- e^{t/\tau_-} x e^{xt} dt}^{\text{average depression}} \\
 &= x \left(\frac{A_+}{\tau_+^{-1} + x} + \frac{A_-}{\tau_-^{-1} + x} \right) \tag{3.1}
 \end{aligned}$$

depicted in Figure 1F. It coincides with the BCM synapse in the sense that low activity results in depression and large activity results in potentiation. Incidentally, the assumption of strictly nearest-neighbor pairs can be relaxed and the result retained if one instead considers one preceding but all succeeding postsynaptic spikes (semi-nearest-neighbor implementation). This also leads to the BCM synapse, as we show in the appendix.

Having an analytic expression for STDP's effects at a given firing rate x (see equation 3.1) allows us to attempt to relate data obtained using STDP and LTP/LTD experimental protocols more quantitatively. In particular, equation 3.1 indicates that the threshold between potentiation and depression (the zero crossing of $C(x)$),

$$\vartheta = -\frac{A_+/\tau_- + A_-/\tau_+}{A_+ + A_-}, \tag{3.2}$$

has positive values when

$$A_+ > |A_-| \quad (\text{potentiation dominates depression for short intervals}),$$

$$|A_-|\tau_- > A_+\tau_+ \quad (\text{depression area is greater than potentiation area}).$$

Using the STDP parameters obtained in layer 2/3 rat visual cortex (Froemke & Dan, 2002), we calculate a threshold value of around 12 Hz, which is near the value of 9 Hz found using LTP/LTD protocols in the same brain area (Kirkwood, Rioult, & Bear, 1996). One key feature of the BCM learning rule is that the threshold does not have a fixed value but slides between higher and lower values as a slow function of postsynaptic activity (Bienenstock et al., 1982). Equation 3.1 can be used to relate the sliding of the threshold to the biophysical processes underlying plasticity. In particular, the potentiation time constant τ_+ most likely depends on the kinetics of NMDA receptors, which in turn depend on their subunit composition. It has been shown (Philpot, Sekhar, Shouval, & Bear, 2001) that low levels of postsynaptic activity due to light deprivation increase the ratio of NR2B subunits to NR2A subunits in NMDA receptors. This in turn increases the time constant of NMDA receptors by up to 20%. As one can see from equation 3.2 and in Figure 1B, increasing τ_+ by as little as 10% results in sliding the calculated threshold by a factor of two, which is in agreement with the experimental data (Kirkwood et al., 1996).

Discussion

The question of how interactions between multiple spike pairs should be treated given an STDP rule has been considered previously (Sjostrom et al., 2001; Froemke & Dan, 2002; Song et al., 2000; van Rossum et al., 2000). In theoretical articles, the choice between applying the rule to all spike pairings or only to nearest-neighbor pairings has been made on a somewhat ad hoc basis, without a real empirical or biological justification. In experimental articles in which the question has been considered, the answers given have been limited in scope. Specifically, they have been limited to explaining data generated by applying STDP experimental protocols. What they have neglected is that a proper approach should take into account both STDP and classical LTP/LTD. Here we have attempted to do so, and our study suggests that the most widely used STDP implementations may not be adequate: the pairings should be restricted to only proximal spike pairs. And we see that a handful of reasonable assumptions generates a simple equation that can link the parameters of STDP to the BCM formulation of LTP/LTD, resulting in a more intuitive picture of how these forms of plasticity are related.

We have been considering two phenomenological models of plasticity, STDP and BCM, and important questions remain about how these models relate to the actual biophysical processes underlying synaptic plasticity. Although considerable data have been gathered about STDP applied to a single pair of spikes, many other aspects of this rule have yet to be worked out experimentally—for example, the issue of multiple spike pairs that we address in this theoretical article, as well as whether changes in synaptic efficacy depend on the size of the synapse. While the BCM formulation of LTP/LTD has been influential, how well it describes the empirical

data is open to question. Tests of the BCM idea have been largely indirect and have not convincingly validated certain aspects of the theory, in particular the idea that only postsynaptic activity determines the sign of plasticity.

4 Conclusion

This article addresses an important theoretical issue: how STDP, a novel form of Hebbian plasticity that has attracted much attention recently, relates to classical long-term potentiation and depression, in the form of the Cooper (or BCM) synapse. It is commonly (though not universally) believed that these two forms of plasticity arise from the same underlying biophysical process. The relationship between the two has been explored by, among others, van Rossum et al. (2000), Senn, Markram, & Tsodyks (2001), Kempter, Gerstner, & van Hemmen (1999, 2001), Castellani, Quinlan, Cooper, & Shouval (2001), and Abarbanel, Huerta, & Rabinovich (2002), though no attempts have been to contrast different implementations of STDP with respect to BCM.

Here we have determined conditions under which one is equivalent to the other, and by making a handful of biologically plausible assumptions, such as Poissonian spiking, we have derived a simple equation that relates one form of plasticity to the other (see equation 3.1 or Figure 1F). We believe that this is an important step in reconciling STDP and classical LTP/LTD and will be of use to a broad range of researchers, ranging from electrophysiologists studying synaptic plasticity to computational neuroscientists interested in understanding dynamics of networks.

Appendix: Different Implementations of STDP

If $\tau = t_j^{\text{post}} - t_i^{\text{pre}}$ is the interval between the i th and the j th spikes of the presynaptic and postsynaptic neurons, then the STDP synaptic modification is given by

$$w_{ij} = \begin{cases} A_+ e^{-\tau/\tau_+} & \text{if } \tau > 0 \\ A_- e^{+\tau/\tau_-} & \text{if } \tau < 0, \end{cases}$$

where $A_+ = 103$, $A_- = -51$, $\tau_+ = .014$ sec, and $\tau_- = .034$ sec are experimentally determined STDP parameters for layer 2/3 visual cortical neurons (Froemke & Dan, 2002). See Figure 1A.

Poisson spike trains with firing rates of 10 Hz and x Hz were generated for presynaptic and postsynaptic neurons, respectively. The postsynaptic firing rate x was varied systematically between 0 and 20 Hz, and the average synaptic modification per presynaptic spike was calculated using eight

STDP implementations, which are described below.

- *Classical additive STDP* adds the effect of all pairs,

$$w = \sum_{i,j} w_{ij}$$

and results in the straight line depicted in Figure 1d, which has a slope equal to the difference between the depression and potentiation areas. Indeed, for every presynaptic spike, all preceding postsynaptic spikes sample the depression curve, and all succeeding postsynaptic spikes sample the potentiation curve, so that the averaged net effect w is the difference between the areas beneath the curves multiplied by the postsynaptic firing rate. More precisely,

$$\begin{aligned} C(x) &= \overbrace{\int_0^{\infty} A_+ e^{-t/\tau_+} x dt}^{\text{average potentiation}} + \overbrace{\int_{-\infty}^0 A_- e^{t/\tau_-} x dt}^{\text{average depression}} \\ &= x(A_+ \tau_+ + A_- \tau_-). \end{aligned}$$

- *Classical multiplicative STDP*,

$$(1 + 0.01w) = \prod_{i,j} (1 + 0.01w_{ij}),$$

is equivalent to the additive model if one considers the logarithm

$$\log(1 + 0.01w) = \sum_{ij} \log(1 + 0.01w_{ij}).$$

This also results in a straight line with a slope equal to the difference between the areas bounded by the curves,

$$\int_0^{\infty} \log(1 + 0.01A_{\pm} e^{-t/\tau_{\pm}}) dt.$$

See Figure 1D.

- *Suppression additive STDP* assigns a suppression weight to each spike that depends on the time elapsed since the previous spike,

$$e_i = 1 - e^{(t_i - t_{i-1})/\tau},$$

where $\tau^{\text{pre}} = 28$ ms and $\tau^{\text{post}} = 88$ ms, so that the net effect is

$$w = \sum_{i,j} e_i^{\text{pre}} e_j^{\text{post}} w_{ij}.$$

This is the model of Froemke and Dan (2002), and the parameters are taken from that model. Because of the interactions between the spikes,

all the spikes are in a chronic state of “suppression,” that is,

$$\langle e \rangle = 1 - \int_0^\infty e^{-t/\tau} x e^{-xt} dt = \frac{1}{1 + x\tau},$$

and the averaged synaptic modification is the curve,

$$C(x) = x(A_+ \tau_+ + A_- \tau_-) \frac{1}{1 + 10\tau^{\text{pre}}} \frac{1}{1 + x\tau^{\text{post}}},$$

depicted in Figure 1D. It saturates at high firing rates but never crosses zero.

- *Classical additive STDP (correlated spike trains)*. When the pre- and postsynaptic neurons have correlated firings, the postsynaptic firing density can be represented in the form

$$p(t) = x + ac(t),$$

where $c(t)$ describes the shape of the peak of the cross-correlogram, as in Figure 1C, and a is a scaling factor, which may depend on the postsynaptic firing rate x . The averaged synaptic modification has the form

$$\begin{aligned} C(x) &= \overbrace{\int_0^\infty A_+ e^{-t/\tau_+} (x + ac(t)) dt}^{\text{average potentiation}} + \overbrace{\int_{-\infty}^0 A_- e^{t/\tau_-} (x + ac(t)) dt}^{\text{average depression}} \\ &= x(A_+ \tau_+ + A_- \tau_-) + ac_0, \end{aligned} \tag{A.3}$$

where

$$c_0 = \int_0^\infty A_+ e^{-t/\tau_+} c(t) dt + \int_{-\infty}^0 A_- e^{t/\tau_-} c(t) dt$$

is a parameter that depends on the exact shape of the peak of the cross-correlogram. If the magnitude of the peak is constant, then $C(x)$ has the same straight-line form as in the uncorrelated case except that it is translated by ac_0 , as in Figure 1D. If the magnitude of the peak is proportional to the firing rate, for example, $a = x$, then $C(x)$ is a straight line with the slope $A_+ \tau_+ + A_- \tau_- + c_0$.

- *Semi-nearest-neighbor additive STDP*. For each presynaptic spike, we consider only one preceding postsynaptic spike and ignore all earlier spikes. The motivation for this is the idea that the preceding postsynaptic spike overrides the effect of the earlier spikes due to sub-linear summation of membrane voltage in dendritic spines. If there is no saturation in Ca^{++} or glutamate/NMDA receptor dynamics, then all subsequent postsynaptic spikes must be considered. As a result, the postsynaptic probability density is exponential before the

presynaptic spike (as in Figure 1E) but flat after the presynaptic spike (as in Figure 1C). The resulting net modification per one presynaptic spike is

$$\begin{aligned}
 C(x) &= \overbrace{\int_0^{\infty} A_+ e^{-t/\tau_+} x dt}^{\text{average potentiation}} + \overbrace{\int_{-\infty}^0 A_- e^{t/\tau_-} x e^{xt} dt}^{\text{average depression}} \\
 &= x \left(A_+ \tau_+ + \frac{A_-}{\tau_-^{-1} + x} \right),
 \end{aligned}$$

which is similar to the curve depicted in Figure 1F.

- *Nearest-neighbor additive STDP.* For each presynaptic spike, only two postsynaptic spikes are considered: the one that occurs before and the one that occurs after the presynaptic firing. The resulting averaged curve,

$$\begin{aligned}
 C(x) &= \overbrace{\int_0^{\infty} A_+ e^{-t/\tau_+} x e^{-xt} dt}^{\text{average potentiation}} + \overbrace{\int_{-\infty}^0 A_- e^{t/\tau_-} x e^{xt} dt}^{\text{average depression}} \\
 &= x \left(\frac{A_+}{\tau_+^{-1} + x} + \frac{A_-}{\tau_-^{-1} + x} \right)
 \end{aligned}$$

is depicted in Figure 1F.

- *Nearest-spike additive STDP.* For each pre-synaptic spike, one finds the nearest postsynaptic spike (model 2 in by Sjostrom et al., 2001), regardless of whether it is before or after the presynaptic firing. Since the conditional postsynaptic probability density is $x e^{\pm 2xt}$, the resulting averaged curve has the form

$$C(x) = x \left(\frac{A_+}{\tau_+^{-1} + 2x} + \frac{A_-}{\tau_-^{-1} + 2x} \right).$$

- *Nearest-spike additive STDP with LTP winning* is similar to the one above except that the LTD due to “post then pre” pairs is not counted if the postsynaptic spike participated in LTP interaction. This corresponds to the model 3 implementation of STDP rule by Sjostrom et al. (2001). The conditional postsynaptic probability for LTP window is as above, that is, $x e^{-2xt}$, $t \geq 0$. The probability for the LTD window depends on the pre- and postsynaptic firing rate, and it can easily be determined numerically. If both rates are the same, then it can be approximated by

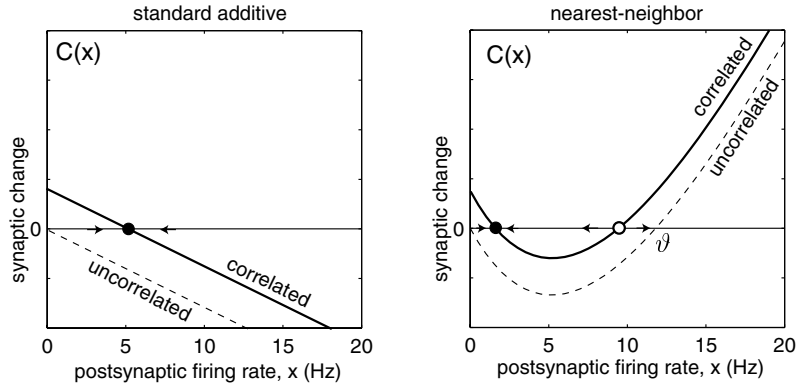


Figure 2: Stability diagram of two implementations of STDP rule.

$2/3xe^{2xt}$, $t < 0$. The resulting averaged curve has the form

$$C(x) = x \left(\frac{A_+}{\tau_+^{-1} + 2x} + \frac{2/3A_-}{\tau_-^{-1} + 2x} \right).$$

A.1 Stability of Various Implementations of STDP. Using the relationship between STDP and BCM rules, we can study some aspects of stability of the former using the latter. Consider the BCM equation

$$\dot{s} = C(x_{\text{post}}(t))x_{\text{pre}}(t),$$

where s is the synaptic weight and x_{pre} and x_{post} are the pre- and postsynaptic firing rates. The form of the function $C(x)$ provides invaluable information about the stability of synaptic modification even when there is no modification of the threshold. The additive, multiplicative, and suppression implementations of STDP without correlations result in negative functions $C(x)$ depicted in Figure 1D provided that $A_+ \tau_+ < |A_- \tau_-|$. In this case, the synaptic weight $s \rightarrow 0$ for any nonzero pre- or postsynaptic firing rate. This result confirms findings of Song et al. (2000) that uncorrelated firing of a postsynaptic neuron always results in synaptic depression provided that the depression area of STDP curves is greater than the facilitation area.

When the postsynaptic firing rate is small, the postsynaptic neuron starts to exhibit correlations with presynaptic neurons (Song et al., 2000). Therefore, to study the dynamics of s for small x_{post} , we need to consider correlated probability density, such as the one in Figure 1C. The classical additive implementation of STDP results in $C(x)$ having form A.3 depicted in Figure 1D and Figure 2, left. It has a nonzero globally stable equilibrium.

Similarly, the nearest-neighbor implementation of STDP with correlations results in

$$C(x) = x \left(A_+ \tau_+ + \frac{A_-}{\tau_-^{-1} + x} \right) + ac_0.$$

Here, the term ac_0 is due to correlations, as in the case of classical additive STDP considered above. The graph of this function is depicted in Figure 2, right, and it has a stable and unstable equilibrium. Small postsynaptic firing rates result in s converging to the stable equilibrium (the black circle in the figure), whereas large firing rates result in runaway dynamics of $s \rightarrow +\infty$. The stable equilibrium corresponds to a balance of potentiation, which is due to the correlations (coincidences) of pre- and postsynaptic firing, and depression, which occurs because the depression area of STDP is greater than the facilitation area.

If the term ac_0 increases, the stable equilibrium in Figure 2, left, persists. In contrast, the stable and unstable equilibria in Figure 2, right, approach and annihilate each other via saddle-node bifurcation. The dynamics of this implementation of STDP becomes unstable: any firing rate of a postsynaptic neuron results in potentiation. Thus, the classical implementation of STDP is unconditionally stable, whereas the nearest-neighbor implementation can be stable or unstable depending on how strong the correlations are.

Finally, notice that we consider the issue of stability assuming that the postsynaptic firing rate does not depend on the strength of the synapse s . The qualitative results persist when x_{post} is a monotone function of s (i.e., the stronger the synapse, the higher the postsynaptic rate). More research is needed to determine the stability of various implementations of STDP when x_{post} is a nonmonotone function of s (i.e., a very strong synaptic drive results in decreased input resistance and/or inactivation of Na conductances and, hence, a smaller firing rate).

Acknowledgments

We thank Jeff Krichmar, Anil Seth, and Jonathan Rubin for reading the first draft of the manuscript and making a number of useful suggestions. This research was supported by the Neurosciences Research Foundation and in part by a grant from the Alafi Family Foundation.

References

- Abarbanel, H. D. I., Huerta, R., & Rabinovich, M. I. (2002). Dynamical model of long-term synaptic plasticity. *PNAS*, *99*, 10132–10137.
- Bear, M. F., Cooper, L. N., & Ebner, F. F. (1987). A physiological basis for a theory of synapse modification. *Science*, *237*, 42–48.
- Bear, M. F., & Malenka, R. C. (1994). Synaptic plasticity: LTP and LTD. *Curr. Opin. Neurobiol.*, *4*, 389–399.

- Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, *18*, 10464–10472.
- 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. *PNAS*, *98*, 12772–12777.
- Debanne, D., Gahwiler, B. H., & Thomson, S. M. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol. (London)*, *507*, 237–247.
- Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*, *27*, 45–56.
- Froemke, R. C., & Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, *416*, 433–438.
- Kempler, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E*, *59*, 4498–4514.
- Kempler, R., Gerstner, W., & van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Comput*, *13*, 2709–2741.
- Kirkwood, A., Dudek, S. M., Gold, J. T., Aizenman, C. D., & Bear, M. F. (1993). Common forms of synaptic plasticity in the hippocampus and neocortex in vitro. *Science*, *260*, 1518–1521.
- Kirkwood, A., Rioult, M. G., & Bear, M. F. (1996). Experience-dependent modification of synaptic plasticity in visual cortex. *Nature*, *381*, 526–528.
- Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, *275*, 213–215.
- Philpot, B. D., Sekhar, A. K., Shouval, H. Z., & Bear, M. F. (2001). Visual experience and deprivation bidirectionally modify the composition and function of NMDA receptors in visual cortex. *Neuron*, *29*, 157–169.
- Senn, W., Markram, H., & Tsodyks, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation*, *13*, 35–67.
- Sjostrom, P. J., Turrigiano, G. G., & Nelson, S. B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, *32*, 1149–1164.
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neurosci.*, *3*, 919–926.
- van Rossum, M. C., Bi, G. Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.*, *20*, 8812–8821.