

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/12354727>

Competitive Hebbian learning through spike timing-dependent plasticity

Article in *Nature Neuroscience* · October 2000

DOI: 10.1038/78829 · Source: PubMed

CITATIONS

2,242

READS

3,299

3 authors, including:



[Sen Song](#)

Tsinghua University

69 PUBLICATIONS 5,723 CITATIONS

[SEE PROFILE](#)



[Kenneth D Miller](#)

Columbia University

122 PUBLICATIONS 9,579 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Optogenetics [View project](#)



Cell type-dependent function in the IC [View project](#)

Competitive Hebbian Learning Through Spike-Timing-Dependent Synaptic Plasticity

Sen Song, Kenneth D. Miller* and L. F. Abbott**

Volen Center for Complex Systems

and

Department of Biology

Brandeis University

Waltham MA 02254-9110

flamingo@volen.brandeis.edu, abbot@brandeis.edu

*Departments of Physiology and Otolaryngology,

Neuroscience Graduate Program,

W.M. Keck Center for Integrative Neuroscience,

Sloan Center for Theoretical Neurobiology

University of California

San Francisco, CA 94143-0444

ken@phy.ucsf.edu

** Corresponding author

March 31, 2000

Abstract

Hebbian models of development and learning require both activity-dependent synaptic plasticity and a mechanism that induces competition between different synapses. Recent experiments have characterized a form of long-term synaptic plasticity that depends on the relative timing of pre- and postsynaptic action potentials, which we call spike-timing-dependent plasticity (STDP). We show, in modeling studies, that this form of synaptic modification can automatically adjust synaptic strengths so that the postsynaptic neuron becomes more sensitive to presynaptic spike timing and operates in a balanced or irregular-firing regime. It has been argued that neurons *in vivo* operate in such a regime, and STDP may thus explain how the required level of excitation arises and is maintained. Despite being synapse specific, STDP generates competition between different synapses because they compete for control of the timing of postsynaptic action potentials. By combining synaptic modification and competition, STDP can serve as a mechanism for competitive Hebbian learning that does not require further assumptions or constraints on synaptic efficacies. This learning mechanism is not sensitive to overall input firing rates or variabilities, but selectively strengthens groups of synapses whose activities are correlated over short time periods.

Introduction

Hebbian learning, the development of neural circuits on the basis of correlated activity, relies on two critical mechanisms. The best known of these is activity-dependent synaptic modification along the lines proposed by Hebb¹. Equally important is a mechanism that forces different synapses to compete with one another, so that when some synapses to a given postsynaptic neuron are strengthened, others are weakened^{2,3}. For example, correlation-based rules of synaptic modification can provide a reasonable account of many aspects of development in visual cortex, but only when they are combined with constraints introduced to ensure competition⁴. While Hebbian synaptic modification has received support from experiments on long-term potentiation and depression^{5–8}, much less is known about the mechanisms that generate competition between synapses.

At first, it might appear that any mechanism that imposes competition among synapses must involve a global intracellular signal that reflects the state of many synapses. The constraints used in many models of Hebbian learning⁹, while not biophysically realistic, are based on this idea. Typically these constraints limit the sum of synaptic strengths received by a cell, or the mean activity of the cell. Experimental evidence concerning such global synaptic scaling is only now beginning to appear^{10–12}. Competition can also arise locally due to synaptic modification mechanisms that equilibrate at a pre-set level of total synaptic innervation or activity⁹. These can be static mechanisms, such as thresholds or negative input correlations^{9,13}, or dynamics mechanisms involving additional non-Hebbian synaptic growth or decay terms⁹ or shifts in the synaptic modification rule itself, as in

the sliding threshold of the BCM model¹⁴. Here we explore an entirely different local mechanism suggested by recent results on the effect of spike timing on long-term synaptic modification^{15–20} in which different synapses compete for control of the timing of postsynaptic action potentials. We show that the dependence of synaptic modification on spike timing provides a mechanism that can lead to competitive Hebbian learning without requiring global intracellular signaling or pre-set activity or synaptic efficacy levels.

Experimental evidence from a number of different preparations^{15–20} suggests that both the sign and degree of synaptic modification arising from repeated pairing of pre- and postsynaptic action potentials depend on their relative timing. In experiments on neocortical slices^{15,20}, hippocampal slice¹⁷ and cell¹⁸ culture, and *in vivo* studies of tadpole tectum¹⁹, long-term strengthening of synapses occurred if presynaptic action potentials preceded postsynaptic firing by no more than about 50 ms. Presynaptic action potentials that followed postsynaptic spikes produced long-term weakening of synapses. The largest changes in synaptic efficacy occurred when the time difference between pre- and postsynaptic action potentials was small, and there was a sharp transition from strengthening to weakening as this time difference passed through zero. We call this form of synaptic plasticity spike-timing-dependent plasticity (STDP).

Synaptic modification by STDP-like rules has been studied previously in models of temporal pattern recognition^{21,22}, temporal sequence learning^{23–25}, coincidence detection^{26,27}, navigation^{28–30}, and direction selectivity^{31–33}. Most of this work is based on earlier data regarding the spike-timing dependence of synaptic modification^{35–37} and, being less con-

strained, involved a range of different time scales for synaptic modification. These examples studied the utility of an STDP-like synaptic modification rule in solving a number of computational tasks. Here we focus instead on the competitive and stabilizing properties of STDP, in a model based on recent experimental data^{15–20}. The competitive nature of STDP has been noted^{26,19}, but not studied in detail previously. Stability of an STDP-like rule in combination with non-Hebbian plasticity has been studied in a linear firing-rate model³⁴, but we find qualitatively new behavior when the timing of individual action potentials and not merely firing-rates is taken into account. In this case, STDP modification alone can lead to stable configurations of synaptic conductances, subject only to a limit on the strengths of individual synapses. Furthermore, the competition and stability that arise in STDP-based synaptic modification yield a close link to ideas concerning the balance of excitation and inhibition and its effect on response variability^{38–50}. STDP automatically forces the postsynaptic neuron into a balanced regime in which it is sensitive to the timing of the presynaptic action potentials it receives. Such sensitivity provides the basis of the competition for the control of spike timing that stabilizes synaptic modification under STDP.

Spike-Timing-Dependent Synaptic Plasticity

The modeling studies we present are based on the spike-time-dependent synaptic modification rule illustrated in Figure 1. The curve drawn in Figure 1 determines the amount of synaptic modification arising from a single pre- and postsynaptic spike pair separated by

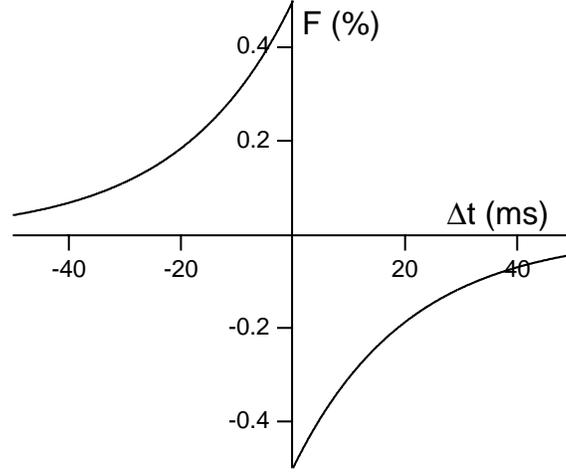


Figure 1: The STDP modification function. The change of the peak conductance at a synapse due to a single pre- and postsynaptic action potential pair is $F(\Delta t)$ times the maximum value \bar{g}_{\max} . Δt is the time of the presynaptic spike minus the time of the postsynaptic spike. In this figure, F is expressed as a percentage.

a time Δt . It is of the form

$$F(\Delta t) = \begin{cases} A_+ \exp(\Delta t/\tau_+) & \text{if } \Delta t < 0 \\ -A_- \exp(-\Delta t/\tau_-) & \text{if } \Delta t > 0, \end{cases} \quad (1)$$

where τ_+ and τ_- determine the ranges of pre- to postsynaptic interspike intervals over which synaptic strengthening and weakening occur. A_+ and A_- determine the maximum amounts of synaptic modification, which occur when Δt is close to zero.

The function $F(\Delta t)$ of equation 1 provides a reasonable approximation of the dependence of synaptic modification on spike timing seen in the experimental data. The experimental results suggest a value for τ_+ in the range of tens of milliseconds and, in the examples we present, we use $\tau_+ = 20$ ms. Data from some preparations indicate that the temporal window for synaptic weakening is roughly the same as that for synaptic strengthening^{15,18,19}, while other results reveal a larger window for synaptic weakening^{17,20}.

We have run simulations under both conditions. For the results we report here, we do not see a significant difference between the two cases, and we use $\tau_- = \tau_+ = 20$ ms throughout.

The amplitude of synaptic modification, which is controlled by the parameters A_+ and A_- , has been adjusted to reflect the modification due to a single pair of pre- and postsynaptic spikes by dividing the total modification measured experimentally for multiple spike pairs by the number of pairs. This assumes that the effects of individual spike pairs sum linearly. At least one contradictory effect has been reported, a dependence of synaptic strengthening on pairing frequency, including a threshold effect and frequency-dependent saturation¹⁵. Our model does not incorporate this finding, but we maintain presynaptic rates above the reported threshold frequency for synaptic strengthening¹⁵. In the simulations we show, we use $A_+ = 0.005$ (the value of A_- is discussed below), except in Figure 2F where we use $A_+ = 0.02$.

In our model, we make the important assumption that synaptic weakening through STDP is, overall, a slightly larger effect than synaptic strengthening²⁷. Specifically, stable competitive synaptic modification requires the integral of the function F to be negative, which assures that uncorrelated pre- and postsynaptic spikes produce an overall weakening of synaptic strength. A negative integral of F requires $A_- \tau_- > A_+ \tau_+$. The data are mixed on this issue. The results that report roughly equal time scales for synaptic strengthening and weakening^{15,18,19} indicate rough equality between the two effects and, in some cases, even suggest a slight dominance of strengthening over weakening. The data

showing a longer temporal window for synaptic weakening^{17,20} more clearly support the dominance of synaptic weakening over strengthening by STDP. In our simulations we use $A_-/A_+ = 1.05$, except for Figure 2D where A_-/A_+ varies.

In the model we study, \bar{g}_a denotes the peak synaptic conductance (the synaptic conductance immediately after an isolated presynaptic spike) at an excitatory synapse labeled by the integer a (with $a = 1, 2, \dots, N$). This conductance must always be positive, and is not allowed to exceed a maximum value \bar{g}_{\max} . A pre- and postsynaptic spike pair separated by a time interval Δt modifies the peak synaptic conductance by an amount $F(\Delta t)\bar{g}_{\max}$. The value $A_+ = 0.005$ thus corresponds to a change of 0.5% of the maximum synaptic strength per spike pair. If this modification rule would push the peak synaptic conductance beyond the allowed range $0 \leq \bar{g}_a \leq \bar{g}_{\max}$, \bar{g}_a is set to the appropriate limiting value. A scheme for implementing this modification rule is presented in the Methods section.

In our modeling studies, we examine how STDP acts on the excitatory synapses driving an integrate-and-fire model neuron with $N = 1000$ excitatory and 200 inhibitory synapses (see Methods). The excitatory synapses are activated by various types of spike trains: uncorrelated spike trains generated by independent Poisson processes at various rates, bursts of action potentials with different latencies, and partially correlated spike trains. The model neuron also receives inhibitory input consisting of Poisson spike trains at a fixed rate of 10 Hz. In the simulations, excitatory synapses are modified on the basis of their pre- and postsynaptic spike timing, while inhibitory synapses are held fixed.

Balanced Excitation

To function properly, a neuron must establish and maintain an appropriate level of excitation so that it can respond to its inputs by firing action potentials at reasonable rates. The firing statistics of cortical neurons suggest that this is achieved through a balance of the currents arising from excitatory and inhibitory synapses and intrinsic membrane conductances^{38–48}. As we will see, STDP provides a mechanism by which this balance can be established and maintained. This results in a state in which presynaptic action potentials can control the timing of postsynaptic spikes, so that competition between synapses can be realized.

To address this issue, we initially set the peak conductances of all the excitatory synapses of the model neuron to \bar{g}_{\max} , which produces a high firing rate. All the excitatory synapses to the model neuron received independent Poisson spike trains with the same average rate. After a period of adjustment, a steady-state condition was achieved in which the firing rate of the postsynaptic neuron and the distribution of peak synaptic conductances remained constant. Figures 2A and 2B show histograms of the resulting distributions of peak synaptic conductances for input firing rates of 10 and 40 Hz. Although all the synaptic conductances started with the same value, there is no stable equilibrium state with a uniform distribution of their values. Instead, most of the peak synaptic conductances have been pushed toward the limiting values of zero or \bar{g}_{\max} . For low input rates, more synapses approach the upper limit (Figure 2A), and for high input rates more are pushed toward zero (Figure 2B). This has the effect of keeping the total synaptic input to the neuron roughly constant, indepen-

dent of the presynaptic firing rates. The split between strong and weak synapses is also affected by the values of \bar{g}_{\max} (fewer strong synapses develop for larger \bar{g}_{\max}) and A_-/A_+ . The initial distribution of synaptic strengths has little effect on the final steady-state distribution as long as the postsynaptic neuron is initially firing action potentials.

Figure 2C shows the firing rate and coefficient of variation (CV) of the interspike intervals for the postsynaptic neuron once the synapses have reached their steady-state distribution. The CV is the standard deviation of the interspike intervals divided by their mean. STDP has a strong regulatory effect on the steady-state firing rate of the postsynaptic neuron, which increases by only about 1 Hz for each 5 Hz increase in the input firing rate. In contrast, if the peak synaptic conductances are held fixed in this model, the firing rate increase is over 100 Hz for a 5 Hz increase in the firing rate of the inputs. Of course, the long-term synaptic changes due to STDP take time to develop, so STDP only regulates the long-term average firing rate. The neuron thus remains highly sensitive to brief changes in the input firing rates.

The coefficient of variation of the postsynaptic spike train is fairly large and remarkably independent of the input firing rate (Figure 2C). This suggests that STDP regulates the variability of the postsynaptic response. The high degree of firing variability is primarily due to an overall balance between inhibitory and excitatory conductances in the model. A reasonable measure of this balance is the ratio of total inhibitory to excitatory current when the membrane potential is at the action-potential threshold. As shown in Figure 2D, STDP adjusts this ratio to be slightly greater than one over the entire range of presynaptic

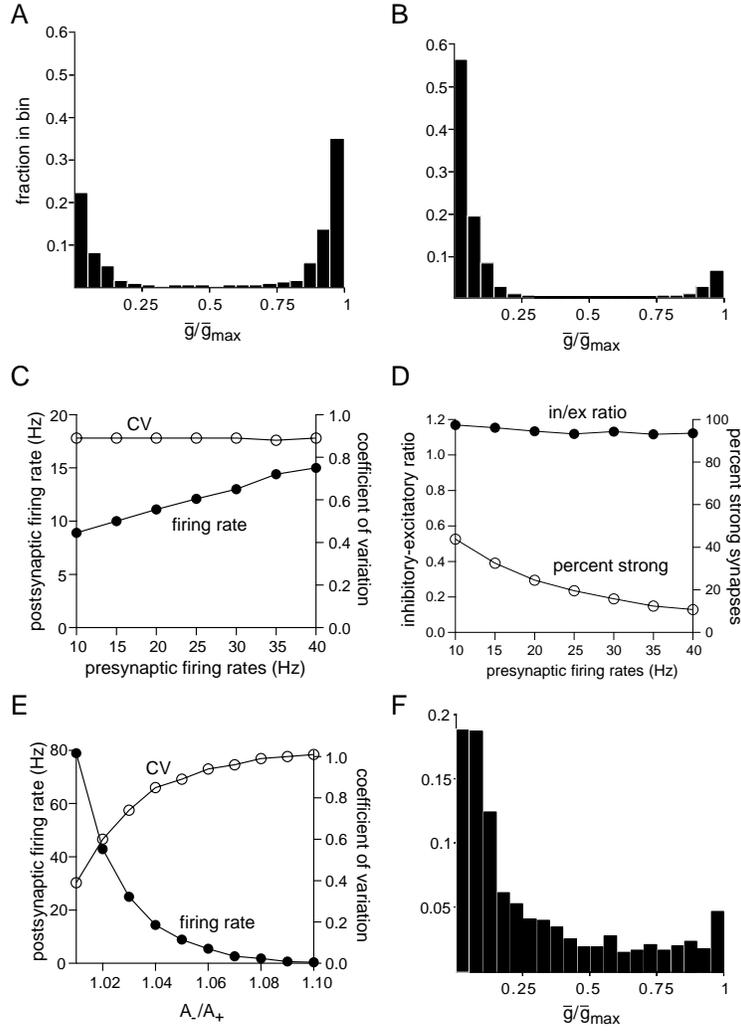


Figure 2: Balanced excitation and irregular firing produced by STDP. A) Histogram of the fraction of synapses taking different peak conductance values ranging from zero to \bar{g}_{\max} . For an input rate of 10 Hz, the peak synaptic conductances tend to the limiting values, but more are near \bar{g}_{\max} than near zero. B) Same as A, but for an input rate of 40 Hz. Now more peak conductances are near zero than near \bar{g}_{\max} . C) The postsynaptic firing rate and CV of the postsynaptic interspike intervals for different input firing rates. D) The ratio of inhibitory to excitatory currents at threshold and the percentage of strong synapses ($\bar{g} \geq 0.8\bar{g}_{\max}$) for different presynaptic firing rate. The leakage conductance is included as an inhibitory current in this ratio because it acts to hyperpolarize the neuron. E) The postsynaptic firing rate and CV of the postsynaptic interspike intervals for input firing rates of 10 Hz but different values of A_-/A_+ , the ratio of the amplitudes of maximal synaptic weakening and strengthening. F) Same as A, but with \bar{g}_{\max} 2.33 times larger and the synaptic modification per spike pair four times larger ($\bar{g}_{\max} = 0.035$, $A_+ = 0.020$, $A_- = 0.021$). The larger value of \bar{g}_{\max} forced more synapses to lower conductance values, while the higher modification rate filled in the distribution.

firing rates considered. This indicates a balanced condition in which, on average, inhibitory effects are slightly dominant at threshold. An additional contribution to firing variability comes from the reduction in the number of strong synapses for high input rates. Figure 2D also shows the number of strong synapses for different presynaptic firing rates, where we define strong synapses as those with $\bar{g} \geq 0.8\bar{g}_{\max}$. For the value of \bar{g}_{\max} we used, roughly half the synapses are strong for a 10 Hz presynaptic rate. The number of strong synapses drops to 10% when the presynaptic rates are set to 40 Hz. In all cases, the balance between inhibition and excitation is the dominant source of variability, but the reduction in the number of strong inputs also contributes when the presynaptic firing rates are high.

Both the firing rate and the coefficient of variation of the postsynaptic neuron depend on the ratio A_-/A_+ , as seen in Figure 2E. If this ratio is slightly larger than one, the firing rate of the postsynaptic neuron is maintained in a reasonable range, and the CV is close to one, indicating an irregular postsynaptic spike train.

The histograms in Figures 2A and B indicate that synaptic conductances tend to be pushed close to the upper and lower limits of their allowed range by the STDP modification rule we are using. This results in a bimodal distribution. A more continuous distribution arises if the degree of synaptic modification per spike pair is increased. For example, Figure 2F shows an example where the equilibrium distribution of synaptic conductances is roughly exponential, except for the a small excess near $\bar{g} = \bar{g}_{\max}$. This more closely matches data on distributions of synaptic strengths estimated from recorded spontaneous synaptic (mini) potentials⁵¹.

Softky and Koch^{38,39} first pointed out that it is difficult to obtain CV values as large as those seen *in vivo* ($CV \approx 1$) in an integrate-and-fire model receiving many independent presynaptic inputs. Correlations of input spike timing, such as synchronization, could contribute to increased CV's⁴⁰. However, a number of authors have noted that a high degree of variability and CV's near one will also arise if the excitatory inputs to a neuron are balanced relative to the inhibitory synaptic and membrane currents^{41–48}. The critical condition is that the mean input to the neuron should only be sufficient to charge the membrane up to a point below, or only slightly above, the threshold for action potential generation, so that spike times are determined primarily by positive fluctuations in input rather than by the mean input. This condition may also play a critical role in explaining a number of fundamental visual cortical response properties⁴⁹. STDP automatically achieves this balanced state for a wide range of input firing rates.

The reason that STDP achieves a balanced state can be understood from basic response characteristics of a neuron integrating many inputs. Such a neuron can operate in two different modes with distinct spike-train statistics and input-output correlations^{44,46}. When excitation is strong, as at the beginning of our simulations, the mean input to the neuron would bring it well above threshold if action potentials were blocked, so the neuron operates in an input-averaging or regular-firing mode. The action potential sequences produced in this mode are significantly more regular than the presynaptic spike trains. The interspike intervals of the postsynaptic responses depend on the total synaptic input, but the absolute timing of individual action potentials is fairly independent of fluctuations in this input. As

a result, there are roughly equal numbers of presynaptic action potentials before and after each postsynaptic spike^{46,50}. This is seen in Figure 3A, which shows the relative frequency of pre- and postsynaptic spike pairs separated by different intervals of time for conditions at the beginning of our simulations. This figure shows only a small excess of presynaptic spikes occurring just before a postsynaptic action potential. The STDP curve from Figure 1 has been overlaid onto Figures 3A and 3B. As we have noted, the area under the synaptic weakening portion of this curve is greater than the area under the strengthening part. Initially in our simulations, there is an overall weakening of the excitatory synapses because the small excess of presynaptic spikes occurring prior to postsynaptic action potentials is not large enough to overcome the excess of synaptic weakening imposed by the STDP rule (Figure 3A).

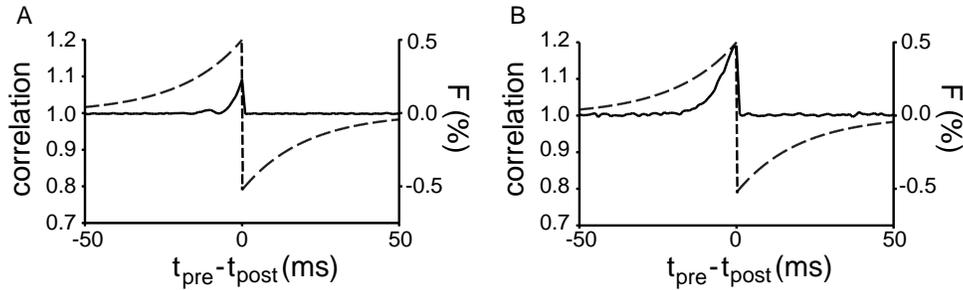


Figure 3: Correlation between pre- and postsynaptic action potentials before and after STDP. The solid curves indicate the relative probability of a presynaptic spike occurring at time t_{pre} when a postsynaptic spike occurs at time t_{post} . A correlation of one is the value due solely to chance occurrences of such pairs. The dashed curves show the STDP modification function from Figure 1. The time-integral of the product of the synaptic modification curve and the correlation function determines whether, on average, the synapses will be strengthened or weakened. A) At the beginning of our simulations, when all peak synaptic conductances are set to their maximal value, there is only a small excess of presynaptic spikes prior to a postsynaptic action potential. B) At the end of the simulations, when STDP has established a steady-state distribution of conductances, there is a larger excess of presynaptic spikes prior to a postsynaptic action potential. In the steady-state, this excess compensates for the asymmetry in the STDP modification curve, *i.e.*, for the fact that $A_-/A_+ > 1$.

As the excitatory synapses get weaker, the postsynaptic neuron enters a balanced mode of operation in which it generates a much more irregular sequence of action potentials. Because the total synaptic input is, on average, near or sub-threshold in the balanced mode, the postsynaptic neuron fires primarily in response to statistical fluctuations in the total input. Such fluctuations occur irregularly, so the postsynaptic firing pattern is highly variable. The absolute timing of action potentials in the irregular firing mode is related to temporal features of the synaptic inputs, namely the fluctuations, and therefore there tend to be more excitatory presynaptic action potentials before than after a postsynaptic response^{45,46,50,52}. As a result, there is a higher degree of correlation between pre- and postsynaptic action potentials in this irregular-firing mode than in the regular-firing mode. Figure 3B shows a larger excess of presynaptic spikes before a postsynaptic action potential than Figure 3A, because it was constructed after STDP had modified the excitatory synapses. The STDP rule achieves a steady-state distribution of peak synaptic conductances when the excess of presynaptic action potentials prior to postsynaptic firing compensates for the asymmetry in areas under the positive and negative portions of the STDP modification curve⁵⁰. STDP thus forces a cell into a stable equilibrium in which it is both reasonably active and reasonably sensitive to input fluctuations.

Latency Reduction

For uncorrelated stochastic presynaptic spike trains, chance determines whether a given synapse will ultimately become weak or strong through STDP. When the presynaptic inputs are correlated in various ways, the fate of individual synapses is controlled in a more

systematic manner. The STDP rule strengthens synapses that fire prior to a postsynaptic spike and weakens those that fire later. If, for example, the postsynaptic neuron receives a barrage of excitatory synaptic input, STDP will strengthen short-latency excitatory inputs while weakening those with long latencies. The ultimate effect of this synaptic modification is to make the postsynaptic neuron respond more quickly.

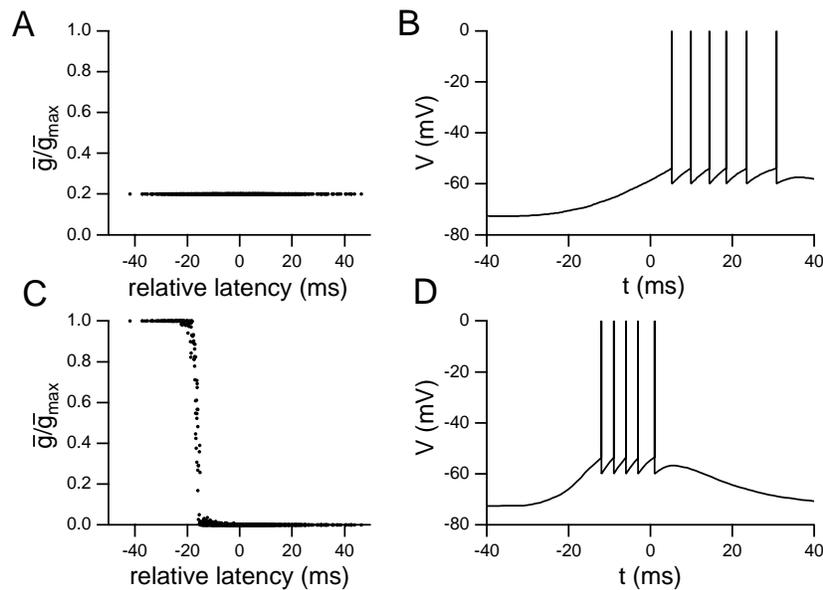


Figure 4: Reduction of latency by STDP. A) The initial peak synaptic conductances plotted as a function of the relative latency of their synaptic inputs. B) The initial postsynaptic response to a barrage of excitatory input with burst onset for each synapse occurring at the time of its relative latency. C) The steady-state peak synaptic conductances plotted as a function of the relative latency of the synaptic input. Short-latency synapses have been strengthened and long-latency synapses have been weakened. D) The response of the postsynaptic neuron to the same input barrage as in B, but after STDP has modified the peak synaptic conductances as in C.

In Figure 4, the inputs to the model neuron were silent except for isolated events represented by bursts of spikes with a Poisson distribution at 100 Hz for 20 ms. Different synapses were not activated precisely synchronously during these events. Instead, each synapse was assigned a relative latency chosen randomly from a Gaussian distribution with

a mean of zero and a standard deviation of 15 ms. The burst of action potentials at a given synapse occurred at a time given by the sum of its relative latency and the mean time of the event.

Initially, all the synapses were set to the same strength of $0.2\bar{g}_{\max}$ (Figure 4A). This produced a response in the postsynaptic cell that began shortly after the time marked zero, which indicates the mean event time, and lasted for about 25 ms (Figure 4B). The input events were then repeated periodically until the STDP rule had established a fixed distribution of peak synaptic conductances. In Figure 4C, the resulting steady-state conductances are plotted as a function of the relative latency of the synapses. Short-latency inputs have been strengthened to the maximum allowed level, \bar{g}_{\max} , while synapses with longer latencies have been weakened to zero. This produces a quicker response in the postsynaptic neuron, which now fires almost 20 ms earlier than it did originally (Figure 4D).

Correlation-Based Hebbian Modification

Factors that enhances the ability of a given synapse to rapidly evoke a postsynaptic response will lead to its strengthening through STDP. Correlating different synaptic inputs so that they are more likely to arrive together in a cluster is an effective way of increasing their ability to evoke postsynaptic action potentials. By cooperatively generating action potentials, such a cluster of synapses can grow stronger, while weakening other synapses that are not part of the cluster. To study this effect, we generated input spike trains at rates that were correlated across synapses (see Methods), and examined how they were affected by STDP.

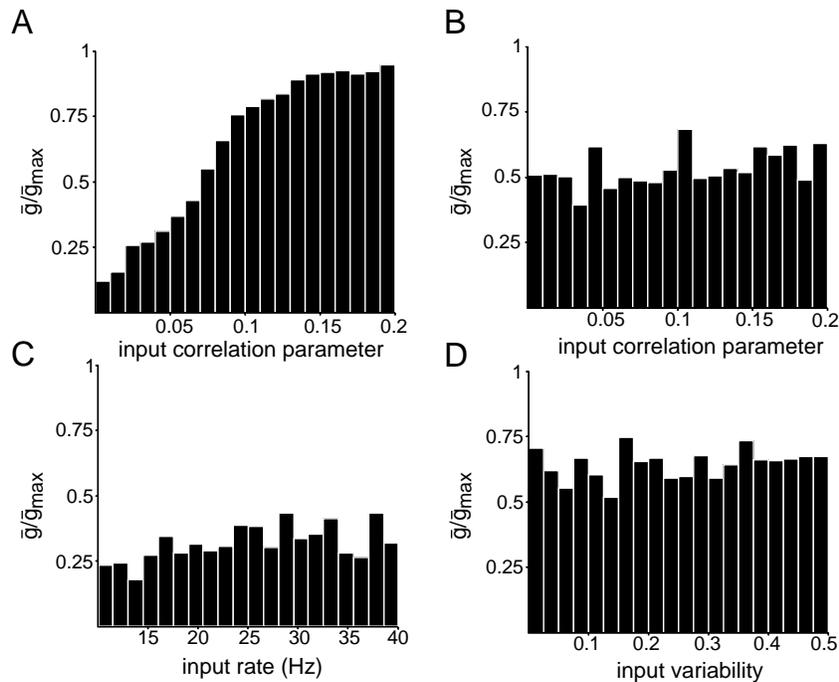


Figure 5: Effects of correlation, firing rate, and variability on peak synaptic conductances. The histograms show the average peak synaptic conductances within 20 bins. These values are the results of averaging bimodal distributions of synaptic strengths within each bin. Different statistical parameters describing the synaptic inputs are binned and displayed on the horizontal axes. A) The synaptic inputs have correlation parameters ranging from zero to 0.2 and have been binned on this basis. The degree correlation between any two inputs is determined by the product of their correlation parameters. The correlation time constant is 20 ms, and the degree of correlation of a synapse has a strong effect on its peak conductance. B) Same as A, but with a correlation time constant of 200 ms. No effect of correlation on synaptic strength is observed. C) The synaptic inputs have different firing rates ranging from 10 to 40 Hz, and this range has been binned. No strong effect of rate on synaptic strength is observed. D) The synaptic inputs have distributions of input firing rates with different standard deviations (labeled input variability) ranging from 0 to 0.5 in units of the mean rate. Binning is based on this input variability. No effect of variability on synaptic strength is observed. In this example, $\tau_c = 20$ ms as in A.

In Figure 5A and 5B, the presynaptic firing rates were generated to have a correlation function that decayed exponentially with time and varied in amplitude across the population of synapses. Specifically, the presynaptic firing rates r_a for $a = 1, 2, \dots, N$ had the correlation function, $\langle r_a(t)r_b(t') \rangle = \bar{r}^2 + \bar{r}^2(\sigma^2\delta_{ab} + (1 - \delta_{ab})c_a c_b) \exp(-|t - t'|/\tau_c)$, where the angle brackets represent an average over the ensemble of rates, $\bar{r} = 10$ Hz, $\sigma = 0.5$,

and c_a , which we call the correlation parameter, varied from zero to 0.2 uniformly across the 1000 excitatory synapses ($c_a = 0.2(a - 1)/(N - 1)$). The variation in c_a values produces a gradation of correlation across the population of synapses. Pairs of synapses with large c values are more highly correlated than synapses with smaller c values. In Figure 5A and 5B, the range of input correlation parameters is divided into 20 bins, and the average values of the steady-state peak synaptic conductances for synapses within these bins are plotted in the form of a histogram. When the correlations decays rapidly, as in Figure 5A ($\tau_c = 20$ ms), there is a marked tendency for more correlated synapses to become stronger. This trend disappears for larger correlation times, as seen in Figure 5B ($\tau_c = 200$ ms). A comparison of Figures 5A and B shows that synapses in Figure 5A have been both strengthened by correlation and weakened by competition. To be strengthened, a group of inputs must fire together long enough to generate a postsynaptic action potential, but must then stop firing so that they are not subsequently weakened. As a result, correlations have a large effect when the correlation time constant is approximately equal to the time constants τ_+ and τ_- that govern the time scales for STDP³⁴.

Although the degree of strengthening produced by STDP is sensitive to correlations, it is not strongly affected by other properties of the presynaptic spike trains. For Figure 5C, the input firing rates were time-independent and uncorrelated, but they varied from 10 to 40 Hz uniformly across the population of synapses ($r_a = 10 + 30(a - 1)/(N - 1)$ Hz). Figure 5C shows average steady-state peak synaptic conductances in bins that specify different ranges of input firing rates. There is little tendency for synapses firing at either

faster or slower rates to be preferentially strengthened or weakened by STDP. Higher firing rates increase the speed at which synaptic modification occurs, but they do not otherwise affect the final equilibrium distribution of maximal synaptic conductance values produced by STDP.

Figure 5D shows that the steady-state peak synaptic conductances are also insensitive to the degree of variability of the presynaptic input. The firing rates for this figure were generated so that $\langle r_a(t)r_a(t') \rangle = \bar{r}^2 + \bar{r}^2\sigma_a^2 \exp(-|t - t'|/\tau_c)$, where $\bar{r} = 10$ Hz and σ_a varied from zero to 0.5 uniformly over the population of synapses ($\sigma_a = 0.5(a - 1)/(N - 1)$). This means that the standard deviation of the input firing rates was different for different synapses. Figure 5D shows the binned average peak synaptic conductances as a function of the standard deviations of their inputs. There is no tendency for synapses with either more or less variable firing rates to be preferentially strengthened or weakened by STDP.

The basic result of these studies is that STDP is insensitive to the average rate or degree of variability of a given synaptic input. It is, however, strongly affected by correlations between different synaptic inputs, provided that they decay rapidly enough as a function of time. Synapses with strong, rapidly decaying temporal correlations will be strengthened as a cluster and will suppress other synapses that are uncorrelated or have temporal correlations that last over longer time periods. STDP thus shows the basic feature of Hebbian learning, the strengthening of correlated groups of synapses, while displaying the desirable features of firing-rate independence and stability, and a novel dependence on correlation decay time.

Discussion

Although Hebbian synaptic plasticity is a powerful concept, it suffers from a number of problems. First, synapses are modified whenever correlated pre- and postsynaptic activity occurs. Such correlated activity can occur purely by chance, rather than reflecting a causal relationship that should be learned. To correct for this, neural network models often use a covariance rather than correlation-based synaptic modification rule⁵³. However, such a rule cannot, in general, achieve competition between synapses. This illustrates a second problem of purely Hebbian modification; it is not competitive, so constraints must be added to obtain interesting results. STDP appears to solve both of these problems. Accidental, non-causal coincidences weaken synapses if, as we have assumed, the integral of the synaptic modification function is negative. Competition arises in a novel way, not due to a global signaling or growth factor, or to an artificially imposed balance of nonspecific synaptic decay and growth terms, but rather through a competition for control of the timing of postsynaptic action potentials. Inputs that consistently predict a postsynaptic response become the strongest inputs to the neuron. Causality is a key element of STDP. As Hebb suggested¹, synapses are only strengthened if their presynaptic action potentials precede, and thus could have contributed to, the firing of the postsynaptic neuron.

STDP automatically leads to a balanced, irregular-firing state in which pre- and postsynaptic spike times are causally correlated. This result depends crucially on considering spike timing and not merely firing rates. In a linear, firing-rate model of STDP³⁴, the cross-correlation between pre- and postsynaptic firing does not change shape with changes in

synaptic efficacy, as it does in Figure 3. This change in shape, in which causal correlations grow stronger relative to acausal correlations as overall synaptic efficacy decreases, is the key to the stabilizing and competitive effects of STDP we have presented.

STPD regulates both the rate and the coefficient of variation of postsynaptic firing over a wide range of input rates. This represents a homeostatic regulatory function of STDP, which is surprising given that, like the Hebb rule, it is destabilizing at individual synapses. STDP also differentially strengthens the shortest-latency inputs evoked by a stimulus. Within the STDP time window, synapses are strengthened in proportion to how predictive they are of events that lead to postsynaptic spiking. There is some experimental evidence suggesting that the reduction of latency illustrated in Figure 4 occurs *in vivo*. A phenomenon analogous to the reduction of latency discussed here predicts that, when a rat moves through a particular region, place cells active for that region should fire earlier after the rat has repeatedly traversed the area^{24,28,29}. This effect has been observed experimentally^{54,30}.

The model of STDP we used involves a number of assumptions and simplifications. We assumed that the effects of spike pairs sum linearly in a time-independent manner, but this is contradicted by some data¹⁵. Our results do not depend strongly on this assumption, but it is impossible to estimate the effects of changing it until the correct alternative is identified by experimental work. We have also ignored delays of several minutes between pairing of pre- and postsynaptic spikes and the resultant induction of synaptic modification that are suggested by experiments¹⁵. If the effect is merely a delay, this has no impact

on our results. If, on the other hand, the process acts as a low-pass filter on the temporal dynamics of weight change (averaging the effects of STDP over a long period of time and changing weights according to this average), this could have a more significant impact. We have re-run our simulations assuming such a low-pass filtering effect. We observed no changes in our results except for the case of Figure 2F, in which individual spike pairings caused larger changes than in the other examples. In this case, the impact of these larger changes is damped by the long-term averaging.

Two other assumptions are more critical. First, stability requires that synaptic weakening by STDP dominates over synaptic strengthening. If this is not true, the results we have reported might nevertheless arise from a combination of STDP and homosynaptic long-term depression (weakening of presynaptic inputs that fire in the absence of a postsynaptic spike^{7,8}). As long as STDP strengthens causally effective inputs, while STDP and/or other forms of long-term plasticity more strongly weaken causally ineffective inputs, the basic results found here should apply. Finally, our model involved hard bounds that kept the synaptic conductances from going below zero or above a maximum value. An alternative is to use soft bounds that smoothly reduce the magnitude of synaptic modification to zero as the limits are approached. Such soft bounds have a much stronger impact on the equilibrium distribution of synaptic conductances than the hard bounds we have used and tend to be highly restrictive. We have not considered them because, in general, they greatly reduce the effectiveness of any synaptic plasticity rule.

STDP may modify both the absolute strength of a synapse and its short-term synaptic

plasticity properties, an effect which has been called synaptic redistribution⁵⁵. We have run simulations in which we couple the strengthening and weakening of synapses through STDP to the degree of synaptic depression exhibited by the synapse in a manner consistent with synaptic redistribution. While this does not change the results we report, it does reveal an interesting interplay between long-term timing-dependent plasticity and short-term plasticity. The most effective way to strengthening a synapse with STDP is to have it release transmitter before a postsynaptic spike and then to stop releasing so that it will not be weakened by subsequent releases after the postsynaptic activity. A high degree of synaptic depression, which is a feature of strong synapses in the redistribution scheme⁵⁵, assures that this occurs. STDP that acts to modify release probability and change the degree of synaptic depression is thus extremely competitive and effective at driving individual synapses to strong or weak limits.

STDP, while making an important and novel contribution to competition, probably cannot be the sole source of plasticity in Hebbian learning situations. Like any other Hebbian modification rule, STDP cannot strengthen synapses in the absence of postsynaptic firing. If for some reason the excitatory synapses to a neuron are too weak to make it fire, STDP cannot rescue them. A non-Hebbian mechanism, such as synaptic scaling^{10–12}, may serve this function instead. Furthermore, in the present implementation of STDP, two input sets that never fire within 100 msec of each other will generate STDP independently and thus will not compete. Experiments suggest that competition occurs even between input sets whose firings are always separated by seconds^{56,57}. Such a result could arise if the STDP

temporal window for synaptic weakening had a long enough tail, if STDP were supplemented or replaced by sufficiently strong heterosynaptic long-term depression⁵⁸, or due to competition induced by synaptic scaling.

The size of the temporal windows over which synaptic strengthening and weakening occur is critical in determining the effects of STDP. It would seem highly advantageous for window sizes to be different in various brain regions, to be modified during stages of development, and perhaps to be dynamically adjustable over shorter time scales as well. This would allow STDP to stay compatible with relevant input correlations. STDP appears to be NMDA-dependent^{15–19}, and NMDA subunit substitution might provide a mechanism for adjusting its time course. For example, the developmental transition from a predominance of NR2B to NR2A subunits leads to a faster decay time of NMDA-receptor-mediated currents⁵⁹. This might be associated with a modification of the STDP window⁶⁰. It will be interesting to see if evidence of such variability in spike-timing plasticity windows is revealed in future experiments.

Methods

The membrane potential of the integrate-and-fire model we use is determined by

$$\tau_m \frac{dV}{dt} = V_{\text{rest}} - V + g_{\text{ex}}(t)(E_{\text{ex}} - V) + g_{\text{in}}(t)(E_{\text{in}} - V). \quad (2)$$

with $\tau_m = 20$ ms, $V_{\text{rest}} = -70$ mV, $E_{\text{ex}} = 0$ mV, and $E_{\text{in}} = -70$ mV. In addition, when the membrane potential reaches a threshold value of -54 mV, the neuron fires an action potential, and the membrane potential is reset to -60 mV (parameters take from

reference [?]). The synaptic conductances g_{ex} and g_{in} are measured in units of the leakage conductance of the neuron and are thus dimensionless, as are the related peak conductances.

Upon arrival of a presynaptic action potential at excitatory synapse a , $g_{\text{ex}}(t) \rightarrow g_{\text{ex}}(t) + \bar{g}_a$. When an action potential arrives at an inhibitory synapse, $g_{\text{in}}(t) \rightarrow g_{\text{in}}(t) + \bar{g}_{\text{in}}$, where \bar{g}_a and \bar{g}_{in} are the peak synaptic conductances. Otherwise, both excitatory and inhibitory synaptic conductances decay exponentially,

$$\tau_{\text{ex}} \frac{dg_{\text{ex}}}{dt} = -g_{\text{ex}} \quad \text{and} \quad \tau_{\text{in}} \frac{dg_{\text{in}}}{dt} = -g_{\text{in}}. \quad (3)$$

We have taken $\tau_{\text{ex}} = \tau_{\text{in}} = 5$ ms, $\bar{g}_{\text{in}} = 0.05$, and $\bar{g}_{\text{max}} = 0.015$ (except for Figure 4, where $\bar{g}_{\text{max}} = 0.02$ and Figure 2D, where $\bar{g}_{\text{max}} = 0.035$). For a $100 \text{ M}\Omega$ input resistance, $\bar{g}_{\text{max}} = 0.015$ corresponds to a peak synaptic conductance of 150 pS .

The synaptic modification given by Figure 1 is realized in the model using $N + 1$ functions, $M(t)$ and $P_a(t)$, for $a = 1, 2, \dots, N$. These are initially set to zero, and decay exponentially,

$$\tau_- \frac{dM}{dt} = -M \quad \text{and} \quad \tau_+ \frac{dP_a}{dt} = -P_a. \quad (4)$$

$M(t)$ is used to decrease synaptic strength. Every time the postsynaptic neuron fires an action potential, $M(t)$ is decremented by an amount A_- . If synapse a receives a presynaptic action potential at time t , its maximal conductance parameter is modified according to $\bar{g} \rightarrow \bar{g} + M(t)\bar{g}_{\text{max}}$. If this makes $\bar{g}_a < 0$, \bar{g}_a is set to zero. $P_a(t)$ is used to increase the strength of synapse a . Every time synapse a receives an action potential, $P_a(t)$ is incremented by an amount A_+ . If the postsynaptic neuron fires an action potential at time t , \bar{g}_a is modified according to $\bar{g}_a \rightarrow \bar{g}_a + P_a(t)\bar{g}_{\text{max}}$. If this makes $\bar{g}_a > \bar{g}_{\text{max}}$, \bar{g}_a is set to

\bar{g}_{\max} .

To generate the presynaptic firing rates for Figure 5, we choose intervals of time from an exponential distribution with mean interval τ_c . For every interval, we generated $N + 1$ random numbers, y and x_a for $a = 1, 2, \dots, N$ from Gaussian distributions with zero mean and standard deviation one and σ_a respectively, where $\sigma_a = \sqrt{\sigma^2 - c_a^2}$. At the start of each interval, the firing rate for synapse a was then set to $r_a = \bar{r}(1 + x_a + c_a y)$, and it was held at this value until the start of the next interval.

To ensure absence of dependence of our results on initial conditions, we have rerun multiple trials of the simulations of starting from a different randomly-generated sets of initial synaptic weights (for example, initial weights chosen from a distribution uniform between 0 and \bar{g}_{\max}), and using different Poisson spike trains each time. There was no detectable change in results. After convergence, the variability in CV and output rate between trials was indistinguishable from that seen in measurements within a trial. There is always a small degree of variability over time after a simulation has converged because statistics are gathered over a finite time, inputs are stochastic, and individual synapses continually change their values although the overall distribution does not appreciably vary. We consider the synaptic distributions to have converged when the output firing rate has stopped changing in a systematic manner. This occurs in about 100 seconds of simulated time. Stability has been checked in some simulations for as long as 100 hours of simulated time, and we have never seen appreciable changes in output rate or CV once convergence is reached. To be assured of convergence, all presented data were collected only after 1000 seconds of

simulated time.

References

1. Hebb, DO (1949) *The Organization of Behavior: A Neuropsychological Theory*. New York:Wiley.
2. Guillery, RW (1972) Binocular competition in the control of geniculate cell growth. *J. Comp. Neurol.* **144**:117-130.
3. Miller, KD (1996) Synaptic Economics: Competition and cooperation in synaptic plasticity. *Neuron* **17**:371-374.
4. Miller, KD (1996) Receptive fields and maps in the visual cortex: Models of ocular dominance and orientation columns. In E Domany, JL van Hemmen & K Schulten, editors, *Models of Neural Networks, III*. New York:Springer-Verlag, 55-78.
5. Bliss TVP, Collingridge GL (1993) A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* **361**:31-39.
6. Malenka, RC & Nicoll, RA (1993) NMDA-receptor-dependent synaptic plasticity: Multiple forms and mechanisms. *Trends Neurosci.* **16**:521-527.
7. Artola A, Singer W (1993) Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci.* **16**:480-487.

8. Bear, MF & Malenka, RC (1994) Synaptic plasticity: LTP and LTD. *Curr. Opin. Neurobiol.* **4**:389-399.
9. Miller, KD & MacKay, DJC (1994) The role of constraints in Hebbian learning. *Neural Computation* **6**:100-126.
10. Turrigiano, GG, Leslie KR, Desai, NS, Rutherford LC & Nelson, SB (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* **391**:892-896.
11. Davis, GW & Goodman, CS (1998) Synapse-specific control of synaptic efficacy at the terminals of a single neuron. *Nature* **392**:82-86.
12. O'Brien, RJ, Kamboj, S, Ehlers, MD, Rosen, KR, Fischbach, GD & Huganir, RL (1998) Activity-dependent modulation of synaptic AMPA receptor accumulations. *Neuron* **21**:1067-1078.
13. Linsker, R (1986) From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proc. Natl. Acad. Sci. USA* **83**:7508-7512.
14. Bienenstock, EL, Cooper, LN & Munro, PW (1982) Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* **2**:32-48.
15. Markram H, Lubke J, Frotscher M & Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**:213-215.

16. Bell CC, Han VZ, Sugawara Y & Grant K (1997) Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* **387**:278-281.
17. Debanne D, Gahwiler, BH & Thompson, SM (1998) Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol.* **507**:237-247.
18. Bi G-q & Poo M-m (1998) Activity-induced synaptic modifications in hippocampal culture: dependence on spike timing, synaptic strength and cell type. *J. Neurosci.* **18**:10464-10472.
19. Zhang LI, Tao, HW, Holt CE, Harris WA & Poo M-m (1998) A critical window for cooperation and competition among developing retinotectal synapses. *Nature* **395**:37-44.
20. Feldman, DE (1999) LTP and LTD induced by action-potential-epsp pairing at vertical inputs to layer II/III pyramids in rat somatosensory cortex. *Soc. Neurosci. Abst.* **25**:223.
21. Herz, A. V. M., Sulzer, B., Kühn, R., and van Hemmen, J. L. (1989) Hebbian learning reconsidered: Representation of static and dynamic objects in associative neural nets. *Biol. Cybern.* **60**, 457-467.
22. Gerstner, W., Ritz, R., and van Hemmen, J. L. (1993) Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.* **69**, 503-515.

23. Minai, AA & Levy, WB (1993) Sequence learning in a single trial. *INNS World Congress of Neural Networks II*:505-508.
24. Abbott, LF & Blum, KI (1996) Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral Cortex* **6**:406-416.
25. Roberts, PD (1999) Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. *J. Computational Neurosci.* **7**:235-246.
26. Gerstner, W, Kempter, R, van Hemmen, JL & Wagner, H (1996) A neuronal learning rule for sub-millisecond temporal coding. *Nature***383**:76-78.
27. Gerstner, W, Kempter, R, van Hemmen, JL & Wagner, H (1997) A developmental learning rule for coincidence tuning in the barn owl auditory system. In Bower, J (editor) *Computational Neuroscience* Plenum:New York. pp. 665-669.
28. Blum, KI & Abbott, LF (1996) A model of spatial map formation in the hippocampus of the rat. *Neural Comp.* **8**:85-93.
29. Gerstner, W & Abbott, LF (1997) Learning navigational maps through potentiation and modulation of hippocampal place cells. *J. Computational Neurosci.* **4**:79-94.
30. Mehta, MR, Quirk, MC & Wilson, M (2000) Experience dependent asymmetric shape of hippocampal receptive fields. *Neuron* (in press).
31. Rao, R & Sejnowski, TJ (2000) Predictive sequence learning in recurrent neocortical circuits. In Solla, SA, Leen, TK & Muller K-b (editors) *Advances in Neural*

Information Processing Systems 12 MIT Press:Cambridge MA.

32. Buchs, NJ, Reutimann, J & Senn, W (1999) Learning direction selectivity through adaptation of the vesicle release probability. *Soc. Neurosci. Abst.* 25:2259.
33. Mehta, MR & Wilson, M (2000) From hippocampus to V1: Effect of LTP on spatio-temporal dynamics of receptive fields. In Bower, J. ed. *Computational Neuroscience, Trends in Research 1999*. Elsevier:Amsterdam.
34. Kempter R, Gerstner W & van Hemmen JL (1999) Hebbian learning and spiking neurons. *Phys. Rev. E* **59**:4498-4514.
35. Levy, W. B. and Steward, D. (1983) Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neurosci.* **8**, 791-797.
36. Gustafsson, B., Wigstrom, H., Abraham, W. C., and Huang, Y. -Y. (1987) Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.* **7**, 774-780.
37. Debanne, D., Gahwiler, B. H., and Thompson, S. M. (1994) Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat hippocampus *in vitro* *Proc. Natl. Acad. Sci. USA* **91**, 1148-1152.
38. Softky, WR & Koch, C (1992) Cortical cells should spike regularly but do not. *Neural Computation* **4**:643-646.

39. Softky, WR & Koch, C (1994) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience* **13**:334-350.
40. Stevens, CF & Zador, AM (1998) Input synchrony and the irregular firing of cortical neurons. *Nature Neuroscience*, **1**:210-7.
41. Shadlen, MN & Newsome, WT (1994) Noise, neural codes and cortical organization. *Current Opinion in Neurobiology* **4**:569-579.
42. Shadlen, MN & Newsome, WT (1998) The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *Journal of Neuroscience* **18**:3870-3896.
43. Tsodyks, M & Sejnowski, TJ (1995) Rapid switching in balanced cortical network models. *Network* **6**:1-14.
44. Troyer, TW & Miller, KD (1997a) Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comp.* **9**:971-983.
45. Troyer, TW & Miller, KD (1997b) Integrate-and-fire neurons matched to physiological F-I curves yield high input sensitivity and wide dynamic range. *Computational Neuroscience, Trends in Research*. JM Boser, ed. New York:Plenum, pp. 197-201.
46. Bugmann, G, Christodoulou, C & Taylor, JG (1997) Role of temporal integration and fluctuation detection in the highly irregular firing of a leaky integrator neuron model

- with partial reset. *Neural Comput.* **9**:985-1000.
47. Amit, DJ & Brunel N (1997) Global spontaneous activity and local structured (learned) delay activity in cortex. *Cerebral Cortex* **7**:237-252.
48. van Vreeswijk, C & Sompolinsky, H (1998) Chaotic balanced state in a model of cortical circuits. *Neural Comput.* **10**:1321-1327.
49. Troyer, TW, Krukowski AE, Priebe NJ & Miller, KD (1998) Contrast-invariant orientation tuning in visual cortex: Thalamocortical input tuning and correlation-based intracortical connectivity. *Journal of Neuroscience* **18**:5908-5927.
50. Abbott, LF & Song, S (1999) Temporally Asymmetric Hebbian Learning, Spike Timing and Neuronal Response Variability. Kearns, MS, Solla, SA & Cohn, DA (editors) *Advances in Neural Information Processing Systems 11* (MIT Press, Cambridge MA) pp. 69-75.
51. Bekkers JM & Stevens CFJ (1996) Cable properties of cultured hippocampal neurons determined from sucrose-evoked miniature EPSCs. *Neurophysiol.* **75**:1250-1255.
52. Softky, WR (1995) Simple codes versus efficient codes. *Curr. Opin. Neurobiol.* **5**:239-247.
53. Sejnowski, TJ (1977) Storing covariance with nonlinearly interacting neurons. *Journal of Mathematical Biology* **4**:303-321.

54. Mehta, MR, Barnes, CA & McNaughton, BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Science USA* **94**:8918-8921.
55. Markram H, Tsodyks MV (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurones. *Nature* **382**:807-809.
56. Stryker, MP & Strickland, SL (1984) Physiological segregation of ocular dominance columns depends on the pattern of afferent electrical activity. *Inv. Ophthalm. Supp.* **25**:278.
57. Stryker, MP (1986) The role of neural activity in rearranging connections in the central visual system. In RJ Ruben, TR Van De Water & EW Rubel, editors, *The Biology of Change in Otolaryngology*. Amsterdam:Elsevier, 211-224.
58. Scanziani, M, Malenka, RC & Nicoll, RA (1996) Role of intercellular interactions in heterosynaptic long-term depression. *Nature* **380**:446-450.
59. Tang, Y-P, Shimizu, E, Dube, GR, Rampon, C, Kerchner, GA, Zhuo, M, & Tsien, JZ (1999) Genetic enhancement of learning and memory in mice. *Nature* **401**:63-69.
60. Yuste, R, Majewska, A, Cash, SS & Denk, W (1999) Mechanisms of Calcium Influx into Hippocampal Spines: Heterogeneity among Spines, Coincidence Detection by NMDA Receptors, and Optical Quantal Analysis. *J. Neurosci.* **19**:1976-1987.

Acknowledgments

Research supported by the Sloan Center for Theoretical Neurobiology at Brandeis University, the National Science Foundation (IBN-9817194), the National Institute of Mental Health (MH58754) and the W.M. Keck Foundation (LA); a Howard Hughes Predoctoral Fellowship (SS); and by R01-EY11001 from the National Eye Institute and an Alfred P. Sloan Research Fellowship (KM). We thank Todd Troyer for very useful discussions.