Competitive Hebbian Learning Through Spike-Timing Dependent Synaptic Plasticity

Sen Song[1], Kenneth D. Miller[2] and L. F. Abbott[1*]

[1]Volen Center for Complex Systems
and
Department of Biology
Brandeis University
Waltham MA 02254-9110
email: flamingo@brandeis.edu, abbott@brandeis.edu

[2]Departments of Physiology and Otolaryngology,
Neuroscience Graduate Program
W.M. Keck Center for Integrative Neuroscience,
Sloan Center for Theoretical Biology
University of California San Francisco CA 94143-0444
email: ken@phy.ucsf.edu

[*]Corresponding author

This is a preprint (final draft) of an article that appeared as
Nature Neuroscience, Vol. 3, pp. 919-926 (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). In modeling studies, we find that this form of synaptic modification can automatically balance synaptic strengths to make postsynaptic firing irregular but more sensitive to presynaptic spike timing. It has been argued that neurons in vivo operate in such a balanced regime. Synapses subject to STDP compete for control of the timing of postsynaptic action potentials. Inputs that fire the postsynaptic neuron with short latency or that act in correlated groups are able to compete most successfully and develop strong synapses, while the synapses of longer latency or less effective inputs are weakened.

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 Hebb[1]. 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 competition[4]. While Hebbian synaptic modification has received support from experiments on long-term potentiation and depression[5], 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 learning[6], 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. Competition can also arise locally if the processes that modify synaptic strengths equilibrate at a preset level of total synaptic innervation or postsynaptic activity. This can be achieved through static mechanisms such as thresholds and negative input correlations[6], dynamic mechanisms involving non-Hebbian synaptic growth or decay terms such as synaptic scaling[7–9], or shifts in the synaptic modification rule itself as in the sliding threshold of the BCM model[10]. Here we explore an entirely different mechanism suggested by experimental results on the effect of spike timing on long-term synaptic modification[11–19] (D.E. Feldman, Soc. Neurosci. Abst. 25, 223, 1999) 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[11–19] 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 recent experiments on neocortical slices[14], hippocampal slice[17] and cell[18] culture, and in vivo studies of tadpole tectum[19], 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 modification spike-timing dependent plasticity (STDP).

Synaptic modification by STDP-like rules has been studied previously in models of temporal pattern recognition[20,21], temporal sequence learning[22–24], coincidence detection[25,26], navigation[27–29], and direction selectivity[30,31] (N.J. Buchs, J. Reutimann & W. Senn, Soc. Neurosci. Abst. 25, 2259, 1999). We focus instead on the competitive and stabilizing properties of STDP. The competitive nature of STDP has been noted[25,19], but not studied in detail previously. Stability of an STDP-like rule in combination with non-Hebbian plasticity has been studied in a linear, stochastically-spiking neuron model[32], but we find qualitatively new behavior when the intrinsic nonlinearity of the spike-generation mechanism is taken into account. We find that STDP alone can lead to stable distributions of synaptic conductances, subject only to a limit on the strengths of individual synapses. The synaptic conductance distributions produced by STDP force the postsynaptic neuron into a balanced, irregularly firing regime[33–42] in which it is sensitive to the timing of the presynaptic action potentials it receives. Such sensitivity leads to competition among inputs for the control of postsynaptic spike timing. This allows STDP to selectively strengthen synapses of inputs with relatively shorter latencies or stronger mutual correlations, while weakening the remaining synapses.

Spike-Timing Dependent Synaptic Plasticity

The modeling studies we present are based on a spike-timing dependent synaptic plasticity rule in which a function $F(\Delta t)$ determines the amount of synaptic modification arising from a single pair of pre- and postsynaptic spikes separated by a time Δt. The function (Fig. 1)

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

provides a reasonable approximation of the dependence of synaptic modification on spike timing seen in the experimental data. The parameters τ_+ and τ_- determine the ranges of pre-to-postsynaptic interspike intervals over which synaptic strengthening

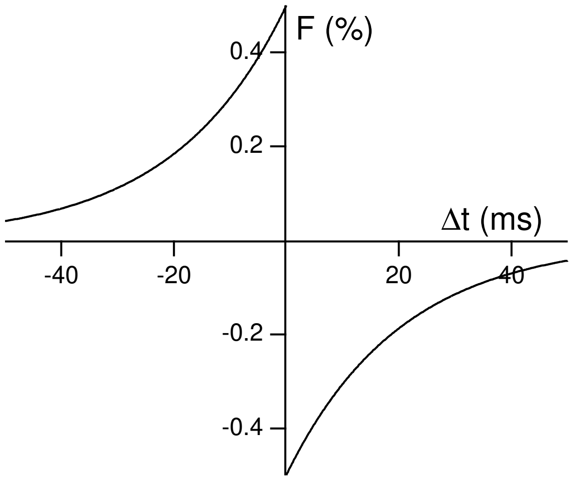

Figure 1: The STDP modification function. The change of the peak conductance at a synapse due to a single pre- and postsynaptic action potential pair is $F(\Delta t)$ times the maximum value $\overline{g}_{\mathrm{max}}$, with Δt the time of the presynaptic spike minus the time of the postsynaptic spike.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

and weakening occur. A_+ and A_-, which are both positive, determine the maximum amounts of synaptic modification, which occur when Δt is close to zero.

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[14,18,19], while other results reveal a larger window for synaptic weakening[17] (D.E. Feldman, Soc. Neurosci. Abst. 25, 223, 1999). 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.

We determine the parameters A_+ and A_- 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 (see Discussion). In our simulations, $A_+ = 0.005$, except in Fig. 2f where $A_+ = 0.02$. To set the value of A_-, we make the important assumption that synaptic weakening through STDP is, overall, a slightly larger effect than synaptic strengthening[26]. 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 synapses. A negative integral of F requires $A_-\tau_- > A_+\tau_+$. The data are mixed on this issue. The results that report similar time scales for synaptic strengthening and weakening[14,18,19] 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] (D.E. Feldman, Soc. Neurosci. Abst. 25, 223, 1999) support the dominance of synaptic weakening over strengthening by STDP. In our simulations we use $A_-/A_+ = 1.05$, except for Fig. 2d where A_-/A_+ varies.

In the model we study, $\overline{g}_a$ denotes the peak synaptic conductance (the synaptic conductance immediately after an isolated presynaptic spike) due to an excitatory synapse labeled by the integer a (with $a = 1, 2, \ldots, N$). This conductance must always be positive, and is not allowed to exceed a maximum value $\overline{g}_{\max}$. A pre- and postsynaptic spike pair separated by a time interval Δt modifies the peak synaptic conductance by an amount $\overline{g}_{\max}F(\Delta t)$. 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 $0 \leq \overline{g}_a \leq \overline{g}_{\max}$, $\overline{g}_a$ is set to the appropriate limiting value. A scheme for implementing this modification rule is presented in the Methods section.

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 with 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. Response variability also provides a constraint on the synaptic inputs to a neuron. The responses of an integrate-and-fire model receiving many independent presynaptic inputs can be considerably less variable than responses seen in vivo[33]. Correlations of input spike timing, such as synchronization, can contribute to increased variability[34]. However, a number of authors have noted that a high degree of variability can also arise if the excitatory inputs to a neuron are balanced relative to the inhibitory synaptic and membrane currents[35–41]. The critical condition is that the mean input to the neuron should only be sufficient to raise the membrane potential 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 the total level of input. As we will see, STDP provides a mechanism by which this balance can be established and maintained over a wide range of input firing rates[42]. This results in a state in which presynaptic action potentials can control the timing of postsynaptic spikes and competition among synapses can arise.

To study the equilibrium distribution of synaptic strengths arising from STDP, we initially set the peak conductances of all the excitatory synapses of the model neuron to $\overline{g}_{\max}$, 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. Although all the peak synaptic conductances started with the same value, there is no stable equilibrium state with a uniform distribution of values. Instead, most of the peak synaptic conductances are pushed toward the limiting values of zero or $\overline{g}_{\max}$ (Figs. 2a and 2b). For low input rates, more synapses approach the upper limit (Fig. 2a), and for high input rates more are pushed toward zero (Fig. 2b). This has the effect of keeping the total synaptic input to the neuron roughly constant, independent of the presynaptic firing rates. The split between strong and weak synapses is also affected by the values of $\overline{g}_{\max}$ (fewer strong synapses develop for larger $\overline{g}_{\max}$) 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.

STDP has a strong regulatory effect on the steady-state firing rate of the postsynaptic neuron, which, for the equilibrium distribution of synaptic strengths, increases by only about 1 Hz for each 5 Hz increase in the input firing rate (Fig. 2c). 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 input firing rates. Synaptic changes due to STDP take time to develop, so STDP only regulates the long-term average firing rate and the neuron remains highly sensitive to transient changes of the input firing rates.

The coefficient of variation (CV) of the postsynaptic spike train, which is the standard deviation of the interspike intervals divided by their mean, is fairly large

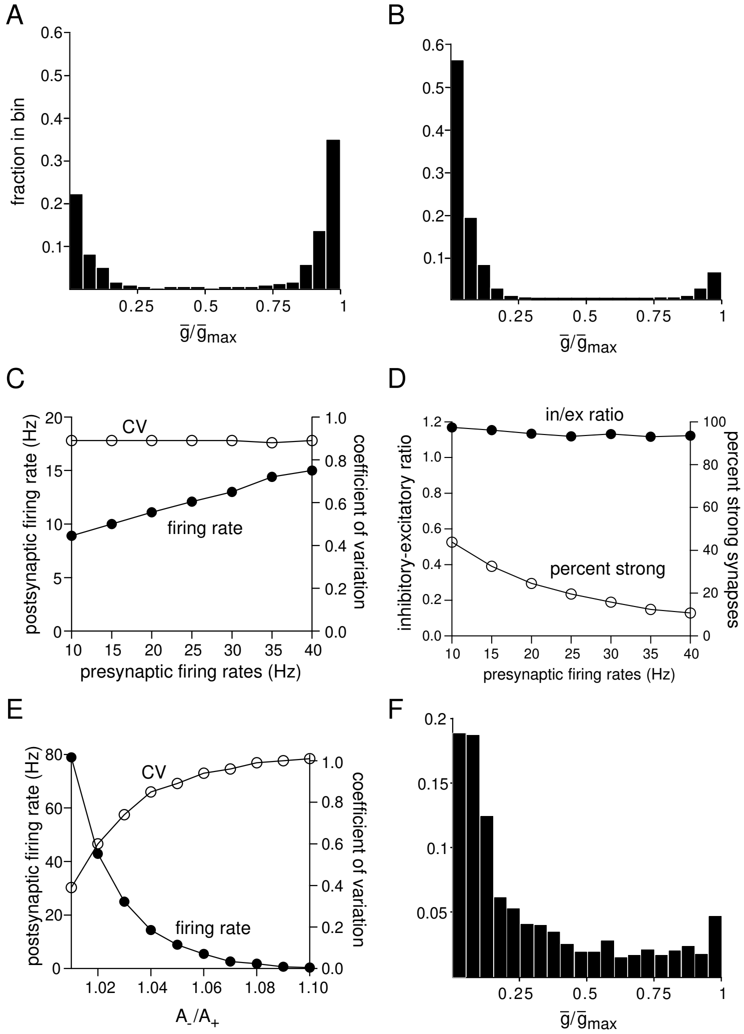

Figure 2: Balanced excitation and irregular firing produced by STDP. A) Histogram of the fraction of synapses taking different peak conductance values ranging from zero to $\overline{g}_{\max}$. For an input rate of 10 Hz, the peak synaptic conductances tend to the limiting values, but more are near $\overline{g}_{\max}$ than near zero. B) Same as A, but for an input rate of 40 Hz. Now more peak conductances are near zero than near $\overline{g}_{\max}$. C) The postsynaptic firing rate and CV (stdev/mean) of the postsynaptic interspike intervals for different input firing rates. D) The ratio of total inhibitory to excitatory currents at threshold and the percentage of strong synapses ($\overline{g} \geq 0.8\overline{g}_{\max}$) for different presynaptic firing rates. The inhibitory or excitatory current at threshold is defined as the current that would flow if all conductances of the given type were simultaneously activated while membrane potential was clamped at threshold. The leak conductance is included as an inhibitory conductance 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 $\overline{g}_{\max}$ 2.33 times larger and the synaptic modification per spike pair four times larger ($\overline{g}_{\max} = 0.035$, $A_+ = 0.020$, $A_=0.021$). The larger value of $\overline{g}_{\max}$ forces more synapses to lower conductance values, while the higher modification rate filled in the distribution.

and remarkably independent of the input firing rate (Fig. 2c) when the distribution of synaptic strengths due to STDP has equilibrated. 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 currents when the membrane potential is at the action-potential threshold. STDP adjusts this ratio to be slightly greater than one over the entire range of presynaptic firing rates considered (Fig. 2d). 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. Fig. 2d shows the number of strong synapses (those with $\overline{g} \geq 0.8\overline{g}_{\mathrm{max}}$) for different presynaptic firing rates. For the value of $\overline{g}_{\mathrm{max}}$ 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_+ (Fig. 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.

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 (Figs. 2a and 2b). This results in a bimodal distribution. A more continuous distribution arises if the degree of synaptic modification per spike pair is increased. Fig. 2f shows an example in which the equilibrium distribution of synaptic conductances is roughly exponential, except for a small excess near $\overline{g} = \overline{g}_{\mathrm{max}}$. The basic features of STDP, regulation of the postsynaptic firing rate and CV, remain in this case, but the synaptic conductance distribution more closely matches the experimentally observed distribution of spontaneous synaptic (mini) potentials[43], which provides one estimate of the distribution of synaptic strengths.

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[38,39,42]. When excitation is strong, as at the beginning of our simula-

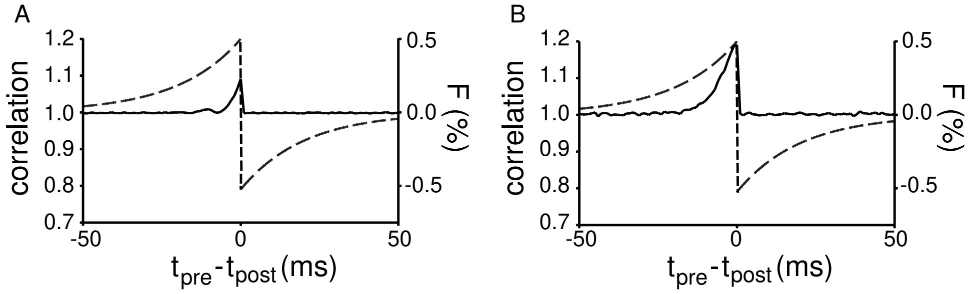

Figure 3: Correlation between pre- and postsynaptic action potentials before and after STDP. The solid curves indicate the relative probability of any presynaptic spike occurring at time t_{pre} when a postsynaptic spike occurs at time t_{post}. A correlation of one is the value due solely to chance 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, synapses are strengthened or weakened. A) At the beginning of our simulations, when all the 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 $A_-/A_+ > 1$.

tions, so that the mean input to the neuron would bring it well above threshold if action potentials were blocked, the neuron operates in an input-averaging or regular-firing mode. The postsynaptic action potential sequences produced in this mode are significantly more regular than the presynaptic spike trains that evoke them. The interspike intervals of the postsynaptic response depend on the total synaptic input, but the absolute timing of individual postsynaptic action potentials is fairly insensitive to presynaptic spike times. As a result, there are roughly equal numbers of presynaptic action potentials before and after each postsynaptic spike[39,42] (Fig. 3a). As we have noted, the area under the synaptic weakening portion of the STDP curve (Fig. 1) 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 (Fig. 3a).

As the excitatory synapses are weakened by STDP, the postsynaptic neuron enters a balanced mode of operation in which it generates a more irregular sequence of action potentials that are more tightly correlated with the presynaptic spikes that evoke them. The total synaptic input in the balanced mode is, on average, near or sub-threshold, so the postsynaptic neuron fires irregularly, primarily in response to statistical fluctuations in the total input. Because action potentials occur preferentially after a random fluctuation, there tend to be more excitatory presynaptic spikes before than after a postsynaptic response[38,39,42] (Fig. 3b). 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 curve[42] (Fig. 3b). If the total excitatory drive were weaker than that provided by this distribution, stronger fluctuations of the total input would be required to cause postsynaptic spikes. This would create an even greater excess of presynaptic action potentials prior to postsynaptic firing, which would lead to an increase in synaptic strengths, driving the system back to the steady-state distribution. STDP thus modifies excitatory synaptic strengths until there is a sufficiently, but not excessively, high probability of a presynaptic action potential occurring prior to a postsynaptic spike . This causes the neuronal response to be sensitive to the timing of 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. STDP strengthens synapses that fire prior to a postsynaptic spike and weakens those that fire later. Suppose, for example, that a stimulus causes a barrage of presynaptic inputs to fire with varying latencies, and that these latencies extend over a period longer than that required to evoke postsynaptic spiking. In this case, STDP will strengthen shorter-latency excitatory inputs while weakening those with longer latencies. The ultimate effect of this synaptic modification is to make the postsynaptic neuron respond more quickly.

To illustrate this latency reduction, we considered a model neuron receiving inputs that are silent except for isolated events represented by bursts of spikes with a Poisson distribution at 100 Hz for 20 ms. Different synapses are 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

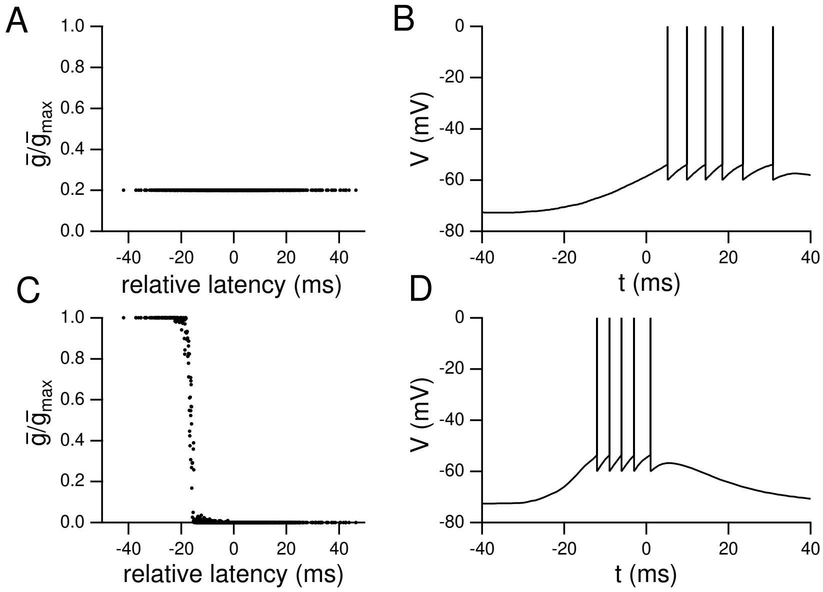

Figure 4: Reduction of latency by STDP. A) The initial peak synaptic conductances plotted as a function of the relative latency of their synaptic inputs. B) The initial postsynaptic response to a barrage of excitatory input with burst onset for each synapse occurring at the time of its relative latency. 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.

standard deviation of 15 ms. The burst of action potentials at a given synapse occurs at a time given by the sum of its relative latency and the absolute latency associated with the event.

Initially, all the synapses were set to the same strength of $0.2\overline{g}_{\rm max}$ (Fig. 4a). This produced a response in the postsynaptic cell that began shortly after the time marked zero, which indicates the mean input latency, and lasted for about 25 ms (Fig. 4b). The input events were then repeated periodically until the STDP rule had established a fixed distribution of peak synaptic conductances. STDP strengthened short-latency inputs to the maximum allowed level, $\overline{g}_{\rm max}$, and weakened synapses with longer latencies to zero (Fig. 4c). This produced a quicker response in the postsynaptic neuron, which fires almost 20 ms earlier than it did originally (Fig. 4d).

Correlation-Based Hebbian Modification

Factors that enhance the ability of a given synapse to rapidly evoke a postsynaptic response lead to its strengthening through STDP. Correlating different synaptic inputs so 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 the effect of STDP.

Presynaptic firing rates were generated to have a correlation function that decayed exponentially with a time constant τ_c and varied in amplitude across the population of synapses (see Methods). Specifically, the correlation between two cells a and b is $c_a c_b$ with c_a and c_b, which we call correlation parameters, varying from zero to 0.2 uniformly across the 1000 excitatory synapses. When the correlations decay rapidly ($\tau_c = 20$ ms, Fig. 5a), more correlated synapses become stronger and less correlated synapses weaken (compare Figs. 5a and 5b). This trend disappears for larger correlation times ($\tau_c = 200$ ms, Fig. 5b). To be strengthened, a group of inputs must fire together long enough to generate a postsynaptic action potential, but must then stop firing so 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 STDP[32].

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. When input firing rates are time-independent and uncorrelated but vary uniformly across the population of synapses, there is little tendency for synapses firing at either faster or slower rates to be preferentially strengthened or weakened by STDP (Fig. 5c). 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.

The steady-state peak synaptic conductances are also insensitive to the degree of variability of the presynaptic input. When we arranged the input firing rates to have a standard deviation that varied uniformly across the inputs, we found no tendency for synapses with either more or less variable firing rates to be preferentially strengthened or weakened by STDP (Fig. 5d).

The basic result of these studies is that STDP is insensitive to the average rate

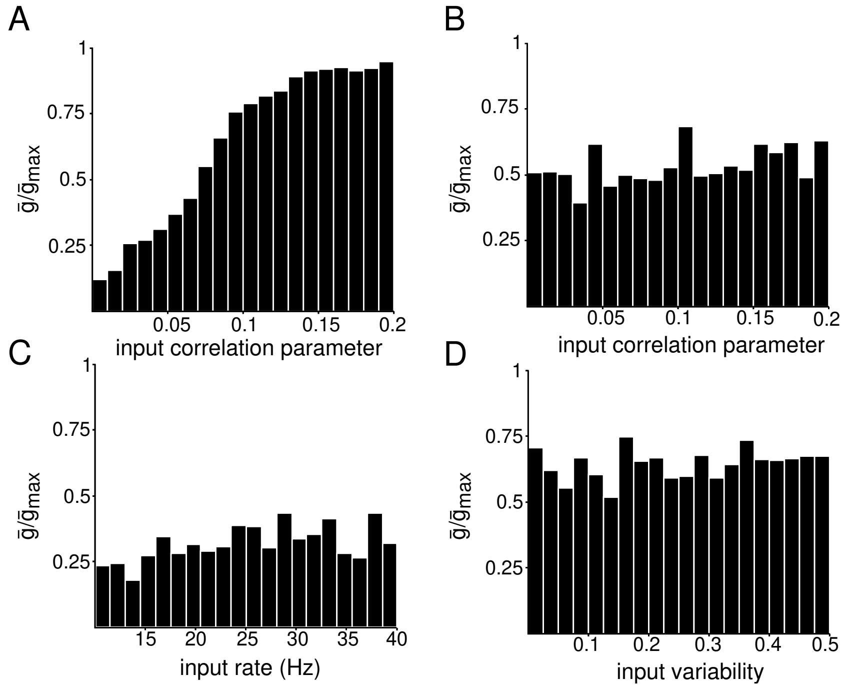

Figure 5: Effects of the input parameters correlation, firing rate, or variability on steady-state peak synaptic conductances. Each input parameter is divided into 20 bins. The histograms show the average peak synaptic conductances for all inputs within a given bin. These values are the results of averaging bimodal distributions of synaptic strengths within each bin. (A) The synaptic inputs have correlation parameters c_a ranging from zero to 0.2 ($c_a = 0.2(a-1)/(N-1)$) The degree of correlation between any two inputs is determined by the product of their correlation parameters. The correlation time constant is 20 ms. 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 r_a ranging from 10 to 40 Hz ($r_a = 10 + 30(a-1)/(N-1)$), 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 ($\sigma_a = 0.5(a-1)/(N-1)$). No effect of variability on synaptic strength is observed. In this example, $\tau_c = 20$ ms, as in (a).

or degree of variability of a given synaptic input. It is, however, strongly affected by correlations between different inputs, provided that they decay rapidly enough as a function of time. Synapses with strong, rapidly decaying temporal correlations are strengthened as a cluster and 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 rule[44]. However, such a rule cannot, in general, achieve competition between synapses[6]. This brings up 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 competition for control of the timing of postsynaptic action potentials. Inputs that consistently the best at predicting 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 the nonlinearity of the spike-generation process. In a model in which the probability of spiking depends linearly on membrane voltage[32], the correlation between pre- and postsynaptic firing does not change shape with overall changes in synaptic efficacy, as it does in Fig. 3. Nonlinear effects, which make causal input-output correlations grow relative to acausal correlations as overall synaptic efficacy decreases, are crucial for producing the stabilizing and competitive effects of STDP that we have discussed.

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 can differentially strengthen the shortest-latency inputs evoked by a stimulus. There is some experimental evidence suggesting that the resulting reduction of latency in the postsynaptic response occurs in vivo. A phenomenon analogous to the reduction of latency discussed here predicts that, when a rat moves through a particular region of space, place cells active for that region should fire earlier after the rat has repeatedly traversed the area[23,27]. This effect has been observed experimentally[45,29].

A key assumption in our model is that synaptic weakening by STDP dominates over synaptic strengthening. This is critical for stability. If this assumption 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[5]). 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.

Our model of STDP involves two additional assumptions. We assumed that the effects of 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 saturation[14]. Our model does not incorporate this finding, but we maintain presynaptic rates above the reported threshold frequency for synaptic strengthening[14]. We also assumed that we could ignore delays of several minutes between pairing of pre- and postsynaptic spikes and the resultant induction of synaptic modification that are suggested by experiments[14]. 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 Fig. 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.

STDP may modify the short-term synaptic plasticity properties of a synapse as well as its efficacy, an effect which has been called synaptic redistribution[46]. We have run simulations in which we coupled 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 STDP and short-term plasticity. The most effective way to strengthen a synapse under STDP is to have it release transmitter before a postsynaptic spike and then stop releasing so that it does not get weakened by subsequent releases occurring after postsynaptic activity. A high degree of synaptic depression, which is a feature of strong synapses in the redistribution scheme[46], assures this. STDP that acts to modify release probability and change the degree of synaptic depression is thus an extremely effective and competitive mechanism for 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[7–9], may serve this function instead. In the model of STDP we use, two sets of inputs that fire at times separated by more that about 100 ms generate STDP independently and thus do not compete. Experiments suggest that competition can nevertheless occur under these conditions[47]. Such a result could arise if the STDP temporal window for synaptic weakening has a long enough tail, or if STDP is supplemented by sufficiently strong heterosynaptic long-term depression[48] or competition induced by synaptic scaling[7–9].

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[14–19], 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 currents[49]. This might be associated with a reduction in width of the STDP window[50]. STDP also depends on postsynaptic back-propagating action potentials[14,15], and modification of their waveforms might also change the timing requirements for synaptic plasticity. Finally, the results we report are sensitive to the ratio of the areas under the strengthening and weakening parts of the STDP curve (Fig. 1) and would be more robust if this ratio were under the dynamic control of the average postsynaptic firing rate. It will be interesting to see if evidence of developmental or activity-dependent meta-plasticity in either the amplitudes or decay times of STDP modification curves is revealed in future experiments.

Methods

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

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

with $\tau_m = 20\,\text{ms}$, $V_{\text{rest}} = -70\,\text{mV}$, $E_{\text{ex}} = 0\,\text{mV}$, and $E_{\text{in}} = -70\,\text{mV}$. In addition, when the membrane potential 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 37). The synaptic conductances g_{ex} and g_{in} and their related peak conductances (see below) are measured in units of the leakage conductance of the neuron and are thus dimensionless.

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

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

We have taken $\tau_{\text{ex}} = \tau_{\text{in}} = 5$ ms, $\overline{g}_{\text{in}} = 0.05$, and $0 \leq \overline{g}_a \leq \overline{g}_{\max}$ with $\overline{g}_{\max} = 0.015$ (except for Figure 4, where $\overline{g}_{\max} = 0.02$ and Figure 2D, where $\overline{g}_{\max} = 0.035$). For a 100 M$\Omega$ input resistance, $\overline{g}_{\max} = 0.015$ corresponds to a peak synaptic conductance of 150 pS.

Synaptic modification is generated in the model through $N+1$ functions, $M(t)$ and $P_a(t)$, for $a = 1, 2, \ldots, N$. These decay exponentially,

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

Every time the postsynaptic neuron fires an action potential, $M(t)$ is decremented by an amount A_- , and every time synapse a receives an action potential, $P_a(t)$

is incremented by an amount A_+. $M(t)$ is used to decrease synaptic strength. If synapse a receives a presynaptic action potential at time t, its maximal conductance parameter is modified according to $\overline{g} \to \overline{g} + M(t)\overline{g}_{\max}$. If this makes $\overline{g}_a < 0$, $\overline{g}_a$ is set to zero. $P_a(t)$ is used to increase the strength of synapse a. If the postsynaptic neuron fires an action potential at time t, $\overline{g}_a$ is modified according to $\overline{g}_a \to \overline{g}_a + P_a(t)\overline{g}_{\max}$. If this makes $\overline{g}_a > \overline{g}_{\max}$, $\overline{g}_a$ is set to $\overline{g}_{\max}$. These modifications are made after the decay in conductance described in the previous paragraph, but changing this order does not modify our results.

The presynaptic firing rates in Figs. 5a and 5b, have the correlation function $\langle r_a(t)r_b(t')\rangle = \overline{r}^2 + \overline{r}^2(\sigma^2\delta_{ab} + (1-\delta_{ab})c_a c_b)\exp(-|t-t'|/\tau_c)$, where the angle brackets represent an average over the ensemble of rates, $\overline{r} = 10$ Hz, and $\sigma = 0.5$. To generate such rates, we chose 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 $a = 1, 2, \ldots, N$, from Gaussian distributions with zero mean and standard deviation one and σ_a respectively, where $\sigma_a^2 = \sigma^2 - c_a^2$. At the start of each interval, the firing rate for synapse a was set to $r_a = \overline{r}(1 + x_a + c_a y)$, and it was held at this value until the start of the next interval. The correlation function for Fig. 5d is $\langle r_a(t)r_b(t')\rangle = \overline{r}^2 + \overline{r}^2\sigma_a^2\delta_{ab}\exp(-|t-t'|/\tau_c)$, and the rates were computed using a similar procedure but with $c_a = 0$ and a variable σ_a.

To ensure that our results do not depend on initial conditions, we ran multiple trials of the simulations starting from different randomly-generated sets of initial synaptic weights. There were no detectable changes. 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 their overall distribution does not significantly change. We consider the synaptic distributions to have converged when the output firing rate stops 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, D.O. The Organization of Behavior: A Neuropsychological Theory, Wiley:New York (1949).

2. Guillery, R.W. Binocular competition in the control of geniculate cell growth. J. Comp. Neurol. 144, 117-130 (1972).

3. Miller, K.D. Synaptic Economics: Competition and cooperation in synaptic plasticity. Neuron 17, 371-374 (1996).

4. Miller, K.D. 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. Springer-Verlag:New York, pp. 55-78 (1996).

5. Bear, M.F. & Malenka, R.C. Synaptic plasticity: LTP and LTD. Curr. Opin. Neurobiol. 4, 389-399 (1994).

6. Miller, K.D. & MacKay, D.J.C. The role of constraints in Hebbian learning. Neural Comp. 6, 100-126 (1994).

7. Turrigiano, G.G., Leslie K.R., Desai, N.S., Rutherford L.C. & Nelson, S.B. Activity-dependent scaling of quantal amplitude in neocortical neurons. Nature 391, 892-896 (1998).

8. Davis, G.W. & Goodman, C.S. Synapse-specific control of synaptic efficacy at the erminals of a single neuron. Nature 392, 82-86 (1998).

9. O'Brien, R.J., Kamboj, S., Ehlers, M.D., Rosen, K.R., Fischbach, G.D. & Huganir, R.L. Activity-dependent modulation of synaptic AMPA receptor accumulations. Neuron 21, 1067-1078 (1998).

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

11. Levy, W.B. & Steward, D. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. Neurosci. 8, 791-797 (1983).

12. Gustafsson, B., Wigstrom, H., Abraham, W.C. & Huang, Y.-Y. Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. J. Neurosci. 7, 774-780 (1987).

13. Debanne, D., Gahwiler, B.H. & Thompson, S.M. Asynchronous pre- and post-synaptic activity induces associative long-term depression in area CA1 of the rat hippocampus in vitro Proc. Natl. Acad. Sci. USA 91, 1148-1152 (1994).

14. Markram H., Lubke J., Frotscher M. & Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275, 213-215 (1997).

15. Magee J.C.& Johnston D. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. Science 275, 209-213 (1997).

16. Bell C.C., Han V.Z., Sugawara Y. & Grant K. Synaptic plasticity in a cerebellum-like structure depends on temporal order. Nature 387, 278-281 (1997).

17. Debanne D., Gahwiler, B.H. & Thompson, S.M. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. J. Physiol. 507, 237-247 (1998).

18. Bi G.-q. & Poo M.-m. Activity-induced synaptic modifications in hippocampal culture: dependence on spike timing, synaptic strength and cell type. J. Neurosci. 18, 10464-10472 (1998).

19. Zhang L.I., Tao, H.W., Holt C.E., Harris W.A. & Poo M.-m. A critical window for cooperation and competition among developing retinotectal synapses. Nature 395, 37-44 (1998).

20. Herz, A.V.M., Sulzer, B., Khn, R. & van Hemmen, J.L. Hebbian learning reconsidered: Representation of static and dynamic objects in associative neural nets. Biol. Cybern. 60, 457-467 (1989).

21. Gerstner, W., Ritz, R. & van Hemmen, J.L. Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. Biol. Cybern. 69, 503-515 (1993).

22. Minai, A.A. & Levy, W.B. Sequence learning in a single trial. INNS World Congress of Neural Networks II, pp. 505-508 (1993).

23. Abbott, L.F. & Blum, K.I. Functional significance of long-term potentiation for sequence learning and prediction. Cerebral Cortex 6, 406-416 (1996).

24. Roberts, P.D. Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. J. Computational Neurosci. 7, 235-246 (1999).

25. Gerstner, W., Kempter, R., van Hemmen, J.L. & Wagner, H. A neuronal learning rule for sub-millisecond temporal coding. Nature 383, 76-78 (1996).

26. Gerstner, W., Kempter, R., van Hemmen, J.L. & Wagner, H. 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 (1997).

27. Blum, K.I . & Abbott, L.F. A model of spatial map formation in the hippocampus of the rat. Neural Comp. 8, 85-93 (1996).

28. Gerstner, W. & Abbott, L.F. Learning navigational maps through potentiation

and modulation of hippocampal place cells. J. Computational Neurosci. 4, 79-94 (1997).

29. Mehta, M.R., Quirk, M.C. & Wilson, M. Experience dependent asymmetric shape of hippocampal receptive fields. Neuron 25, 707-715 (2000).

30. Rao, R. & Sejnowski, T.J. Predictive sequence learning in recurrent neocortical circuits. In Solla, S.A., Leen, T.K. & Muller K.-b. (editors) Advances in Neural Information Processing Systems 12 MIT Press:Cambridge MA (2000).

31. Mehta, M.R. & Wilson, M. From hippocampus to V1: Effect of LTP on spatio-temporal dynamics of receptive fields. In Bower, J. (editor) Computational Neuroscience, Trends in Research 1999. Elsevier:Amsterdam (2000).

32. Kempter R., Gerstner W. & van Hemmen J.L. Hebbian learning and spiking neurons. Phys. Rev. E59, 4498-4514 (1999).

33. Softky, W.R. & Koch, C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. J. Neurosci. 13, 334-350 (1994).

34. Stevens, C.F. & Zador, A.M. Input synchrony and the irregular firing of cortical neurons. Nature Neurosci. 1, 210-217 (1998).

35. Shadlen, M.N. & Newsome, W.T. Noise, neural codes and cortical organization. Curr. Opin Neurobiol. 4, 569-579 (1994).

36. Tsodyks, M. & Sejnowski, T.J. Rapid switching in balanced cortical network models. Network 6, 1-14 (1995).

37. Troyer, T.W. & Miller, K.D. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. Neural Comp. 9:971-983 (1997).

38. Troyer, T.W. & Miller, K.D. Integrate-and-fire neurons matched to physiological F-I curves yield high input sensitivity and wide dynamic range. In Bower, J. (editor) Computational Neuroscience, Trends in Research, Plenum:New York, pp. 197-201 (1997).

39. Bugmann, G., Christodoulou, C. & Taylor, J.G. Role of temporal integration and fluctuation detection in the highly irregular firing of a leaky integrator neuron model with partial reset. Neural Comp. 9, 985-1000 (1997).

40. Amit, D.J .& Brunel N. Global spontaneous activity and local structured (learned) delay activity in cortex. Cerebral Cortex 7, 237-252 (1997).

41. van Vreeswijk, C. & Sompolinsky, H. Chaotic balanced state in a model of cortical circuits. Neural Comp. 10, 1321-1327 (1998). ???

42. Abbott, L.F. & Song, S. Temporally Asymmetric Hebbian Learning, Spike Timing

and Neuronal Response Variability. In Kearns, M.S., Solla, S.A. & Cohn, D.A. (editors) Advances in Neural Information Processing Systems 11, MIT Press:Cambridge MA, pp. 69-75 (1999).

43. Bekkers J.M. & Stevens C.F.J. Cable properties of cultured hippocampal neurons determined from sucrose-evoked miniature EPSCs. Neurophysiol. 75, 1250-1255 (1996).

44. Sejnowski, T.J. Storing covariance with nonlinearly interacting neurons. J. Math. Biol. 4, 303-321 (1977).

45. Mehta, M.R., Barnes, C.A. & McNaughton, B.L. Experience-dependent, asymmetric expansion of hippocampal place fields. Proc. Natl. Acad. Sci. USA 94, 8918-8921 (1997).

46. Markram H., Tsodyks M.V. Redistribution of synaptic efficacy between neocortical pyramidal neurones. Nature 382, 807-809 (1996).

47. Stryker, M.P. The role of neural activity in rearranging connections in the central visual system. In Ruben, R.J., Van De Water, T.R. & Rubel, E.W. (editors) The Biology of Change in Otolaryngology. Elsevier:Amsterdam, pp. 211-224 (1986).

48. Scanziani, M., Malenka, R.C. & Nicoll, R.A. Role of intercellular interactions in heterosynaptic long-term depression. Nature 380, 446-450 (1996).

49. Tang, Y.-P., Shimizu, E., Dube, G.R., Rampon, C., Kerchner, G.A., Zhuo, M. & Tsien, J.Z. Genetic enhancement of learning and memory in mice. Nature 401, 63-69 (1999).

50. Yuste, R., Majewska, A., Cash, S.S . & Denk, W. Mechanisms of Calcium Influx into Hippocampal Spines: Heterogeneity among Spines, Coincidence Detection by NMDA Receptors, and Optical Quantal Analysis. J. Neurosci. 19, 1976-1987 (1999).

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.