Spike-timing dependent plasticity

J. Sjöström and W. Gerstner
Scholarpedia, 5(2):1362 (2010)

Spike Timing Dependent Plasticity (STDP) is a temporally asymmetric form of Hebbian learning induced by tight temporal correlations between the spikes of pre- and postsynaptic neurons. As with other forms of synaptic plasticity, it is widely believed that it underlies learning and information storage in the brain, as well as the development and refinement of neuronal circuits during brain development (e.g. Bi and Poo, 2001; Sjöström et al., 2008). With STDP, repeated presynaptic spike arrival a few milliseconds before postsynaptic action potentials leads in many synapse types to long-term potentiation (LTP) of the synapses, whereas repeated spike arrival after postsynaptic spikes leads to long-term depression (LTD) of the same synapse. The change of the synapse plotted as a function of the relative timing of pre- and postsynaptic action potentials is called the STDP function or learning window and varies between synapse types. The rapid change of the STDP function with the relative timing of spikes suggests the possibility of temporal coding schemes on a millisecond time scale.

Contents

1.1 Experimental STDP Protocol

In a typical STDP protocol (Markram and Sakmann, 1997; Bi and Poo, 1998; Sjöström et al., 2001), a synapse is activated by stimulating a presynaptic neuron (or presynaptic pathway) shortly before or shortly after making the postsynaptic neuron fire by injection of a short current pulse. The pairing is repeated for 50-100 times at a fixed frequency (for example 10 Hz). The weight of the synapse is measured as the amplitude (or initial slope) of the postsynaptic potential. The change of the synaptic weight is plotted as a function of the relative timing between presynaptic spike arrival and postsynaptic firing, see Fig.1. The resulting plot is the STDP function or learning window. It is worth noting that different synapse types can have quite different forms of STDP function (Abbott and Nelson, 2000; Bi and Poo, 2001). Compared to many other synaptic plasticity induction protocols, STDP is especially attractive since it is believed to be biologically plausible. In the intact brain, action potentials are often quite precisely timed to stimuli in the outside world, although this is not true for all

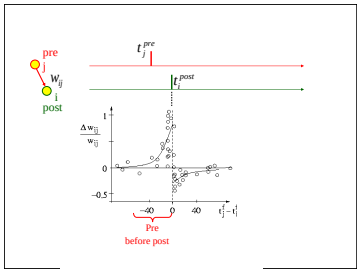

Fig. 1.1: The STDP function shows the change of synaptic connections as a function of the relative timing of pre- and postsynaptic spikes after 60 spike pairings. Schematically redrawn after Bi and Poo (1998).

brain regions and cell types. Nevertheless, STDP is very likely to be induced under such circumstances and many studies provide strong evidence that this is indeed the case (Zhang et al., 1998; Allen et al., 2003; Jacob et al., 2007; Meliza and Dan, 2006).

1.2 Basic STDP Model

The weight change Δw_j of a synapse from a presynaptic neuron j depends on the relative timing between presynaptic spike arrivals and postsynaptic spikes. Let us name the presynaptic spike arrival times at synapse j by t_j^f where $f = 1, 2, 3, ...$ counts the presynaptic spikes. Similarly, t^n with $n = 1, 2, 3, ...$ labels the firing times of the postsynaptic neuron. The total weight change Δw_j induced by a stimulation protocol with pairs of pre- and postsynaptic spikes is then (Gerstner and al. 1996, Kempter et al. 1999)

$$\Delta w_j = \sum_{j=1}^{N} \sum_{n=1}^{N} W(t_i^n - t_j^f) \tag{1.1}$$

where $W(x)$ denotes one of the STDP functions (also called learning window) illustrated in Fig.1.1.

A popular choice for the STDP function $W(x)$

$$W(x) = A_+ \exp(-x/\tau_+) \qquad for \quad x > 0 \tag{1.2}$$

$$W(x) = -A_- \exp(x/\tau_-) \qquad for \quad x < 0 \tag{1.3}$$

which has been used in fits to experimental data (Zhang et al. 1998) and models (e.g, Song et al. 2000). The parameters A_+ and A_- may depend on the current value of the synaptic weight w_j. The time constants are on the order of $\tau_+ = 10ms$ and $\tau_- = 10ms$

1.3 Variants of STDP Models

1.3.1 Online implementation of STDP models

Spike-Timing Dependent Plasticity with an STDP function as in Eq. (1.2) can be implemented in an on-line update rule using the following assumptions. Each presynaptic spike arrival leaves a trace $x_j(t)$ which is updated by an amount $a_+(x)$ at the moment of spike arrival and decays exponentially in the absence of spikes:

$$\tau_+ \frac{dx_j}{dt} = -x + a_+(x) \sum_j \delta(t - t_j^f) \tag{1.4}$$

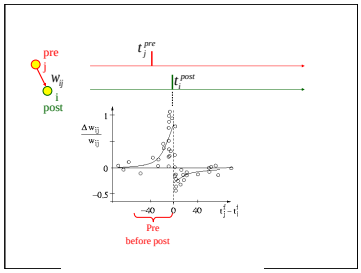

Fig. 1.2: Spike-Timing Dependent Plasticity can be implemented by local variables. Top: A presynaptic spike leaves a trace $x_j(t)$ which is read out (arrow) at the moment of the postsynaptic spike. The weight change is proportional to that value $x_j(t_n)$ Bottom: A postsynaptic spike leaves a trace $y(t)$ which is read out (arrow) at the moment of a presynaptic spike.

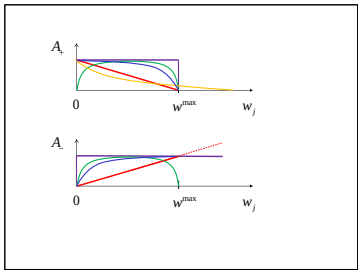

Fig. 1.3: Top: Potentiation remains bounded if the parameter $A_+(w_j)$ has hard bounds (magenta), linear soft bounds (red), a nonlinear soft-bound at w^{max}(blue), or two-sided nonlinear soft bounds (green). With the yellow weight-dependence, potentiation is the smaller the larger the weight, but does not have a fixed upper bound. Bottom: Analogously, depression is stopped at zero weights choosing some bounds for $A_-(w_j)$.

The biophysical nature of the variable x need not to be specified, but potential candidates are the amount of glutamate bound to postsynaptic receptors; or the fraction of NMDA receptors in the open state. Similarly, each postsynaptic spike leaves a trace y

$$\tau_- \frac{dy}{dt} = -y + a_-(x) \sum_n \delta(t - t^n) \tag{1.5}$$

which increases by an amount $a_-(y)$ at the moment of postsynaptic spikes. This trace could possibly be interpreted as the voltage at the synapse caused by a backpropagating action potential; or by calcium entry due to a backpropagating action potential. The weight change is then

$$\frac{dw_j}{dt} = A_+(w_j)x(t) \sum_n \delta(t - t^n) - A_-(w_j)y(t) \sum_f \delta(t - t_j^f) \tag{1.6}$$

Thus, the weight is increased at the moment of postsynaptic firing by an amount that depends on the value of the trace x left by the presynaptic spike. Similarly, the weight is depressed at the moment of presynaptic spikes by an amount proportional to the trace y left by previous postsynaptic spikes. Integration of Eq. (1.6) yields Eq. (1.2). For an illustration see Fig.1.2

1.3.2 Weight dependence: hard bounds and soft bounds

For biological reasons, it is desirable to keep the synaptic weights in a range $w^{min} < w_j < w^{max}$. This can be achieved by an appropriate choice of the functions $A_+(w_j)$

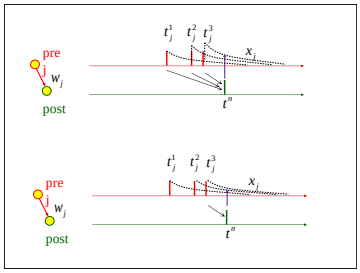

Fig. 1.4: Top: All three presynaptic spikes interact (black arrows) with a later postsynaptic spike (all-to-all interaction). This can be implemented by a trace $x_j(t)$ which accumulates. Bottom: Only the most recent presynaptic spike interacts (black arrow) with a later postsynaptic spike (nearest-neighbor interaction). This can be implemented by a trace $x_j(t)$ which starts after each presynaptic spike at the same value.

and $A_-(w_j)$. For the sake of simplicity, the lower bound is set in most models to zero, $w^{min} = 0$. A choice: $A_+(w_j) = (w^{max} - w_j)\eta_+$

and $A_-(w_j) = w_j\eta_-$ with positive constants η_+ and η_- is called soft bounds or multiplicative weight dependence. The choice :$A_+(w_j) = \Theta(w^{max} - w_j)\eta_+$ and $A_-(w_j) = \Theta(-w_j)\eta_-$ is called hard bounds, see Fig.1.3. Here $\Theta(x)$ denotes the Heaviside step function. In practice, hard bounds mean that an update rule with fixed parameters η_+ and η_- is used until the bounds are reached (Gerstner et al., 1996, Kempter et al., 1999, Roberts et al., 1999, Song et al., 2000). Soft bounds mean that, for large weights, synaptic depression dominates over potentiation (Kistler and van Hemmen, 2000; van Rossum et al., 2000, Rubin et al., 2001). It is possible to interpolate between the two cases (Guetig et al., 2003).

1.3.3 Temporal all-to-all versus nearest-neighbor spike-interaction

If the sum in Eq. (1.2) goes over all presynaptic spike arrivals and all postsynaptic spikes, then all spike pairs contribute equally. This case has been called all-to-all spike interaction (Fig. 1.4). It is also possible to restrict the interactions so that only nearest spikes interact. In the mechanistic update rule of Eq. (1.6), nearest-neighbor interaction can be implemented as follows. The potentiation at the moment of the postsynaptic spike should depend only on the time since the most recent presynaptic spike. To achieve this, suppose that the trace variable x is increased at the moment of presynaptic spikes by an amount $a_+(x) = 1 - x_-$ where x_- denotes the value of the variable x just before the update. In other words, the update of x is not cumulative but goes always to a fixed value of one, so that the influence of previous spikes is cancelled; see Morrison et al. (2008) for a review.

1.3.4 Triplet rule of STDP

Pair-wise interaction between spikes as in Eq. (1.1) would predict that 60 repetitions of pre-post pairings (say, presynaptic spikes 10 ms before postsynaptic ones) give the same result independent of whether the pairing is repeated at 1 Hz or 5Hz. At frequencies above 25 Hz, a pair-wise interaction model would predict a reduced potentiation, since in addition to the pre-post pair at 10ms virtual post-pre pairs at 30ms are created - that should lead to depression. However, in experiments the opposite is observed (Senn et al., 2001; Sjöström et al., 2001). The frequency-dependence of STDP experiments can be accounted for if one assumes that the basic building block of potentiation during STDP experiments is not a pair-wise interaction as assumed in Eq. (1.1), but a triplet interaction between two postsynaptic spikes and one presynaptic spike (see Fig.1.5).

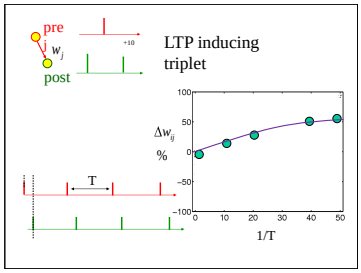

Fig. 1.5: Top: In a triplet model, the elementary building block of LTP is not a pair, but a combination of 1 pre and 2 postsynaptic spikes. Bottom: Frequency dependence of LTP. The same number of pre-post pairings at 10ms is repeated at different intervals T. On the right, the amount of LTP is given as a function of the repetition frequency 1/T, redrawn after Sjöström et al. (2001)

Such a triplet interaction can be implemented in the mechanistic model if one works with two postsynaptic traces y_1 and y_2 with two different time constants, rather than a single trace (Pfister and Gerstner, 2006). Such a model is also compatible with explicit triplet experiments (Wang et al., 2005) while a pair-based model is not.

1.3.5 Homeostatic terms

In addition to the pair-based and triplet-based STDP effects mentioned above, one can also consider STDP models where an isolated postsynaptic or presynaptic spike induces a small change of the synaptic weight, even if not paired with another spike. These terms can be used in models to yield a homeostatic control of the total input into the postsynaptic neuron (Kempter et al., 1999, van Rossum, 2001).

Another possibility to implement homeostasis into STDP models is by making the parameter A_- in Eq. (1.2) depend on the mean firing rate calculated as a running average over a time scale of seconds (Pfister and Gerstner, 2006).

1.3.6 Voltage dependence

Experiments and models of Spike-Timing Dependent Plasticity suggest that synaptic weight changes are caused by the tight temporal correlations between pre- and postsynaptic spikes. However, other experimental protocols where presynaptic spikes are paired with a fixed depolarization of the postsynaptic neuron (e.g. under voltage clamp) show that postsynaptic spikes are not necessary to induce long-term potentiation and depression of the synapse (Artola et al., 1990; Ngezahayo et al., 2000; Sjöström et al., 2004). Moreover, the voltage of the postsynaptic neuron just before generation of action potentials influences the direction of change of the synapse, even if the spike timing is held fixed (Sjöström et al., 2001), suggesting that postsynaptic voltage is more fundamental than spike timing. Indeed, a model of synaptic plasticity that postulates pairing between presynaptic spike arrival and postsynaptic voltage contains STDP models as a special case (Brader et al., 2007, Clopath et al., 2008).

1.3.7 Biophysical models

Since signaling chains involved in the induction of Long-Term Potentiation and Depression are partially unknown, most models of STDP are phenomenological models. However, some models attempt to identify variables such as the traces x and y in the above mechanistic model with specific biophysical quantities. A few examples:

* Senn-Markram-Tsodyks model. The model shares features with the mechanistic triplet model above and identifies some of the variables with internal states of the

NMDA receptor and unspecified second messengers (Senn et al., 1997, 2001).

* Karmarkar-Buonomano model. The model emphasizes the fact that the pathways for upregulation and downregulation are independent and give interpretations of internal variables in terms of NMDA receptor, calcium, and backpropagating action potentials (Karmarkar and Buonomano, 2002).

* Shouval model. The model of Shouval starts from the hypothesis that the intracellular calcium concentration in the vicinity of the synapse controls the up- and downregulation of synaptic weights (Shouval et al., 2002).

* Rubin et al. model. The model gives a detailed account of some of the signaling steps translating the calcium time course into synaptic weight changes (Rubin et al., 2005).

* Lisman model. The model focuses on the autophosphorylation of CaMKII as a critical step for memory formation (Lisman and Zhabotinsky, 2001; Lisman 2003). The calcium based model can be simplified and shows STDP (Graupner and Brunel, 2007).

1.4 Relation of STDP to other learning rules

1.4.1 STDP and Hebbian learning rules

STDP can be seen as a spike-based formulation of a Hebbian learning rule. Hebb formulated that a synapse should be strengthened if a presynaptic neuron 'repeatedly or persistently takes part in firing' the postsynaptic one (Hebb 1949). This formulation suggests a potential causal relation between the firing of the two neurons. Causality requires that the presynaptic neuron fires slightly before the postsynaptic one. Indeed, in standard STDP experiments on synapses onto pyramidal neurons, potentiation of the synapse occurs for pre-before-post timing, in agreement with Hebb's postulate. Hebb did not, however, postulate the existence of synaptic weakening (Hebb 1949). The existence of a temporal window for weakening of connective strength in the typical STDP learning rule is in a sense an extension to the Hebbian postulate.

The existence of synaptic weakening, however, was postulated long before the discovery of STDP. Stent argued already in 1973 that the input from a presynaptic cell that is consistently not co-active with the postsynaptic cell should be weakened (Stent, 1973). One important distinction as compared to STDP is that the Stentian extension to Hebb's postulate does not emphasize temporal contrast, only persistent lack of coincidence (Stent 1973), and it is therefore more akin to heterosynaptic LTD than to STDP (Sjöström et al., 2008). Standard STDP, on the other hand, possesses a characteristic temporal asymmetry (Fig. 1; Caporale and Dan, 2008; Abbott and Nelson, 2000; Sjöström et al., 2008).

1.4.2 STDP versus Rate based learning rules

Under the assumption of stationary Poisson statistics for the firing times of pre- and postsynaptic neurons, the most relevant aspect of the STPD function is its integral and an STDP model can mapped to an equivalent rate-based learning rule (Kempter et al., 1999). Under the assumption of independence between pre- and postsynaptic firing, the total weight change is $\Delta w_{ij} = \alpha f_i(t) f_j(t)$ where $f_j(t)$ and $f_i(t)$ denote the firing rate of pre- and postsynaptic neurons averaged over some time T and $\alpha = \int W(s) ds$ is the integral over the learning window. If the integral is positive, STDP is identical to standard rate-based Hebbian learning. For negative integral, as often used in modeling, STDP corresponds to a anti-Hebbian rate rule.

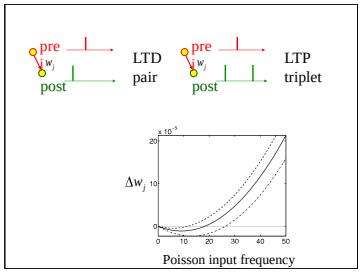

Fig. 1.6: A STDP rule (top) where post-pre pairs cause LTD and post-pre-post triplets cause LTP generates for Poisson input a frequency dependence of weight changes as in the BCM model (bottom).

However, the assumption of independence of pre- and postsynpatic firing is obviously wrong since it neglects the causal correlations generated by the interaction of the two neurons. A more precise mapping to rate models can be achieved if the postsynaptic neuron is described as an inhomogeneous Poisson Process with a rate $f_i(t) = \gamma \sum_j \sum_f \epsilon(t - t_j^f)$ where t_j^f denotes the spike times of a presynaptic neuron j generated by a Poisson process of rate $f_j(t)$ and $\epsilon(s)$ for $s > 0$ describes the time course of a postsynpatic potential. The total weight change in a period T is then $\Delta w_{ij} = \alpha f_i(t) f_j(t) + \beta f_j(t)$ where $\beta = \gamma \int_0^\infty \epsilon(s) W(s) ds$ is the integral over the 'causal' part of the learning window, i.e., over all times with 'pre-before-post' relation (Kempter et al. 1999). For standard STDP models $\beta > 0$, i.e., presynaptic spike arrival leads on average to a positive change of the synapse, because it is likely to cause postsynaptic firing. This is then often combined with a negative integral over the STDP function $\alpha < 0$ so that random pairings of pre- and postsynaptic firings leads to a decrease of the synapse (Gerstner et al., 1996, Song et al., 2000). The functional consequences of such a choice are discussed below (see Rate normalization).

1.4.3 STDP and Bienenstock-Cooper-Munro (BCM) rule

STDP can also be related to a nonlinear rate model where the weight change depends linearly on the presynaptic rate, but nonlinearly on the postsynaptic rate (Bienenstock et al., 1982). This can be achieved in two different ways. The first possibility is to implement standard STDP with nearest-neighbor instead of all-to all coupling (see above). This leads to a nonlinearity consistent with the BCM rule (Izikhevich and Desai, 2003). The second possibility is to use a triplet STDP rule (see above) instead of the standard pair-based rule (Fig. 1.6). If potentiation requires a triplet of two postsynaptic spike and one presynaptic spike (with post-pre-post or pre-post-post firings in temporal proximity) while depression is modeled by the interaction of a post-pre-pair, then the equivalent rate model under a Poisson firing assumption as above is $\Delta w_{ij} = a_+ [f_i]^2 f_j - a_- f_i f_j = \phi(f_i - \vartheta) f_j$ where ϑ describes the minimal postsynaptic frequency for potentiation and ϕ is a quadratic function (Pfister and Gerstner, 2006). If the amount a_- of depression increases with the mean postsynaptic frequency, then the threshold shifts with the mean postsynaptic rate. In this case the triplet rule of STDP becomes identical to BCM rule (Bienenstock et al. 1982).

1.5 Functional Consequences

As described above, STDP models can be related to rate models under the assumption of Poisson firing of both pre- and postsynaptic neurons. Hence STDP rules inherit functional consequences known for rate models. In particular, the potential of synaptic

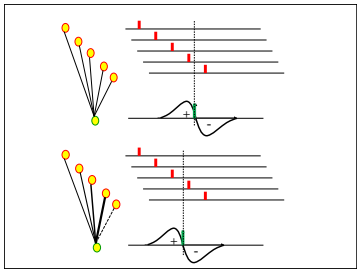

Fig. 1.7: Top: Presynaptic neurons are firing one after the other and cause the postsynaptic neuron to fire a single action potential (green vertical bar). The STDP function will strengthen those synapses that have been activated just before the postsynaptic spike. Bottom: If the stimulation pattern repeats the new synaptic weights make the postsynaptic neuron fire earlier.

learning to principal component analysis; to receptive field development; to clustering and map formation does not change fundamentally if one switches from rate-based to spike-based models (Kempter et al., 1999; Song and Abbott, 2001). In the rest of this section we focus on aspects that are specific to STDP and go beyond known features of rate-based learning.

1.5.1 Spike-spike correlations

The postsynaptic depolarization caused by spike arrival at an excitatory synapse makes the postsynaptic neuron more likely to fire. In all spiking neuron models (including Poisson models driven by presynaptic input) this leads to a correlation of the spikes of pre- and postsynaptic neurons on the timescale of milliseconds. These spike-spike correlations contribute to learning in STDP models (Kempter et al., 1999), but are completely neglected in standard rate models of learning. See the section 'STDP versus rate based learning rules'.

1.5.2 Reduced latency

Suppose a postsynaptic neuron is connected to N presynaptic neurons that fire one after another in a sequence 1-2-3-...-N over several milliseconds; see Fig. 1.7. Suppose that the synaptic input makes the postsynaptic neuron fire between the firings of presynaptic neurons N-1 and N. As a result of STDP the connection from neuron N to the postsynaptic neuron is weakened (because of the post-before-pre timing) whereas the connections from neurons N-1, N-2, N-3 ... to the postsynaptic neuron are reinforced (because of appropriate pre-before-post firing). After several repetitions of the same stimulus, the postsynaptic neuron fires earlier, i.e. with reduced latency, because of the stronger input. Hence the timing of the postsynaptic spike shifts forward in time (Song et al., 2000, Mehta et al., 2000).

1.5.3 Temporal coding

Since STDP is sensitive to spike timing on the millisecond rate, it can be used in temporal coding paradigms. Examples include tuning of synaptic connections in a model of sound source localisation in the auditory pathway (Gerstner et al., 1996); learning of spatio-temporal spike patterns in a model of associative memory (Gerstner et al., 1993); suppression of predictable signals in a model of the electrosensory system of electric fish (Roberts and Bell, 2000); learning time-order codes (Guyonneau et al. 2005); amongst others.

1.5.4 Rate normalization

Rate-based Hebbian learning is intrinsically unstable: synaptic inputs that drive the neuron to a high firing rate will be strengthened further. On one hand, such an instability is necessary to make the neuron detect, and become sensitive, to weak correlations in the input. On the other hand, this leads not only to a growth of individual synapses, but also to an explosion of the firing rate of the postsynaptic neuron. In practice, in rate based learning the growth of synapses and firing rates is controlled by (i) introducing upper and lower bounds for individual weights and (ii) renormalization of the weights of each time step or each episode. Renormalization can alternatively be implemented online by a rate-dependent decay term of the weights (Oja 1982). Surprisingly, STDP models with an appropriate set of parameters do not need such an explicit normalisation step (Kempter et al., 1999, Song et al., 2000, Kempter et al., 2001).

As discussed above in section 'STDP versus Rate based learning rules', the equivalent rate model of a standard STDP rule is $\Delta w_{ij} = \alpha f_i(t) f_j(t) + \beta f_j(t)$. For a choice of parameter where the integral over the STDP function is negative ($\alpha < 0$) and pre-before-post firings lead to potentiation ($\beta > 0$), the firing rate of the postsynaptic neuron has a stable fixed point, while the learning rule is sensitive to the temporal correlations between pre- and postsynaptic neurons (Kempter et al., 2001).

1.6 Experimental results and open questions

1.6.1 Diversity of STDP

STDP varies tremendously across synapse types and brain regions (Abbott and Nelson, 2000). Even so, it is worth recollecting that the temporal asymmetry of classical STDP is also remarkably well preserved and is found in species as different as rat, frog, locust, zebra finch, cat, and probably also humans (reviewed in Sjöström et al., 2008; Caporale and Dan, 2008). In mammals, STDP has also been uncovered in multiple brain regions, such as prefrontal, entorhinal, somatosensory, and visual cortices, hippocampus, striatum, the cochlear nucleus, and the amygdala (cf. Sjöström et al., 2008; Caporale and Dan, 2008). The activity requirements that govern STDP at many of these different synapses, however, is variable. For example, the width of the temporal windows for LTD and LTP appear to be roughly equal at hippocampal excitatory synapses (Bi and Poo, 1998; Nishiyama et al., 2000; Zhang et al., 1998), whereas the LTD timing window is considerably wider than that of LTP at several neocortical synapses (Feldman et al., 2000; Sjöström et al., 2001).

For some synapses, the STDP timing windows is inverted as compared to the classical form of STDP, so that pre-before-post timings result in LTD whereas the opposite temporal order results in LTP. This is the case at e.g. inhibitory connections onto neocortical L2/3 pyramidal neurons (Holmgren and Zilberter, 2001), at corticostriatal synapses (Fino et al., 2005) as well as in the electrosensory lobe of the mormyrid electric fish (Bell et al., 1997). The timing requirements for STDP at connections between spiny stellate cells in rat somatosensory cortex are yet again different: Here, synapses undergo LTD seemingly regardless of temporal order (Egger et al., 1999). In neocortical layer-5 pyramidal neurons, the timing requirements also depend critically on synapse location in the dendritic tree: Whereas proximal inputs undergo classical STDP, distal synapses are subject to a "temporally inverted" STDP rule (Letzkus et al., 2006). These same inputs also undergo non-Hebbian LTD or Hebbian LTP depending on the state of depolarization of the apical dendrite (Sjöström and Häusser, 2006).

The activity requirements of STDP may thus vary considerably not only across brain regions and synapse types, but also within a cell, in different dendritic compart-

ments. One open question is what this variability is good for. Since it is well established that synaptic plasticity underpins neural circuit development (Katz and Shatz, 1996), this implies that the STDP rules engaged during development determine circuit functionality in the mature brain. In other words, this variability of STDP is most likely not coincidental.

1.6.2 Biophysical and biochemical mechanisms

Both LTP and LTD depend on intracellular calcium transients: LTP is triggered by brief and strong postsynaptic calcium events, whereas LTD is induced by smaller, more prolonged calcium elevations, a concept known as the calcium hypothesis in synaptic plasticity (Sjöström et al., 2008). This calcium dependence of plasticity is critically dependent on the activation of postsynaptic NMDA receptors residing in the spine: These NMDA receptors detect the coincidence of glutamate release due to the presynaptic spike and depolarization due to the postsynaptic spike, resulting in a supralinear rise in postsynaptic calcium during LTP (Yuste and Denk, 1995; Koester and Sakmann, 1998; Schiller et al, 1998). The calcium hypothesis, however, is probably an oversimplification, since additional sources of calcium such as voltage-dependent calcium channels (Bi and Poo, 1998; Magee et al, 1997) and other signalling mechanisms such as metabotropic glutamate receptors (Nevian et al, 2006) also contribute to STDP.

Downstream to the calcium influx is calmodulin, which may provide a watershed readout mechanism (DeMaria et al, 2001) to distinguish between LTP and LTD-promoting calcium signals (Sjöström and Nelson, 2002; Sjöström et al, 2008). Eventually, the enzyme calcium/calmodulin-dependent kinase II, or CaMKII, is affected by the calcium transient. This enzyme has been hypothesized to encode synaptic weight through gradations in the fraction of active subunits (Lisman, 1985; 1989; Lisman and Zhabotinsky 2001; Lisman et al, 2002).

Although there is a relative consensus regarding the mechanisms underlying potentiation, it is less clear-cut with depression. In one view, sublinear calcium summation triggers LTD (Koester and Sakmann, 1998), perhaps because postsynaptic NMDA receptors are suppressed during STDP timings inducing LTD (Froemke et al, 2005). It is, however, becoming increasingly clear that the presynaptic terminal too actively participates in the induction of STDP. In particular, presynaptically located NMDA receptors trigger timing-dependent LTD (Sjöström et al, 2003; Rodriguez-Moreno and Paulsen, 2008). The calcium hypothesis is thus clearly flawed and more work is required to elucidate the biophysical and biochemical mechanisms that underpin STDP.

1.6.3 Role of backpropagating action potentials

The NMDA-receptor-based spine coincidence detector described above requires that an action potential backpropagating from the initiation zone near the axon hillock into the dendritic tree makes it all the way to the synapse. If, however, a synapse is so distal from the soma that backpropagating action potentials fail and propagate passively, then it may not sufficiently depolarize NMDA receptors in the spine to allow opening and calcium influx (Golding and Spruston, 2002; Sjöström and Häusser, 2006). The prevalence of such failures of action potential backpropagation depend on the biophysical properties and morphology of the dendritic arbor, on sub and suprathreshold activity patterns, as well as on neuromodulatory state (Sjöström et al, 2008). For example, dendritic depolarization may boost otherwise failing backpropagating action potentials, thus promoting LTP of suitably timed synaptic inputs in the distal dendritic tree (Sjöström and Häusser, 2006). This is not to say that backpropagating action potentials are necessarily critical for NMDA-receptor-based LTP; local dendritic spikes

may entirely replace them, at least under some circumstances (Golding and Spruston, 2002). STDP, however, is by definition dependent on relatively global action potentials in the postsynaptic cell.

1.6.4 Voltage dependence and cooperativity in STDP

Classically, NMDA-receptor-based synaptic plasticity is closely connected to the degree of activation of the postsynaptic cell: Moderate depolarization only partially opens NMDA receptors, resulting in relatively low calcium levels and in LTD, whereas strong depolarization results in more massive calcium responses and in LTP (Artola et al 1990; also see above). In keeping with this, pairing presynaptic spikes with subthreshold postsynaptic depolarization results in timing-dependent depression (Sjöström et al, 2004; Markram et al, 1997).

Another hallmark feature of classical LTP is cooperativity; the notion that high-frequency stimulation of a weak pathway results in LTP only if in synchrony with a stronger pathway (McNaughton et al, 1978). Some have argued that this requirement for cooperativity among inputs to reach the threshold for LTP is a reflection of a need to reach the postsynaptic spike threshold. If this line of reasoning were true, then STDP should not exhibit a cooperativity requirement, since there is by definition always postsynaptic spiking. It turns out, however, that neocortical STDP does require a sufficient number of inputs to be co-activated in order to elicit LTP, even in the presence of postsynaptic spiking (Sjöström et al, 2001). It has been demonstrated that this cooperativity requirement in STDP arises due to a voltage dependence, so that large depolarizations (e.g. due to a large number of synchronous inputs) enable potentiation, whereas small ones fail to do so (Sjöström et al, 2001). This voltage dependence is at least in part due to the fact that action potentials backpropagate decrementally into the dendritic tree unless they are boosted by a relatively depolarized dendritic state (Sjöström and Häusser, 2006). In other words, in STDP, the backpropagating action potential may require help to make it to the synapse, especially for synapses far from the soma, otherwise it cannot sufficiently depolarize the spine coincidence detector to trigger potentiation.

Although it is relatively well-established that STDP is voltage dependent and that backpropagating action potentials are crucial, it is unclear to what extent other voltage dependent mechanisms contribute. For example, calcium influx directly mediated by voltage-dependent calcium channels may contribute. Another open question is if a form of STDP exists for local dendritic spikes, i.e. in the absence of postsynaptic spiking output (cf. Golding and Spruston, 2002).

1.6.5 Induction versus expression of Long-Term Potentiation

By and large, STDP refers to an experimental plasticity induction protocol. It may thus be tempting to conclude that the controversial question of how synaptic plasticity is expressed–in particular the so-called pre versus post debate in LTP (Malenka and Nicoll, 1999)–is not relevant to STDP models. Such a conclusion, however, may be hasty.

Although it has been disputed on good grounds (e.g. Bolshakov and Siegelbaum, 1994), the canonical view remains that LTP at Schaeffer collateral inputs to CA1 pyramidal cells is postsynaptically expressed (Malenka and Nicoll, 1999). In this view, potentiation is a simple synaptic gain control that underlies information storage in the brain.

With presynaptic LTP, however, the situation is considerably more complex, because presynaptic LTP does not only change the synaptic gain, it also affects infor-

mation transfer across the synapse. At excitatory connections of neocortical layer 5, LTP is apparently expressed through an upregulation of the probability of release, thus resulting in an increase of short-term depression, a concept known as Redistribution of Synaptic Efficacy (RSE; Markram and Tsodyks, 1996). Consistent with this finding, the induction of timing-dependent depression at layer-5 synapses in visual cortex results in a long-term down regulation of short-term depression through an apparent decrease of the probability of release (Sjöström et al., 2003), which is in effect anti-RSE.

Since short-term depression effectively differentiates rates on a given input with respect to time, this may lead to brief onset and offset bursts of activity in the postsynaptic cell as input rates change (Abbott et al 1997). A cell with short-term depressing inputs thus becomes an efficient coherence detector (Abbott et al 1997; Markram and Tsodyks, 1996). By extension, the presence of RSE at its synaptic inputs may thus make a neuron more sensitive to input coherence (Markram and Tsodyks, 1996), while anti-RSE (Sjöström et al, 2003) may do the opposite.

The up or down regulation of short-term depression in a network of connected neurons would presumably alter the statistics of spike timing dramatically. Given the acute sensitivity of STDP to spike timing, it thus follows that a network with STDP triggering RSE may result in complex loops between STDP and network activity. The ensuing dynamics of activity are likely to be quite different from those in a network where STDP does not trigger RSE. To our knowledge, this possibility and its functional consequences for network coding has not yet been explored, neither theoretically nor experimentally. In fact, most models introduce the synaptic weight as formal parameter that corresponds to the amplitude of the EPSP or the maximum conductance during synaptic transmission. However, if one combines an STDP model with a model of short-term plasticity with several parameters, the term 'synaptic weight' is not precise enough, since long-term plasticity may affect the parameters of short-term plasticity differentially. The basic question of what we mean by synaptic weight thus remains to be addressed properly.

1.6.6 Maintenance of Long-Term Potentiation

The focus of STDP as an experimental paradigm (and therefore of this article) is the induction of plasticity by suitable protocols. The question of how the changes induced by synaptic plasticity are maintained over a period of hours, weeks, or even years as expected for long-term memory is the topic of Maintenance of synaptic plasticity.

1.6.7 Influence of neuromodulators

STDP depends on the presence or absence of neuromodulators such as dopamine (Pawlak and Kerr, 2008, Zhang et al. 2009). These studies suggest that neuromodulators are more than simple switches that turn plasticity on or off. Neuromodulation cannot be considered as a simple multiplicative factor. Rather the presence of neuromodulators changes the temporal profile of STDP (Pawlak and Kerr, 2008, Zhang et al. 2009).

The modulation of STDP by a third factor such as dopamine has potentially interesting functional consequences that turn STDP from unsupervised learning into a reward-based learning paradigm (Izhikevich 2007, Florian 2007, Pfister 2006, Farries and Fairhall 2008, Legenstein et al. 2008).

1.6.8 Discrete or continuous synapses

Most models describe the synaptic weight as a continuous variable, although it is quite conceivable that weights are coded in discrete jumps. In fact, certain benefits would arise from such discrete synaptic weights. For example, bistability of individual synapses would help to assure the long-term stability of synapses over weeks or years in the presence of molecular turnover (Lisman 1985). While this is a strong argument in favor of discrete synapses it does not preclude that, on a shorter time scale, synaptic weights undergo continuous synaptic depression or facilitation which would be overlaid on the discrete long-term dynamics.

Since in the typical STDP experiment, results are averaged across several synapses, the question of whether single synapses respond to plasticity protocols with discrete or continuous changes cannot readily be answered. There are at least two studies suggesting that synaptic weights onto CA1 pyramidal neurons in the hippocampus are altered in discrete steps (Petersen et al, 1998; O'Connor et al, 2005). Other studies, however, appear to be in disagreement with this view. For example, recent glutamate uncaging experiments suggest that weights change continuously (Tanaka et al, 2008).

Finally, even if weights are discrete, it is difficult to provide conclusive experimental evidence showing stepwise changes in plasticity. The stochastic nature of neurotransmitter release, for example, hampers such experiments, by adding noise to the point that stepwise changes might be masked, although glutamate uncaging would help address this problem. Furthermore, the situation might be complex, and plasticity may be for example be discrete postsynaptically and continuous presynaptically. Last but not least, synapses are at different distances from the soma and their corresponding postsynaptic potentials are therefore subjected to different amounts of dendritic filtering as they propagate toward the soma. Since most connections in the brain are made up of more than one synaptic contact, which are made onto different dendritic compartments at different electrotonic distances from the soma, the net result is that any discrete steps that might exist would be exceedingly difficult to find conclusive evidence for experimentally. Even if weights were discrete, synaptic weight distributions would seem continuous and discrete plasticity would appear continuous.

Whether synapses typically are discrete or continuous thus remains an open but very intriguing question.

1.7 History of STDP

The first experiments with precisely timed pre- and postsynaptic spikes at a millisecond temporal resolution were performed by Markram et al. (1995,1997) soon followed by others (Bell et al., 1997, Bi and Poo, 1998, Debanne et al., 1998, Zhang et al., 1998). While the first publications on true STDP experiments came out in the mid-nineties temporal requirements for the coincidence of pre- and postsynaptic activity had already been investigated in 1983 in experiments by Levy and Stuart albeit with a lower temporal resolution using bursts of spikes rather than individual action potentials (Levy and Stuart, 1983). These early experiments can be understood as temporally asymmetric Hebbian learning under a rate coding hypothesis, but also as precursors of modern STDP experiments.

The first model using an STDP function with potentiation and depression at a millisecond resolution was published in 1996 (Gerstner et al. 1996). Modelers and theoreticians interested in Hebbian learning have been interested in temporally asymmetric forms of Hebbian learning in the context of sequence learning with pre-before-post LTP for spike patterns (Gerstner et al. 1993) and asymmetric LTP/LTD for behavioral se-

quences (Abbott and Blum, 1996) and even earlier, already in the mid 1980s, in the context of associative memories (Sompolinsky and Kanter, 1986, Herz et al., 1988, Kleinfeld and Sompolinky 1988). In standard attractor networks of memory, the learning rule includes terms of the Hebbian form $f_i(t)f_j(t)$ where $f_j(t)$ and $f_i(t)$ denote the firing rate of pre- and postsynaptic neurons in the rate pattern present at time step t. In order to learn sequences of patterns (or non-stationary attractors) the learning rule should contain terms of the form $f_i(t)f_j(t-1)$, i.e., a form of temporally asymmetric Hebbian learning that correlates the firing rate of the postsynaptic neuron i at time t with that of a presynaptic neuron j during the time step $t-1$ (Sompolinsky and Kanter, 1986, Herz et al., 1988, Kleinfeld and Sompolinky 1988). In 1993 Gerstner and van Hemmen started to translate ideas from sequence learning in discrete-time rate models to the case of spiking neurons in continuous time and formulated a learning rule where presynaptic spikes arrival a few milliseconds before postsynaptic firing leads to a potentiation of synapses (Gerstner et al., 1993). Depression of synapses was unspecific and not part of the spike-timing dependent learning rule. For purely theoretical reasons Gerstner and colleagues postulated in a paper submitted to Nature in 1995 that presynaptic spike arrival before postsynaptic firing should lead to potentiation whereas the reverse timing should lead to depression. Referees of that paper asked whether there was any experimental support for this speculation. In the mean time Markram et al. published an abstract in the Society of Neuroscience meeting of 1995, which was then cited by Gerstner et al. so as to convince the referees and the theory paper was published in Nature in 1996 (Gerstner et al., 1996). To our knowledge, this is the first paper that plotted synaptic plasticity as a function of the relative timing of individual pre- and postsynaptic action potentials. The work of Markram that came out in 1995 at the Society of Neuroscience meeting as an abstract and in 1997 as a full article (Markram et al. 1997) was in fact conducted in 1992-1993 while Henry Markram was a postdoctoral fellow in the laboratory of Bert Sakmann.

2.1 References

Abbott, L.F., and Nelson, S.B. (2000). Synaptic plasticity: taming the beast. Nature Neuroscience 3, 1178-1183.

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

Allen, C.B., Celikel, T., and Feldman, D.E. (2003). Long-term depression induced by sensory deprivation during cortical map plasticity in vivo. Nat Neurosci 6, 291-299.

Artola, A., Bröcher, S. and Singer, W. (1990). Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. Nature 347, 69-72.

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

Bi, G. Q. and Poo, M. M. (1998). Synaptic modifications in cultured Hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci, 18:10464-72.

Bi, G. and Poo, M. (2001). Synaptic modification of correlated activity: Hebb's postulate revisited. Ann. Rev. Neurosci., 24:139-166.

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

Caporale, N., and Dan, Y. (2008). Spike timing-dependent plasticity: a Hebbian learning rule. Annu Rev Neurosci 31, 25-46.

DeMaria, C.D., Soong, T.W., Alseikhan, B.A., Alvania, R.S. and Yue, D.T. (2001). Calmodulin bifurcates the local Ca2+ signal that modulates P/Q-type Ca2+ channels. Nature 411, 484-9.

Egger, V., Feldmeyer, D., and Sakmann, B. (1999). Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. Nat Neurosci 2, 1098-1105.

Farries, M.A., Fairhall A.L. (2007) Reinforcement learning with modulated spike timing-dependent synaptic plasticity. J Neurophysiol 98: 3648-3665.

Feldman, D.E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. Neuron 27, 45-56.

Fino, E., Glowinski, J., and Venance, L. (2005). Bidirectional activity-dependent plasticity at corticostriatal synapses. J Neurosci 25, 11279-11287.

Florian, R.V. (2007) Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. Neural Computation 19: 1468-1502.

Froemke, R., Poo, M., and Dan, Y. (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. Nature, 434:221-5.

Frey U, Morris R (1997) Synaptic tagging and long-term potentiation. Nature 385:533 536.

Gamble, E. and Koch, C. (1987). The dynamics of free calcium in dendritic spines in response to repetitive synaptic input. Science, 236:1311-1315.

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.

Gerstner, W., Kempter R., van Hemmen J.L., and Wagner H. (1996). A neuronal learning rule for sub-millisecond temporal coding. Nature, 386:76-78.

Golding, N.L., Staff, N.P. and Spruston, N. (2002). Dendritic spikes as a mechanism for cooperative long-term potentiation. Nature 418, 326-31.

Graupner, M. and Brunel, N. (2007). STDP in a Bistable Synapse Model Based on CaMKII and Associated Signaling Pathways. PLoS Comput Biol 3(11): e221. doi:10.1371/journal.pcbi.0030221

Gutig, R.,Aharonov, R., Rotter, S. and Sompolinsky, H. (2003). Learning Input Correlations through Nonlinear Temporally Asymmetric Hebbian Plasticity. J. Neurosci. 23:3697-3714

Guyonneau, R., VanRullen, R., and Thorpe, S.J. (2005) Neurons tune to the earliest spikes through STDP. Neural Computation, 17:859-879, 2005.

Hebb, D. O. (1949). The Organization of Behavior; a neuropsychological theory. Wiley, New York.

Herz, A. V. M., Sulzer, B., Kuhn, R., and van Hemmen, J. L. (1988). The Hebb rule: Representation of static and dynamic objects in neural nets. Europhys. Lett., 7:663-669.

Holmes, W. and Levy, W. (1990). Insights into associative long- term potentiation from computational models of NMDA receptor-mediated calcium influx and intracellular calcium concentration changes. J Neurophysiol, 63:1148-68.

Holmgren, C.D., and Zilberter, Y. (2001). Coincident spiking activity induces long-term changes in inhibition of neocortical pyramidal cells. J Neurosci 21, 8270-8277.

Izhikevich, E. and Desai, N. (2003). Relating STDP to BCM. Neural Comput, 15:1511-23.

Izhikevich, E. (2007) Solving the distal reward problem through linkage of STDP and dopamine signaling. Cerebral Cortex 17: 2443-2452.

Jacob, V., Brasier, D.J., Erchova, I., Feldman, D., and Shulz, D.E. (2007). Spike timing-dependent synaptic depression in the in vivo barrel cortex of the rat. J Neurosci 27, 1271-1284.

Karmarkar, U. and Buonomano, D. V. (2002). A model of spike-timing dependent plasticity: one or two coincidence detectors? J Neurophysiol, 88:507-13.

Katz, L.C., and Shatz, C.J. (1996). Synaptic activity and the construction of cortical circuits. Science 274, 1133-1138.

Kempter, R., Gerstner, W., and van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. Phys. Rev. E, 59:4498-4514.

Kempter, R., Gerstner, W., and van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. Neural Comput. 13:2709-2741

Kleinfeld, D. and Sompolinksy, K. (1988). Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. Biophysical J. 54:1039-1051

Koester, H.J. and Sakmann, B. (1998). Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials. Proceedings of the National Academy of Sciences of the United States of America 95, 9596-601.

Legenstein, R., Pecevski, D., and Maass W. (2008) A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. PLoS Computational Biology 4(10): e1000180.

Letzkus, J.J., Kampa, B.M., and Stuart, G.J. (2006). Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. J Neurosci 26, 10420-10429.

Lisman, J.E. (1985). A mechanism for memory storage insensitive to molecular turnover: a bistable autophosphorylating kinase. Proc Natl Acad Sci U S A 82, 3055-7.

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

Lisman, J. and Zhabotinsky, A. (2001). A model of synaptic memory: A CaMKII/PP1 switch that potentiates transmission by organizing an AMPA receptor anchoring assembly, Neuron 31:191-201

Lisman, J., Schulman, H., Cline, H. (2002) The molecular basis of CamKII function in synaptic and behavioural memory. Nat Rev Neurosci 3:175190.

Lisman, J. (2003) Long-term potentiation: outstanding questions and attempted synthesis. Phil Trans R Soc Lond B: Biological Sciences 358:829 842.

Magee, J.C. and Johnston, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. Science 275, 209-13.

Markram, H. and Sakmann, B. (1995). Action potentials propagating back into dendrites triggers changes in efficacy. Soc. Neurosci. Abs. 21,

Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science, 275:213-5.

McNaughton, B.L., Douglas, R.M. and Goddard, G.V. (1978). Synaptic enhancement in fascia dentata: cooperativity among coactive afferents. Brain Research 157, 277-93.

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

Meliza, C.D., and Dan, Y. (2006). Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. Neuron 49, 183-189.

Nevian, T. and Sakmann, B. (2006). Spine Ca2+ signaling in spike-timing-dependent plasticity. J Neurosci 26, 11001-13.

Nishiyama, M., Hong, K., Mikoshiba, K., Poo, M.M., and Kato, K. (2000). Calcium stores regulate the polarity and input specificity of synaptic modification. Nature 408, 584-588.

OConnor, D., Wittenberg, G., and Wang, S.H. (2005) Graded bidirectional synaptic plasticity is composed of switch-like unitary events. Proc Natl Acad Sci USA 102:96799684.

Pawlak, V. and Kerr, J.N.D. (2008) Dopamine receptor activation is required for corticostriatal spike-timing-dependent plasticity. J Neurosci 28: 24352446.

Petersen, C., Malenka, R., Nicoll, R., and Hopfield, J. (1998) All-or-none potentiation of ca3-ca1 synapses. Proc Natl Acad Sci USA 95:47324737.

Pfister, J.P., Toyoizumi, T., Barber, D., and Gerstner, W. (2006) Optimal spike-timing dependent plasticity for precise action potential firing in supervised learning. Neural Computation 18: 1309-1339.

Pfister, J.P. and Gerstner, W. (2006) Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. J. Neurosci. 26:9673-9682

Roberts, P.D. and Bell, C.C. (2000). Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. J. Comput. Neurosci., 9:67-83.

Rodriguez-Moreno, A. and Paulsen, O. (2008). Spike timing-dependent long-term depression requires presynaptic NMDA receptors. Nat Neurosci 11, 744-5.

Rubin, J., Gerkin, R., Bi, G., and Chow, C. (2005). Calcium time course as a signal for spike-timing-dependent plasticity. J Neurophysiol, 93:2600-13.

Schiller, J., Schiller, Y. and Clapham, D.E. (1998). NMDA receptors amplify calcium influx into dendritic spines during associative pre- and postsynaptic activation. Nature Neuroscience 1, 114-8.

Senn, W., Tsodyks, M., and Markram, H. (1997). An algorithm for synaptic modification based on exact timing of pre- and post-synaptic action potentials. In Gerstner, W., Germond, A., Hasler, M., and Nicoud, J.-D., editors, Artificial Neural Networks - ICANN '97, pages 121-126. Springer.

Senn, W., Markram, H., and Tsodyks, M. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and post-synaptic spike timing. Neural Comput, 13:35-67.

Shouval, H. Z., Bear, M. F., and Cooper, L. N. (2002). A unified theory of NMDA receptor-dependent bidirectional synaptic plasticity. Proc. Natl. Acad. Sci. USA, 99:10831-6.

Sjöström, P.J., Turrigiano, G.G., and Nelson, S.B. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 32, 1149-1164.

Sjöström, P.J. and Nelson, S.B. (2002). Spike timing, calcium signals and synaptic plasticity. Curr Opin Neurobiol 12, 305-14.

Sjöström, P.J., Turrigiano, G.G., and Nelson, S.B. (2003). Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. Neuron 39, 641-654.

Sjöström, P.J., Turrigiano, G.G., and Nelson, S.B. (2004). Endocannabinoid-dependent neocortical layer-5 LTD in the absence of postsynaptic spiking. J Neurophysiol 92, 3338-3343.

Sjöström, P.J., and Häusser, M. (2006). A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. Neuron 51, 227-238.

Sjöström, P.J., Rancz, E.A., Roth, A., and Häusser, M. (2008). Dendritic Excitability and Synaptic Plasticity. Physiological Reviews 88, 769-840.

Sompolinsky, H. and Kanter, I. (1986). Temporal association in asymmetric neural networks. Phys. Rev. Lett., 57:2861-2864.

Song, S., Miller, K.D., and Abbott, L.F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat Neurosci 3, 919-926.

Song, S., and Abbott, L.F. (2001). Cortical development and remapping through spike timing-dependent plasticity. Neuron 32, 339-350.

Stent, G.S. (1973). A physiological mechanism for Hebb's postulate of learning. PNAS 70, 997-1001.

Tanaka, J., Horiike, Y., Matsuzaki, M., Miyazaki, T., Ellis-Davies, G.C. and Kasai, H. (2008). Protein synthesis and neurotrophin-dependent structural plasticity of single dendritic spines. Science 319, 1683-7.

Yuste, R. and Denk, W. (1995). Dendritic spines as basic functional units of neuronal integration. Nature 375, 682-4.

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

Zhang, J.C., Lau P.M., Bi G.Q. (2009) Gain in sensitivity and loss in temporal contrast of stdp by dopaminergic modulation at hippocampal synapses. Proc Natl Acad Sci U S A 106: 1302813033.

2.2 Recommended Reading

Bi, G. and Poo, M. (2001). Synaptic modification of correlated activity: Hebb's postulate revisited. Ann. Rev. Neurosci., 24:139-166.

Gerstner, W. and Kistler, W. (2002). Spiking Neuron Models, Cambridge University Press, Chapters 10-12

Dan, Y. and Poo, M.-m. (2004) Spike Timing-Dependent Plasticity of Neural Circuits, Neuron, Vol. 44:2330, 2004

Sjöström, P.J., Rancz, E.A., Roth, A., and Häusser, M. (2008). Dendritic Excitability and Synaptic Plasticity. Physiological Reviews 88, 769-840.