

Extending the effects of spike-timing-dependent plasticity to behavioral timescales

Patrick J. Drew*[†] and L. F. Abbott[‡]

*Neurobiology Section, Division of Biology, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0357; and [‡]Center of Neurobiology and Behavior, Department of Physiology and Cellular Biophysics, Columbia University College of Physicians and Surgeons, Kolb Research Annex, 1051 Riverside Drive, New York, NY 10032-2695

Edited by Charles F. Stevens, The Salk Institute for Biological Studies, La Jolla, CA, and approved April 19, 2006 (received for review January 26, 2006)

Activity-dependent modification of synaptic strengths due to spike-timing-dependent plasticity (STDP) is sensitive to correlations between pre- and postsynaptic firing over timescales of tens of milliseconds. Temporal associations typically encountered in behavioral tasks involve times on the order of seconds. To relate the learning of such temporal associations to STDP, we must account for this large discrepancy in timescales. We show that the gap between synaptic and behavioral timescales can be bridged if the stimuli being associated generate sustained responses that vary appropriately in time. Synapses between neurons that fire this way can be modified by STDP in a manner that depends on the temporal ordering of events separated by several seconds even though the underlying plasticity has a much smaller temporal window.

computational model | conditioning | learning

Linking behavior to underlying physiological mechanisms is one of the central goals of neuroscience, but a major impediment is the large difference between behavioral and physiological timescales. A good example of this discrepancy is provided by an experiment with *Drosophila* in which Tanimoto *et al.* (1) found that an odor becomes an aversive stimulus if it is followed by a shock during training but becomes attractive if it is presented after the shock. This finding is tantalizingly reminiscent of spike-timing-dependent plasticity (STDP; reviewed in ref. 2) in which a synapse is strengthened if presynaptic spikes are paired with subsequent postsynaptic spikes but weakened if the order is reversed. Indeed, the plot of learning index versus time interval between odor and shock in Tanimoto *et al.* (ref. 1; reproduced in Fig. 1*B*) has a shape and form similar to curves showing the amount of STDP as a function of the time interval between paired pre- and postsynaptic spikes (Fig. 1*A*). The problem is that the timescale is vastly different, seconds in the case of the behavioral experiment and milliseconds in the case of STDP. As in many similar cases, to account for the behavioral data on the basis of synaptic physiology, we must find a mechanism to span this large gap in timescales. We propose that responses to stimuli that are sustained, but slowly decay, could provide such a bridging mechanism.

The ability of animals to detect temporal correlations and use them to make predictions is a basic result of classical conditioning. More recent experiments have found neural correlates of prediction. Dopaminergic neurons initially respond to rewards such as food or juice, but over repeated pairings, they lose these direct reward responses and instead begin responding to stimuli that predict reward or even to stimuli that predict stimuli that predict reward (3). In the hippocampus, the receptive fields of place cells shift and expand after repeated traversals of a path, making them respond to positions that precede their initial place field during the run (4, 5) and leading to a predictive map of space (6). To connect such predictive activity to synaptic plasticity, we consider the fate of a single neuron within the complex circuitry involved in actual conditioning behavior. The key issue is whether a form of synaptic plasticity that depends on spike

timing at the millisecond timescale can induce modifications in the response of this model neuron that depend on temporal sequences occurring over vastly longer timescales.

It is well known that correlations between pre- and postsynaptic firing can lead to synaptic potentiation or depression through spike-timing-dependent forms of long-term synaptic plasticity (6–11). Our purpose is to explore whether this mechanism and the firing patterns it requires are a plausible way of bridging the gap in time scales between synaptic plasticity and behavior.

Results

The amount of synaptic potentiation or depression due to STDP is typically measured by pairing a number of pre- and postsynaptic action potentials with a specified time interval between them (2, 12–17). This procedure results in an STDP “window function” such as that shown in Fig. 1*A*, which indicates the amount of synaptic modification caused by pairing pre- and postsynaptic spikes as a function of the time interval between them. If firing rates are low enough so that only a single spike pair typically appears within the STDP window at any given time, the effects of individual spike pairs appear to sum fairly linearly, although this breaks down at higher firing rates (15, 16). Within this linear regime, the total impact of a sequence of spike pairings can be computed by calculating the integral of the STDP window function multiplied by the cross-correlation function of the pre- and postsynaptic spike trains (6–11). This calculation is equivalent to determining the total amount of synaptic modification over a trial by multiplying the amount of synaptic change in Fig. 1*A* for a particular interspike interval by the probability of a pre-/postsynaptic spike pair occurring with that interval, and then integrating it over all possible intervals. Experimental work supports the validity of this method for calculating the total amount of potentiation or depression produced by STDP (14, 17).

For STDP to produce net potentiation over time, pre- and postsynaptic spike sequences must be correlated in such a way that it is more likely for a pre-then-post temporal ordering to occur than a post-then-pre sequence. Such correlations require an appropriate relationship between time-dependent pre- and postsynaptic firing rates. An obvious requirement, shared by any Hebbian mechanism of plasticity, is that pre- and postsynaptic firing must overlap in time, at least to within the interval covered by the STDP window function. When two stimuli are separated by seconds, as is the case between the conditioned stimulus (CS) and the unconditioned stimulus (US) in the classical conditioning paradigm, we consider that this is not an easy requirement to satisfy. Our approach to this problem is to assume that both

Conflict of interest statement: No conflicts declared.

This paper was submitted directly (Track II) to the PNAS office.

Freely available online through the PNAS open access option.

Abbreviations: STDP, spike-timing-dependent plasticity; CS, conditioned stimulus; US, unconditioned stimulus; LTP, long-term potentiation; LTD, long-term depression.

[†]To whom correspondence should be addressed. E-mail: pjdrew@ucsd.edu.

© 2006 by The National Academy of Sciences of the USA




Fig. 1. STDP and conditioned responses. (A) The fractional strengthening or weakening of a synapse per spike pair as a function of the time difference between the pre- and postsynaptic action potentials, also known as the STDP window function. (B) The learning index, which indicates a fly's attraction to or avoidance of an odor that previously had been paired with a shock as a function of the time interval between the presentation of the odor and the shock during training. [Adapted with permission from ref. 1 (Copyright 2004, Macmillan Publishers, Ltd.).] (C) The amount of synaptic modification as a function of the interstimulus interval produced by each single trial, from the analytic formula given in the text. Parameter values are $A_+ = A_- = 0.5\%$, $\tau_+ = \tau_- = 20$ ms, $R_{pre} = R_{post} = 50$ Hz, $\tau_{pre} = 30$ s, and $\tau_{post} = 80$ s.

the US and CS generate persistent responses that decay over time but nevertheless outlast these stimuli long enough to span the interval between them. Modeling the mechanism that sustains these responses and makes them decay slowly is beyond the scope of our work, but we discuss possible mechanisms in *Discussion*.

We begin our analysis by looking at cross-correlations between spike sequences generated by neurons excited by stimuli that then exhibit sustained but slowly decaying activity. Fig. 2 *Inset* shows the pattern of pre- and postsynaptic firing we consider. We

assume that the postsynaptic neuron and the presynaptic afferents are driven by separate stimuli. The stimulus that excites the presynaptic afferents corresponds to the CS in a classical conditioning paradigm, whereas the stimulus that excites the postsynaptic neuron is the US. Both stimuli drive their targets to fire at a constant rate of 45 Hz for 1 s. The interval between the US and the CS is defined as T , so that the CS precedes the US if $T < 0$ and follows it if $T > 0$. After stimulation, we assume that the firing rates are sustained but decay slowly back to zero with an exponential time constant of 2 s (Fig. 2 *Inset*).

Fig. 2 shows cross-correlations between pre- and postsynaptic spike trains with rates as shown in Fig. 2 *Inset* that decay exponentially after a sustained plateau, with the decay constants for the pre- and postsynaptic firing rates given by τ_{pre} and τ_{post} . Slowly decaying firing rates create asymmetries in these cross-correlations over the time range relevant for STDP, even when the stimuli that excite the neurons are separated by many seconds. Furthermore, the ordering of the initiation of firing is preserved in the ordering of spikes, as revealed by the different slopes in Fig. 2. When presynaptic firing is initiated before postsynaptic firing ($T < 0$), presynaptic spikes tend to lead postsynaptic spikes more often than the other way around, even on a millisecond timescale (Fig. 2 *Left*). The situation is reversed if postsynaptic firing is initiated before presynaptic activity ($T > 0$; Fig. 2 *Right*). This effect is due to the slow decay of the firing rate, not simply to the order of firing. Because the correlations in Fig. 2 are normalized and the firing rate decays are exponential, the same plots apply independent of the magnitude of T , but the normalization factor is proportional to $\exp(-|T|/\tau_{pre})$ when $T < 0$ and $\exp(-T/\tau_{post})$ when $T > 0$. Thus, the effect gets very small if T is large relative to τ_{pre} and τ_{post} . In general, these asymmetries are small, as in Fig. 2, but they are sufficient to evoke significant levels of synaptic plasticity.

To examine the effect of these cross-correlations on STDP, we drove a single integrate-and-fire model postsynaptic neuron with 1,000 presynaptic afferents. The excitatory synapses were subject to STDP given by the window function in Fig. 1A. Presynaptic afferents fired in response to stimulus 1 as in Fig. 2 *Inset* (i.e., for 1 s at 45 Hz followed by an exponential decay to zero with a time constant $\tau_{pre} = 2$ s). In this case, the US and the firing rate decay following it were represented by a current injected into the postsynaptic neuron that was constant for 1 s, producing firing at 45 Hz, and then decayed exponentially to zero with a time constant $\tau_{post} = 2$ s. This current thus contains contributions arising directly from the US and from whatever mechanism sustains the response.

For Fig. 3, we activated the CS 5 s before the US (as in the Fig. 2 *Inset*). Fig. 3A shows the situation at the beginning of our simulations. All synapses began at zero strength (histogram in Fig. 3A *Right*), so the postsynaptic neuron only responded to the US (membrane potential trace in Fig. 3A *Left*). However, after 20 trials, the synapses were potentiated by STDP due to the asymmetry in pre-post spike correlations shown in Fig. 2. As seen in Fig. 3B, the broadened distribution of synaptic strengths leads to a depolarization of the postsynaptic potential in response to the CS acting through the strengthened synapses. After 40 trials (Fig. 3C), enough synapses have been strengthened to generate robust responses to both stimuli. Due to its response to the CS, the postsynaptic neuron has "learned" the association between the two stimuli.

Fig. 4 shows the effect of interstimulus timing and ordering on the association learning. When the interstimulus interval is shortened to 4 s, an even stronger response to the CS is generated after 40 trials (Fig. 4A). When the interstimulus interval is lengthened to 6 s, there is insufficient potentiation after 40 trials to drive spikes, although the postsynaptic neuron is depolarized by the CS (Fig. 4B). Finally, if the order of the stimuli is reversed, so that the US drives the postsynaptic neuron before the CS excites its afferents, no strength-




Fig. 2. Cross-correlations resulting from decaying firing patterns. *Inset* shows the firing rates of two neurons (one presynaptic and one postsynaptic), and the two panels indicate the resulting cross-correlations, normalized to the value for synchronous spikes. (*Left*) When the decay of the firing of the presynaptic neuron coincides with the activation of the postsynaptic neuron, the pre–post cross-correlation shows more pre-before-post than post-before-pre spike ordering. (*Right*) When the decay of the firing of the postsynaptic neuron coincides with the activation of the presynaptic neuron, the pre–post cross-correlation shows post-before-pre more than pre-before-post spike ordering.

ening of the synapses occurs, and no postsynaptic response to the CS is generated (Fig. 4C). The reverse ordering produces weakening rather than strengthening of the synapses, but this is not evident in this simulation because the synapses started out at zero strength and could not be weakened.

The sustained but slowly decaying component of the firing of the postsynaptic neuron is not important for the effects shown in Figs. 3 and 4. Virtually identical effects can be produced by having the postsynaptic neuron fire only for the

1-s interval while the US is being presented. We include the postsynaptic firing decay, however, because it plays an important role in accounting for the results of Fig. 1B, as discussed below (see Fig. 1C). The period of constant firing we have assumed is not necessary, but increases the asymmetry in the cross-correlation and consequently increases the amount of synaptic change per trial. Finally, we have assumed that the effects of separate spike pairs due to STDP sum linearly. Nonlinear effects have been observed when both pre- and




Fig. 3. Development of a conditioned response. A CS drives the afferents to a postsynaptic neuron at 45 Hz from time 1s to 2s, followed by an exponential decay of afferent rates with a decay time constant of 2 s. The postsynaptic neuron is driven by a US, represented by a constant current from time 6–7 s that drives it at 45 Hz, followed by a 2-s time-constant exponential decay of the current to zero. (*Left*) Shows the membrane potential of the postsynaptic neuron. For clarity, the dots over the action potentials indicate every 11th action potential. (*Right*) Shows distributions of the strengths of the 1,000 excitatory synaptic conductances between the afferents and the postsynaptic neuron relative to their maximal allowed value. (A) At the beginning of the simulation, all synapses are set to zero strength, and the postsynaptic neuron responds only to the US. (B) After 20 trials, some of the synapses have strengthened causing the postsynaptic neuron to depolarize in response to the CS. (C) After 40 trials, the synapses have grown strong enough to make the postsynaptic neuron fire in response to both stimuli.




Fig. 4. Effect of timing and order on the conditioned response. The format and procedures are identical to those in Fig. 3 except for the timing and ordering between the CS and the US. All results are after 40 trials. (A) When the CS precedes the US by 4 s (rather than 5 s as in Fig. 3), synaptic strengths are stronger, and the postsynaptic neuron responds more strongly to the CS than in Fig. 3. (B) When the CS precedes the US by 6 s, the postsynaptic neuron is depolarized but fails to fire in response to the CS after 40 trials. Synaptic strengths are weaker than in Fig. 3 and A. (C) When the US precedes the CS by 5 s, no synaptic strengthening occurs.

postsynaptic neurons fire simultaneously >40 Hz (15, 16), which does not happen in our examples. Although it is not clear how to describe the nonlinear summation, we have tried to model it with saturating long-term potentiation (LTP) (16), and the effects shown in Figs. 3 and 4 were still found. In addition, we can reproduce these results with significantly lower rates as well. We can also introduce jitter into the postsynaptic response and preserve the effects. The essential correlations can even be induced in a Poisson model, as in Fig. 2.

To further clarify the phenomenon being studied, we computed the amount of synaptic modification analytically as a function of the interstimulus interval T , the pre- and postsynaptic firing rate decay constants, τ_{pre} and τ_{post} , and the parameters of the STDP window function (Fig. 1A). Details of this calculation are presented in *Methods*. To make these calculations tractable, we assume that the synaptic strengths are not close to their lower or upper limits, and we do not include the effects of the afferents on postsynaptic firing. Under these assumptions, when the stimulus driving the presynaptic neuron (the CS) precedes that driving the postsynaptic neuron (the US) ($T < 0$), synapses strengthen by an amount:

$$(A_+ + A_-)R_{\text{pre}}R_{\text{post}}\tau_+\tau_-\left(\frac{\tau_{\text{post}}}{\tau_{\text{pre}} + \tau_{\text{post}}}\right)\exp\left(-\frac{|T|}{\tau_{\text{pre}}}\right). \quad [1]$$

When the stimulus to the postsynaptic neuron (the US) is activated before that to the presynaptic neuron (the CS) ($T > 0$), synapses weaken by an amount:

$$-(A_+ + A_-)R_{\text{pre}}R_{\text{post}}\tau_+\tau_-\left(\frac{\tau_{\text{pre}}}{\tau_{\text{pre}} + \tau_{\text{post}}}\right)\exp\left(-\frac{T}{\tau_{\text{post}}}\right). \quad [2]$$

These equations indicate that the effect of stimulus pairing with a time interval T falls off exponentially as $|T|$ increases. This

dependence is similar to the falloff of STDP as the interspike interval increases, as seen in Fig. 1A. There is a critical difference, however. Whereas the falloff of STDP is governed by the millisecond-scale time constants τ_+ and τ_- , the falloff of the effect of the two stimuli on synaptic efficacy is governed by the much larger time constants τ_{pre} and τ_{post} that govern the slowly decaying responses of the pre- and postsynaptic neurons. The time constant in the exponential factor that determines how these changes vary with the interstimulus interval is, in fact, the firing decay time constant of the neuron that fired first (τ_{pre} if $T < 0$ and τ_{post} if $T > 0$). The factor in front of the exponential terms indicates that, in agreement with the experiment (Fig. 1B), an effect with a slower decay has a smaller magnitude. Matching the results of Fig. 1B requires considerably slower firing rate decays than those considered in Figs. 3 and 4, but with appropriate parameter choices, the amount of synaptic modification predicted by the above equations, plotted as a function of T (Fig. 1C), matches the data from Tanimoto *et al.* (1) quite well (Fig. 1B).

Discussion

In our model, asymmetric cross-correlations allow a fast correlation-based mechanism of synaptic plasticity to react to associations that involve timescales relevant to behavior. In the case we considered, sustained activity that decays away slowly allows associations to be made between events across timescales that are far greater than the width of the STDP window. In our simulations, we do not explicitly model the mechanism for sustaining activity or for making it decay. Reverberations in a network with high gain could generate such a pattern of activity if slightly detuned from the conditions required to sustain constant firing (18, 19). Slow cellular or synaptic mechanisms

could also contribute (20–22). Persistent activity lasting several seconds or more in response to a transient input has been observed in single neurons (23), in slice culture (24), in the locust antennal lobe (25), and in the prefrontal cortex of monkeys (26). Thus, it is not unlikely that an odor or shock will generate a long-lasting response.

There are a number of other ways to produce the asymmetric cross-correlations needed in the model. Up and down states (27, 28) might generate asymmetries if they occurred in an appropriate temporal sequence. More complicated cellular or network mechanisms that, for example, change the latency or order of spikes could generate correlations that favor potentiation or depression depending on the history of stimulation. Network mechanisms could also cause changes in the spontaneous background activity, altering correlations, as a form of off-line learning (29). Whatever the source, long time correlations that could support the type of learning we have presented have been observed in spike trains (30, 31), spike counts (32), and local field potentials (33) from many different areas.

Our model requires a relative fine balance between LTP and long-term depression (LTD) for the synaptic weights to respond to the asymmetric rather than the symmetric component of the unsubtracted cross-correlation. We found that reliable formation of temporally sensitive associative plasticity in the example of Fig. 3 requires the ratio of the area under the LTD part of the STDP window to the area under the LTP part ($A_{-}\tau_{-}/A_{+}\tau_{+}$) to be between 1.00 and 1.01. Such tuning is a universal feature in such cases when a small asymmetry (the asymmetry of the cross-correlation function) must be detected on top of a larger background. If tight control is not maintained on the balance between synaptic potentiation and depression due to STDP, uncorrelated background firing will have a dominate effect, and synapses will be pushed to the upper or lower limits of strength independent of the temporal correlations of the pre- and postsynaptic inputs. Mechanisms for dynamically regulating this balance might control it automatically without requiring fine tuning of parameters (34).

The model we studied requires tens of trials for the association to be made, but this number can be decreased simply by increasing the amplitude of the STDP. We have used a form of STDP in which hard boundaries are imposed at the limits of the range of allowed synaptic conductances, but other models with soft bounds have been proposed and studied (35–38). The upper boundary played no role in our simulations because it was never reached, but the lower boundary is needed to prevent synaptic conductances from becoming negative. It is not essential that this boundary is hard, but it is critical that the pre–postsynaptic correlations, not the boundary constraints, are the dominant factor in determining the fate of the synapses through STDP.

We considered only one synaptic relay, but the same mechanism active during the presentation of reliably repeated sequences of events to more complicated networks should lead to the formations of synfire chains (39, 40). The position of a neuron in such a chain will then correspond to the position of the event in the driving sequence. In this case, the mechanism that we propose would link the timescale of behavioral events to the much faster neuronal firing sequences of the synfire chain, resulting in a temporally compressed representation (41–43).

Methods

Correlations. The correlation shown in Fig. 2 is:

$$C(t) = \int_{-\infty}^{\infty} dt' r_{\text{post}}(t') r_{\text{pre}}(t' + t). \quad [3]$$

Simulations. The model we use for all of the simulations is similar to that described in ref. 10. The postsynaptic neuron shown in

Figs. 3 and 4 is described by an integrate-and-fire model with membrane potential V and driving current I satisfying:

$$\tau_m \frac{dV}{dt} = V_{\text{rest}} - V + g_{\text{ex}}(t)(E_{\text{ex}} - V) + I, \quad [4]$$

with $\tau_m = 20$ ms, $V_{\text{rest}} = -60$ mV, and $E_{\text{ex}} = 0$ mV. The synaptic conductance g_{ex} is the total conductance arising from 1,000 excitatory synapses, measured in units of the leakage conductance of the neuron. When a presynaptic spike arrives at synapse a , where $a = 1, 2, \dots, 1,000$, and $g_{\text{ex}}(t) \rightarrow g_{\text{ex}}(t) + \bar{g}_a$, where \bar{g}_a is the peak synaptic conductance of synapse a and falls in the range between zero and \bar{g}_{max} . In addition, g_{ex} decays exponentially to zero according to:

$$\tau_{\text{ex}} \frac{dg_{\text{ex}}}{dt} = -g_{\text{ex}}, \quad [5]$$

with $\tau_{\text{ex}} = 5$ ms.

During the simulations for Figs. 3 and 4 (with obvious modifications), the postsynaptic neuron was driven by a current:

$$\begin{aligned} I(t) &= 0 && \text{for } t < T + 1 \text{ s} \\ I(t) &= 15 \text{ mV} && \text{for } T + 1 \text{ s} < t < T + 2 \text{ s} \\ I(t) &= (15 \text{ mV}) \cdot \exp(-(t - T - 2 \text{ s})/\tau_{\text{pre}}) && \text{for } t > T + 2 \text{ s} \end{aligned} \quad [6]$$

(a factor of the membrane resistance has been absorbed into the definition of I , which is why these results are given in millivolts). Presynaptic spikes were generated by Poisson processes acting independently at each of the 1,000 synapses at a rate given by:

$$\begin{aligned} r(t) &= 0 && \text{for } t < 1 \text{ s} \\ r(t) &= 45 \text{ Hz} && \text{for } 1 \text{ s} < t < 2 \text{ s} \\ r(t) &= (45 \text{ Hz}) \exp(-(t - 2 \text{ s})/\tau_{\text{post}}) && \text{for } t > 2 \text{ s} \end{aligned} \quad [7]$$

with $\tau_{\text{pre}} = \tau_{\text{post}} = 2$ s. In all cases, identical trials, described as above, were repeated in blocks lasting 10 s each.

STDP is modeled as ref. 10 by introducing the functions $M(t)$ and $P_a(t)$ for $a = 1, 2, \dots, N$ satisfying:

$$\tau_{-} \frac{dM}{dt} = -M \quad \text{and} \quad \tau_{+} \frac{dP_a}{dt} = -P_a, \quad [8]$$

with $\tau_{+} = \tau_{-} = 20$ ms. $M(t)$ is decremented by an amount A_{-} every time the postsynaptic neuron fires an action potential, and $P_a(t)$ is incremented by an amount, A_{+} , every time synapse a receives an action potential. As shown in Fig. 1A, we took $A_{+} = A_{-} = 0.005$. In addition, when synapse a receives a presynaptic action potential at time t , its maximal conductance parameter is updated by $\bar{g}_a \rightarrow \bar{g}_a + M(t)\bar{g}_{\text{max}}$. If this change makes $\bar{g}_a < 0$, \bar{g}_a is set to zero. If the postsynaptic neuron fires an action potential at time t , \bar{g}_a is incremented by $\bar{g}_a \rightarrow \bar{g}_a + P_a(t)\bar{g}_{\text{max}}$. If this change would make $\bar{g}_a > \bar{g}_{\text{max}}$, \bar{g}_a would be set to \bar{g}_{max} , but this situation never actually happens in our simulations.

Analytic Calculation. For pre- and postsynaptic firing rates $r_{\text{pre}}(t)$ and $r_{\text{post}}(t)$, the amount of LTP over one trial is given by:

$$\begin{aligned} \text{LTP} &= A_{+} \int_{-\infty}^{\infty} dt_{\text{post}} r_{\text{post}}(t_{\text{post}}) \\ &\cdot \int_{-\infty}^{t_{\text{post}}} dt_{\text{pre}} r_{\text{pre}}(t_{\text{pre}}) \exp(-(t_{\text{post}} - t_{\text{pre}})/\tau_{+}). \end{aligned} \quad [9]$$

Similarly, the amount of LTD is:

$$\text{LTD} = A_- \int_{-\infty}^{\infty} dt_{\text{post}} r_{\text{post}}(t_{\text{post}}) \cdot \int_{t_{\text{post}}}^{\infty} dt_{\text{pre}} r_{\text{pre}}(t_{\text{pre}}) \exp(-(t_{\text{pre}} - t_{\text{post}})/\tau_-). \quad [10]$$

The total change is the difference between these two terms.

We take the presynaptic rate, r_{pre} , to be exponentially decaying with time constant τ_{pre} from an initial value R_{pre} and zero for $t < 0$. We take the postsynaptic rate to be zero for $t < T$ and, for $t \geq T$, $r_{\text{post}}(t) = R_{\text{post}} \exp(-(t - T)/\tau_{\text{post}})$. The analytic result reported in the text comes from integrating the equations for LTP and LTD given above with these rates. To simplify the resulting equations, we have taken the area under the STDP curve to be zero ($A_- \tau_- = A_+ \tau_+$), and we have removed terms that are negligibly small when, as is the case here, τ_{pre} , τ_{post} , and $|T|$ are all much larger than τ_+ and τ_- . The condition $A_- \tau_- =$

$A_+ \tau_+$ is not required for the results we report; the requirement is actually $\tau_+ \tau_- (A_+ + A_-) > \tau_{\text{pre}} (A_- \tau_- - A_+ \tau_+)$, which requires $A_- \tau_- - A_+ \tau_+$ to be small but not necessarily zero.

A useful form of the above results (8, 9, 11) can be obtained if the pre- and postsynaptic firing rates vary slowly on the scale of τ_+ and τ_- . Then, by Taylor expanding r_{pre} and r_{post} , we find that the total plasticity $\Delta = \text{LTP} - \text{LTD}$ is given by:

$$\Delta = (A_+ \tau_+ - A_- \tau_-) \int_{-\infty}^{\infty} dt r_{\text{post}}(t) r_{\text{pre}}(t) - (A_+ \tau_+^2 + A_- \tau_-^2) \int_{-\infty}^{\infty} dt r_{\text{post}}(t) r'_{\text{pre}}(t), \quad [11]$$

where r'_{pre} is the time derivative of r_{pre} .

We thank S. Fusi and S. Nelson for valuable discussions and J. Lisman for encouragement. This research was supported by National Institutes of Health (NIH) Grant MH58754, the Swartz Foundation, and an NIH Director's Pioneer Award, part of the NIH Roadmap for Medical Research, through Grant 5-DP1-OD114-02.

1. Tanimoto, H., Heisenberg, M. & Gerber, B. (2004) *Nature* **430**, 983.
2. Bi, G. Q. & Rubin J. (2005) *Trends Neurosci.* **28**, 222–228.
3. Schultz, W. & Dickinson, A. (2000) *Annu. Rev. Neurosci.* **23**, 473–500.
4. Mehta, M. R., Barnes, C. L. & McNaughton B. L. (1997) *Proc. Natl. Acad. Sci. USA* **94**, 8918–8921.
5. Mehta, M. R., Quirk, M. C. & Wilson M. (2000) *Neuron* **25**, 707–715.
6. Blum, K. I. & Abbott, L. F. (1996) *Neural Comput.* **8**, 85–93.
7. Abbott, L. F. & Blum K. I. (1996) *Cereb. Cortex* **6**, 406–416.
8. Kempter, R., Gerstner, W. & van Hemmen, J. L. (1999) *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.* **59**, 4498–4514.
9. Roberts, P. (1999) *J. Comput. Neurosci.* **7**, 235–246.
10. Song, S., Miller, K. M. & Abbott, L. F. (2000) *Nat. Neurosci.* **3**, 919–926.
11. Xie, X. H. & Seung, H. S. (2000) *Adv. Neural Inf. Process. Sys.* **12**, 199–205.
12. Markram, H., Lubke, J., Frotscher, M. & Sakmann, B. (1997) *Science* **275**, 213–215.
13. Bi, G. Q. & Poo, M. M. (1998) *J. Neurosci.* **18**, 10464–10472.
14. Feldman, D. E. (2000) *Neuron* **27**, 45–56.
15. Sjostrom, P. J., Turrigiano, G. & Nelson, S. B. (2001) *Neuron* **32**, 1149–1164.
16. Froemke, R. C. & Dan, Y. (2002) *Nature* **416**, 433–438.
17. Celikel, T., Szostak, V. & Feldman, D. E. (2004) *Nat. Neurosci.* **7**, 534–541.
18. Seung, H. S. (1996) *Proc. Natl. Acad. Sci. USA* **93**, 13339–13344.
19. Wang, X. J. (2001) *Trends Neurosci.* **24**, 455–463.
20. Thorson, J. & Biederman-Thorson, M. (1974) *Science* **183**, 161–172.
21. Schwindt, P. C., Spain, W. J. & Crill, W. E. (1992) *J. Neurophysiol.* **67**, 216–226.
22. Zucker, R. S. & Regehr, W. G. (2002) *Annu. Rev. Physiol.* **64**, 355–405.
23. Egorov, A. V., Hamam, B. N., Franssen, E., Hasselmo, M. E. & Alonso, A. A. (2002) *Nature* **420**, 133–134.
24. Lau, P. M. & Bi, G. Q. (2005) *Proc. Natl. Acad. Sci. USA* **102**, 10333–10338.
25. Mazor, O. & Laurent, G. (2005) *Neuron* **48**, 661–673.
26. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. (1989) *J. Neurophys.* **61**, 331–349.
27. Galan, R. F., Weidert, M., Menzel, R., Herz, A. V. & Galizia, C. G. (2006) *Neural Comput.* **18**, 10–25.
28. Wilson, C. J. & Kawaguchi, Y. (1996) *J. Neurosci.* **16**, 2397–2410.
29. Anderson, J., Lampl, I., Reichova, I., Carandini, M. & Ferster, D. (2000) *Nat. Neurosci.* **3**, 617–621.
30. Lowen, S., Ozaki, T., Kaplan, E., Saleh, B. E. & Teich, M. C. (2001) *Methods* **24**, 377–394.
31. Lowen, S. & Teich, M. (1996) *J. Acoust. Soc. Am.* **99**, 3585–3591.
32. Bair, W., Zohary, E. & Newsome, W. T. (2001) *J. Neurosci.* **21**, 1676–1697.
33. Leopold, D., Murayama, Y. & Logothetis, N. K. (2003) *Cereb. Cortex* **13**, 422–433.
34. Kepecs, A., van Rossum, M. C., Song, S. & Tegner, J. (2002) *Biol. Cybern.* **87**, 446–458.
35. van Rossum, M. C., Bi, G. Q. & Turrigiano, G. G. (2000) *J. Neurosci.* **20**, 8812–8821.
36. Rubin, J. E. (2001) *Network* **12**, 131–140.
37. Rubin, J., Lee, D. D. & Sompolinsky, H. (2001) *Phys. Rev. Lett.* **86**, 4958–4961.
38. Gutig, R., Aharonov, R., Rotter, S. & Sompolinsky, H. (2003) *J. Neurosci.* **23**, 3697–3714.
39. Buonomano, D. V. (2003) *Proc. Natl. Acad. Sci. USA* **100**, 4897–4902.
40. Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D. & Yuste, R. (2004) *Science* **304**, 559–564.
41. Buzsaki, G. (1989) *Neuroscience* **31**, 551–570.
42. Wilson, M. A. & McNaughton, B. L. (1994) *Science* **265**, 676–679.
43. Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J. & Buzsaki, G. (1999) *J. Neurosci.* **19**, 9497–9507.