REVIEW

Biological
Cybernetics

Phenomenological models of synaptic plasticity based on spike timing

Abigail Morrison · Markus Diesmann ·
Wulfram Gerstner

Abstract Synaptic plasticity is considered to be the biological substrate of learning and memory. In this document we review phenomenological models of short-term and long-term synaptic plasticity, in particular spike-timing dependent plasticity (STDP). The aim of the document is to provide a framework for classifying and evaluating different models of plasticity. We focus on phenomenological synaptic models that are compatible with integrate-and-fire type neuron models where each neuron is described by a small number of variables. This implies that synaptic update rules for short-term or long-term plasticity can only depend on spike timing and, potentially, on membrane potential, as well as on the value of the synaptic weight, or on low-pass filtered (temporally averaged) versions of the above variables. We examine the ability of the models to account for experimental data and to fulfill expectations derived from theoretical considerations. We further discuss their relations to teacher-based rules (supervised learning) and reward-based rules (reinforcement learning). All models discussed in this paper are suitable for large-scale network simulations.

Keywords Spike-timing dependent plasticity · Short term plasticity · Modeling · Simulation · Learning

A. Morrison · M. Diesmann
Computational Neuroscience Group,
RIKEN Brain Science Institute, Wako City, Japan

M. Diesmann
Bernstein Center for Computational Neuroscience,
Albert-Ludwigs-University, Freiburg, Germany

W. Gerstner (✉)
Laboratory of Computational Neuroscience, LCN,
Brain Mind Institute and School of Computer and Communication
Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL),
Station 15, 1015 Lausanne, Switzerland
e-mail: wulfram.gerstner@epfl.ch

1 Introduction

Synaptic changes are thought to be involved in learning, memory, and cortical plasticity, but the exact relation between microscopic synaptic properties and macroscopic functional consequences remains highly controversial. In experimental preparations, synaptic changes can be induced by specific stimulation conditions defined through presynaptic firing rates (Bliss and Lomo 1973; Dudek and Bear 1992), postsynaptic membrane potential (Kelso et al. 1986; Artola et al. 1990), calcium entry (Lisman 1989; Malenka et al. 1988), or spike timing (Markram et al. 1997; Bi and Poo 2001).

Whereas detailed biophysical models are crucial to understand the biological mechanisms underlying synaptic plasticity, phenomenological models which describe the synaptic changes without reference to mechanism are generally more tractable and less computationally expensive. Consequently, phenomenological models are of great use in analytical and simulation studies. In this manuscript, we will examine a number of phenomenological models with respect to their compatibility with both experimental and theoretical results. In all cases, we consider a synapse from a presynaptic neuron j to a postsynaptic neuron i. The strength of a connection from j to i is characterized by a weight w_{ij} that quantifies the amplitude of the postsynaptic response, typically measured as the height of the postsynaptic potential or the slope of the postsynaptic current at onset. The conditions for synaptic changes as well as their directions and magnitudes can be formulated as 'synaptic update rules' or 'learning rules'. Such rules can be developed from purely theoretical considerations, or to account for macroscopic phenomena such as the development of receptive fields, or based on findings from electrophysiological experiments manipulating firing rate or voltage. In this manuscript, however, we restrict our scope to rules which have been developed to account for the results

of experiments in which synaptic plasticity was observed as a result of pre- and postsynaptic spikes (for more general reviews, see Dayan and Abbott 2001; Gerstner and Kistler 2002; Cooper et al. 2004).

For the classification of the synaptic plasticity rules, it is important to specify the time necessary to *induce* such a change as well as the time scale of *persistence* of the change. For both short-term and long-term plasticity, changes can be induced in about 1 s or less. In short-term plasticity (see Sect. 3), a sequence of eight presynaptic spikes at 20 Hz evokes successively smaller (depression) or successively larger (facilitation) responses in the postsynaptic cell. The characteristic feature of short-term plasticity is that this change does not persist for more than a few hundred milliseconds: the amplitude of the postsynaptic response recovers to close-to-normal values within less than a second (Markram et al. 1998; Thomson et al. 1993).

In contrast to short-term plasticity, long-term potentiation and depression (LTP and LTD) refer to persistent changes of synaptic responses (see Sect. 4). Note that the time necessary for induction can still be relatively brief. For example, in spike-timing-dependent plasticity (Bi and Poo 2001; Sjostrom et al. 2001), a change of the synapse can be induced by 60 pairs of pre- and postsynaptic spikes with a repetition frequency of 20 Hz; hence stimulation is over after 3 s. However, this change can persist for more than one hour. The final stabilization of, say, a potentiated synapse occurs only thereafter, called the late phase of LTP (Frey and Morris 1997). An additional aspect is that neurons in the brain must remain within a sustainable activity regime, despite the changes induced by LTP and LTD. This is achieved by homeostatic plasticity, an up- or down-regulation of all synapses converging onto the same postsynaptic neuron which occurs on the time scale of minutes to hours (Turrigiano and Nelson 2004).

The phenomenological models discussed in this manuscript can be classified from a theoretical point of view as unsupervised learning rules. There is no notion of a task to be solved, nor is there any notion of the change being 'good' or 'bad' for the survival of the animal; learning consists simply of an adaptation of the synapse to the statistics of the activity of pre- and postsynaptic neurons. This is to be contrasted with reward-based learning, also called reinforcement learning (Sutton and Barto 1998). In reward-based learning the direction and amount of change depends on the presence or absence of a success signal, that may reflect the current reward or the difference between expected and received reward (Schultz et al. 1997). Reward-based learning rules are distinct from supervised learning since the success signal is considered as a global and unspecific feedback signal, that often comes with a delay, whereas in supervised learning the feedback is much more specific. In the theoretical literature, there exists a large variety of update rules that can be classified as supervised, unsupervised or reward based learning rules.

In this paper, we start with a review of some basic experimental facts that could be relevant for modeling, followed by a list of theoretical concepts arising from fundamental notions of learning and memory formation (Sect. 2). We then review models of short-term plasticity in Sect. 3 and models of long-term potentiation/depression (LTP/LTD), in particular the spike-timing dependent form, in Sect. 4. Throughout the review we discuss spike-based plasticity rules from a computational perspective, giving implementations that are appropriate for analytical and simulation approaches. In the final sections we briefly mention reward driven learning rules for spiking neurons (Sect. 5) and provide an outlook toward current challenges for modeling. The relevance of molecular mechanisms and signaling chains (Lisman 1989; Malenka et al. 1988) for models of synaptic plasticity (Lisman and Zhabotinsky 2001; Shouval et al. 2002; Rubin et al. 2005; Badoual et al. 2006; Graupner and Brunel 2007; Zou and Destexhe 2007), as well as the importance of the postsynaptic voltage (Kelso et al. 1986; Artola et al. 1990; Sjostrom et al. 2001), is acknowledged but not further explored.

2 Perspectives on plasticity

Over the last 30 years, a large body of experimental results on synaptic plasticity has been accumulated. The most important discoveries are summarized in Sect. 2.1. Simultaneously, theoreticians have investigated the role of synaptic plasticity in long-term memory, developmental learning and task-specific learning. The most important concepts arising from this research are described in Sect. 2.2. Many of the plasticity models employed in the theoretical approach were inspired by Hebb's (1949) postulate that describes how synaptic connections should be modified:

> When an axon of cell A is near enough to excite cell B or repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.

In classical Hebbian models, this famous postulate is often rephrased in the sense that modifications in the synaptic transmission efficacy are driven by correlations in the firing activity of pre- and postsynaptic neurons. Even though the idea of learning through correlations dates further back in the past (James 1890), correlation-based learning is now generally called Hebbian learning. Most classic theoretical studies represented the activity of pre- and postsynaptic neurons in terms of rates, expressed as continuous functions. This has led to a sound understanding of rate-based Hebbian learning. However, rate-based Hebbian learning neglects the fine temporal structure between pre- and postsynaptic spikes.

Spike-based learning models for temporally structured input need to take this timing information into account (e.g. Gerstner et al. 1993) which leads to models of spike-timing dependent plasticity (STDP) (Gerstner et al. 1996; Kempter et al. 1999; Roberts 1999; Abbott and Nelson 2000) that can be seen as a spike-based generalization of Hebbian learning. The first experimental reports showing both long-term potentiation and depression induced by causal and acausal spike timings on a time scale of 10 ms were published by Markram and Sakmann (1995) and Markram et al. (1997), slightly after the theoretical work, however potentiation induced by the pairing of EPSPs with postsynaptic depolarization on a time scale of 100 ms was demonstrated considerably earlier (Gustafsson et al. 1987). Timing in rate-based Hebbian learning (although not spike-based) can be traced even further back in the past (Levy and Steward 1983). From a conceptual point of view, all spike-based and rate-based Hebbian learning rules share the feature that only variables that are *locally* available at the synapse can be used to change the synaptic weight. These local elements that can be used to construct such rules are listed in Sect. 2.3.

2.1 Experimental results

The most important results from experiments on synaptic plasticity with respect to the modeling of synaptic plasticity are as follows:

(i) Short-term plasticity depends on the sequence of presynaptic spikes on a time scale of tens of milliseconds (Markram et al. 1998; Thomson et al. 1993).

(ii) Long-term plasticity is sensitive to the presynaptic firing rate over a time scale of tens or hundreds of seconds. For example 900 presynaptic stimulation pulses at 1 Hz (i.e. 15 min of total stimulation time) yield a persistent depression of the synapses, whereas the same number of pulses at 50 Hz yields potentiation (Dudek and Bear 1992).

(iii) Long-term plasticity depends on the exact timing of the pre- and postsynaptic spikes on the time scale of milliseconds (Markram et al. 1997; Bi and Poo 2001). For example LTP is induced if a presynaptic spike precedes the postsynaptic one by 10 ms, whereas LTD occurs if the order of spikes is reversed. In this context it is important to realize that most experiments are done with repetitions of 50–60 pairs of spikes whereas a single pair has no effect.

(iv) STDP depends on the repetition frequency of the pre-post spike-pairings. In fact, 60 pairings pre-before-post at low frequency have no effect, whereas the same number of pairs at a repetition frequency of 20 Hz gives strong potentiation (Sjostrom et al. 2001).

(v) Plasticity depends on the postsynaptic potential (Kelso et al. 1986; Artola et al. 1990). If the postsynaptic neuron is clamped to a voltage slightly above rest during presynaptic spike arrival, the synapses are depressed, while at higher depolarization the same stimulation leads to LTP (Artola et al. 1990; Ngezahayo et al. 2000).

(vi) On a slow time scale of hours, homeostatic changes of synapses may occur in form of rescaling of synaptic response amplitudes (Turrigiano et al. 1994). These changes can be useful to stabilize neuronal firing rates.

(vii) Also on the time scale of hours, early phase LTP is consolidated into late phase LTP. During the consolidation phase heterosynaptic interactions may take place, probably as a result of synaptic tagging and competition for scarce protein supply (Frey and Morris 1997). Consolidation is thought to lead to long-term stability of the synapses.

(viii) Distributions of synaptic strength (e.g., the EPSP amplitudes) in data collected across several pairs of neurons are reported to be unimodal (Sjostrom et al. 2001). At a first glance, this seems to be at odds with experimental data suggesting that single synaptic contacts are in fact binary (Petersen et al. 1998; O'Connor et al. 2005).

(ix) Synapses do not form a homogeneous group, but different types of synapse have different plasticity properties (Abbott and Nelson 2000; Thomson and Lamy 2007). In fact, the same presynaptic neuron makes connections to different types of target neurons with different plasticity properties for short-term (Markram et al. 1998) and long-term plasticity (Lu et al. 2007).

Many other experimental features could be added to this list, e.g., the role of intracellular calcium, of NMDA receptors, etc., but we will not do so; see Bliss and Collingridge (1993) and Malenka and Nicoll (1993) for reviews. We emphasize that, given the heterogeneity of synapses between different brain areas (plasticity has mainly been studied in visual or somatosensory cortex and hippocampus) and between different neuron and synapse types, we cannot expect that a single theoretical model can account for all experimental facts. In the next section, we will instead consider which theoretical principles could guide our search for suitable plasticity rules.

2.2 Theoretical concepts

Synaptic plasticity is held to be the basis for long-term memory, developmental learning, and task-specific learning. From a theoretical point of view, synaptic learning rules should therefore provide:

(i)   sensitivity to correlations between pre- and postsynaptic neurons (Hebb 1949) in order to respond to correlations in the input (Oja 1982). This is the essence of all unsupervised learning rules

(ii) a mechanism for the development of input selectivity such as receptive fields (Bienenstock et al. 1982; Miller et al. 1989), in the presence of strong input features. This is the essence of developmental learning

(iii) a high degree of stability (Fusi et al. 2005) in the synaptic memories whilst remaining plastic (Grossberg 1987). This is the essence of memory formation and memory maintenance

(iv) the ability to take into account the quality of task performance mediated by a global success signal (e.g. neuromodulators, Schultz et al. 1997). This is the essence of reinforcement learning (Sutton and Barto 1998).

These items are not necessarily exclusive, and the relative importance of a given aspect may vary from one subsystem to the next; for example, synaptic memory maintenance might be more important for a long-term memory system than for primary sensory cortices. There is so far no rule which exhibits all of the above properties; moreover, theoretical models which reproduce some aspects of experimental findings are generally incompatible with other findings. For example, traditional learning rules that have been proposed as an explanation of receptive field development (Bienenstock et al. 1982; Miller et al. 1989), exhibit a spontaneous separation of synaptic weights into two groups, even if the input shows no or only weak correlations. This is difficult to reconcile with experimental results in visual cortex of young rats where a unimodal distribution was found (Sjostrom et al. 2001). Moreover model neurons that specialize early in development on one subset of features cannot readily re-adapt later on. On the other hand, learning rules that do produce a unimodal distribution of synaptic weights (van Rossum et al. 2000; Rubin et al. 2001; Gütig et al. 2003; Morrison et al. 2007) do not lead to long-term stability of synaptic changes, as the trajectories of individual synaptic weights perform random walks. Hence it appears that long-term stability of memory requires a multimodal synapse distribution (Toyoizumi et al. 2007; Billings and van Rossum 2008) or additional mechanisms to stabilize the synaptic weights contributing to the retention of a memory item.

2.3 Locally computable measures

In a typical spiking network model, a neuron is characterized by its voltage (subthreshold) and its firing times (spikes/superthreshold). Mesoscopic measures such as population activity are not accessible for the individual synapses. As a
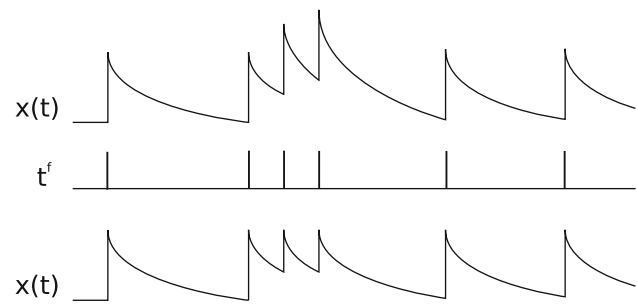
Fig. 1 Implementation of plasticity by local variables: each spike contributes to a trace $x(t)$. The update of the trace is either by a fixed value (*top*) or to a fixed value (*bottom*)

consequence, spike-based plasticity models may be constructed from a combination of the following terms:

(i)   spontaneous growth or decay (a non-Hebbian zero-order term)—this could be a small effect that leads to slow 'homeostatic' scaling of weights in the absence of any activity

(ii) effects caused by postsynaptic spikes alone independent of presynaptic spike arrival (a non-Hebbian first-order term). This could be an additional realization of homeostasis: if the postsynaptic neuron spikes at a high rate over hours, all synapses are down-regulated

(iii) effects caused by presynaptic spikes, independent of postsynaptic variables (another non-Hebbian first-order term). This is typically the case for short-term synaptic plasticity

(iv) effects caused by presynaptic spikes in conjunction with postsynaptic spikes (STDP) or in conjunction with postsynaptic depolarization (Hebbian terms)

(v)   all of the above effects may depend on the current value of the synaptic weight. For example, close to a maximum weight synaptic changes could become smaller.

We note that the changes induced by pre- or postsynaptic spikes need not necessarily immediately affect the synaptic weight. Alternatively, they may lead to an update of an internal hidden variable which evolves with some time constant τ. Hence, the hidden variable implements a low-pass filter. For example, let us denote by $\delta(t - t^f)$ a spike of a neuron occurring at time t^f. Then an internal variable x can be defined with dynamics:

$$\frac{dx}{dt} = -\frac{x}{\tau} + \sum_{t^f} A\,\delta\left(t - t^f\right), \qquad (1)$$

such that it is updated with each spike by an amount A and decays between spikes with a time constant τ (see Fig. 1, top). If the time constant is sufficiently long and $A = 1/\tau$, the hidden variable gives an online estimate of the mean firing rate in the spike train. Other variations in the formulation of

such a 'trace' left by a spike are possible that do not scale linearly with the rate. First, instead of updating by the same amount each time, we may induce saturation,

$$\frac{dx}{dt} = -\frac{x}{\tau} + \sum_{t^f} A \, (1 - x_-) \, \delta\left(t - t^f\right). \qquad (2)$$

For $A < 1$ the amount of increase gets smaller as the variable x before the update (denoted by x_-) approaches its maximal value of 1. Hence the variable x stays bounded in the range $0 \leq x \leq 1$. An extreme case of saturation is given by $A = 1$, in which case the reset is always to the value of 1, regardless of the value of x just before. In this case, the value of the trace x depends only on the time since the most recent spike (see Fig. 1, bottom). We will see in the following sections that the idea of traces left by pre- or postsynaptic spikes plays a fundamental role in algorithmic formulations of short-term and long-term plasticity. For example, in the case of Hebbian long-term potentiation, traces left by presynaptic spikes need to be combined with postsynaptic spikes, whereas short-term plasticity can be seen as induced by traces of presynaptic spikes, independent of the state of the postsynaptic neuron.

In principle, voltage dependence could be treated in a similar fashion, see, e.g., Brader et al. (2007), but we will focus in the following on learning rules for short-term and long-term plasticity that use spike timing as the relevant variable for inducing postsynaptic changes.

3 Short-term plasticity

Biological synapses have an inherent dynamics, which controls how the pattern of amplitudes of postsynaptic responses depends on the temporal pattern of the incoming spike train. Notably, each successive spike can evoke a response in the postsynaptic neuron that is smaller (depression) or larger (facilitation) than the previous one. Its time scale ranges from 100 ms to about a second. Fast synaptic dynamics is firmly established in biological literature (Markram et al. 1998; Gupta et al. 2000), and well-accepted models exist for it (Abbott et al. 1997; Tsodyks et al. 1998). Neurotransmitter is released in quanta of fixed size, each evoking a contribution to the postsynaptic potential of fixed amplitude; this is known as the quantal synaptic potential (Kandel et al. 2000). The release of an individual quantum is known to be stochastic, but the details of the mechanism underlying this stochasticity remain unclear. However, the following two phenomenological models describe the average response and are therefore entirely deterministic. Both models use the idea of a 'trace' left by presynaptic spikes (see previous section), but in slightly different formulations.

3.1 Markram–Tsodyks Model

One well-established phenomenological model for fast synaptic dynamics was originally formulated for depression only in Tsodyks and Markram (1997) and later extended to facilitating dynamics in Markram et al. (1998). Here, we discuss the formulation of the model presented in Tsodyks et al. (2000).

If neuron i receives a synapse from neuron j (see Fig. 2b), the synaptic current (or conductance) in neuron i is $w_{ij} y_{ij}(t)$, where w_{ij} is the absolute strength and $y_{ij}(t)$ is a scaling factor that describes the momentary input to neuron i. Dropping the indices for the rest of this discussion, y evolves according to:

$$\frac{dx}{dt} = \frac{z}{\tau_{rec}} - u_+ x_- \delta\left(t - t_j^f\right) \qquad (3)$$

$$\frac{dy}{dt} = -\frac{y}{\tau_I} + u_+ x_- \delta\left(t - t_j^f\right) \qquad (4)$$

$$\frac{dz}{dt} = \frac{y}{\tau_I} - \frac{z}{\tau_{rec}} \qquad (5)$$

where x, y and z are the fractions of synaptic resources in the recovered, active, and inactive states respectively, t_j^f gives the timing of presynaptic spikes, τ_I is the decay constant of PSCs and τ_{rec} is the recovery time from synaptic depression. These equations describe the use of synaptic resources by each presynaptic spike—a fraction u_+ of the available resources x is used by each presynaptic spike. The variable u_+ therefore describes the effective use of the synaptic resources of the synapses, which is analogous to the probability of release in the model described in (Markram et al. 1998). The notation x_- in the update equations (3) and (4) is intended to remind the reader that the value of x just *before* the update is used. In facilitating synapses, u_+ is not a fixed parameter, but derived from a variable u which is increased with each presynaptic spike and returns to baseline with a time constant τ_{fac}:

$$\frac{du}{dt} = -\frac{u}{\tau_{fac}} + U(1 - u_-) \delta\left(t - t_j^f\right) \qquad (6)$$

where the parameter U determines the increase in u with each spike. We note that the update is equivalent to the saturated trace (2). The notation u_- indicates that the value of u is taken just before the update caused by presynaptic spike arrival. However, in (3) and (4) we use the value u_+ just *after* the update of the variable u. If $\tau_{fac} \to 0$, facilitation is not exhibited, and u_+ is identical to U after each spike, as is the case with depressing synapses between excitatory pyramidal neurons (Tsodyks and Markram 1997). The model described by Eqs. 3–6 gives a very good fit to experimental results: compare Fig. 2a and c. However, it should be noted that the values for the model parameters, including the time constants, are quite heterogeneous, even within a single
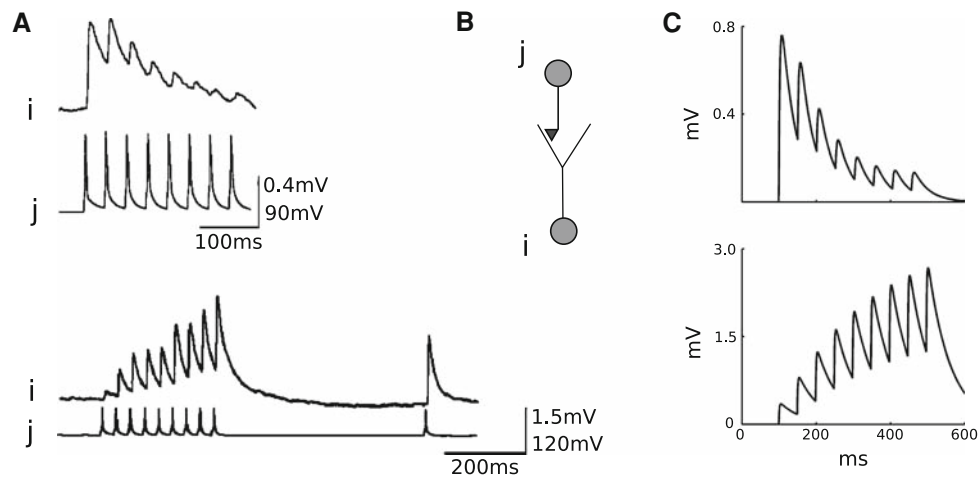
Fig. 2 Short-term plasticity in experiment and simulation. **a** Experimental results from rat cortex in slice. The average amplitude of the evoked postsynaptic potential in neuron i varies with each successive spike of the presynaptic neuron j. *Top panel*: depression; bottom panel: facilitation. **b** Schematic representation of the neuron configuration. **c** Simulation results. The membrane potential of the postsynaptic neuron i is shown for repeated stimulations mediated by a synapse implementing the Markram–Tsodyks model. Experimental results adapted from Markram et al. (1998), simulations performed with NEST (Gewaltig and Diesmann 2007)

neural circuit (Markram et al. 1998). The biophysical causes of the heterogeneity are still largely unclear.

We note that not only the usage variable u, but also the variable y in (4) is essentially a 'trace' very similar to the one defined in the preceding section. To see this we eliminate the variable x which is possible since the total amount of synaptic resources is fixed ($x + y + z = 1$). Hence (4) becomes:

$$\frac{dy}{dt} = -\frac{y}{\tau_I} + u_+(1 - y_- - z_-)\,\delta\left(t - t_j^f\right),$$

which is a modification of the saturated trace in (2), the difference stemming from the fact that an additional 'inactive' state has been introduced. Let us now suppose that the life-time of the 'inactive' state is short, i.e. $\tau_{rec} \ll \tau_I$. Then z decays rapidly back to zero and the above equation becomes the standard saturated trace. Since $x = 1 - y$ and $dx/dt = -dy/dt$, the available synaptic resources have the dynamics:

$$\frac{dx}{dt} = \frac{1 - x}{\tau_I} - u_+ x_-\,\delta\left(t - t_j^f\right), \tag{7}$$

which implies that the variable x is reduced at each presynaptic spike and, in the absence of spikes, approaches an asymptotic value of unity with time constant τ_I.

The model defined in (3)–(6) can be solved using the technique of exact integration (Rotter and Diesmann 1999) by exploiting the following observations. The system of differential equations (3)–(6) is essentially linear, because all products of state variables are multiplied with delta functions. Therefore, between each presynaptic spike the system can be integrated linearly, and on the occurrence of each spike

the system is reset to a new initial condition. Moreover, the amount of synaptic resources is conserved - note that the right-hand sides of (3)–(5) add up to 0—thus z can be eliminated from the system. Let the state of the synapse be given by:

$$s = \begin{pmatrix} u \\ x \\ y \end{pmatrix},$$

with the dynamics of u given by (6), that of y by (4) and that of x given by $dx/dt = (1 - x - y)/\tau_{rec} - u_+ x_-\,\delta\left(t - t_j^f\right)$. Between two successive presynaptic firing times t' and t'' the state of the synapse evolves linearly. At t'', the state of the synapse without the effects of the new spike can be calculated as:

$$\widetilde{s}\left(t''\right) = P_{\Delta t}\begin{pmatrix} s\left(t'\right) \\ 1 \end{pmatrix},$$

where $\Delta t = t'' - t'$ is the time difference of the two spikes and $(s(t'), 1)^T$ is a four-dimensional vector. The closed form expression of the propagator matrix is:

$$P_{\Delta t} = \begin{pmatrix} e^{-\frac{\Delta t}{\tau_{fac}}} & 0 & 0 & 0 \\ 0 & e^{-\frac{\Delta t}{\tau_{rec}}} & \frac{\tau_I(e^{-\frac{\Delta t}{\tau_{rec}}} - e^{-\frac{\Delta t}{\tau_I}})}{\tau_I - \tau_{rec}} & 1 - e^{-\frac{\Delta t}{\tau_{rec}}} \\ 0 & 0 & e^{-\frac{\Delta t}{\tau_I}} & 0 \end{pmatrix}.$$

The state of the synapse at t'' is the sum of the linear evolution since t' and the non-linear modification of the state due to the new spike:

$$s\left(t''\right) = \widetilde{s}\left(t''\right) + s_0,$$

where the initial conditions are given by:

$$s_0 = \begin{pmatrix} U\left(1 - \widetilde{u}\right) \\ -\widetilde{x}\left[\widetilde{u} + U\left(1 - \widetilde{u}\right)\right] \\ \widetilde{x}\left[\widetilde{u} + U\left(1 - \widetilde{u}\right)\right] \end{pmatrix}.$$

Note that the updated value of u is used to update the variables x and y. This reflects the assumption that the effectivity of resource use is determined not just by the history of the synapse but also by the arrival of the new presynaptic spike, thus ensuring a non-zero response to the first spike (Tsodyks et al. 1998).

In many simulation systems synapse models are constrained to transmit a synaptic weight rather than a continuous synaptic current. In such cases, the synaptic weight transmitted to the postsynaptic neuron is $w_{ij}y_0$, assuming the postsynaptic neuron reproduces the dynamics of the y variable. It is not necessary for the neuron to reproduce the y dynamics for each individual synapse; due to the linearity of y between increments, all synapses with the same τ_I can be lumped together. This is the implementation used in NEST (Gewaltig and Diesmann 2007). If the postsynaptic neuron also implements an exact integration scheme (for a worked example see Morrison et al. 2007), the dynamics of y can be incorporated into the propagator of the dynamics of the postsynaptic neuron.

3.2 Abbott model

A simpler model was developed by Abbott et al. (1997), for a complete description see Dayan and Abbott (2001). In this model, synaptic conductance is expressed as $g_s = \overline{g}_s P_s P_{rel}$, where $\overline{g}_s$ is the maximum conductance, P_s is the fraction of open postsynaptic channels and P_{rel} is the fraction of presynaptic sites releasing transmitter. P_s generates the shape of the postsynaptic conductance, and will not be further considered here. Facilitation and depression can both be modeled as presynaptic processes that modify P_{rel}. In both cases, between presynaptic action potentials P_{rel} decays exponentially with a time constant τ_P back to its 'resting' level P_0. In the case of facilitation, a presynaptic spike causes P_{rel} to be increased by $f_F\left(1 - P_{rel}\right)$:

$$\frac{dP_{rel}}{dt} = \frac{P_0 - P_{rel}}{\tau_P} + f_F\left(1 - P_{rel}\right)\delta\left(t - t_j^f\right), \qquad (8)$$

where t_j^f is the timing of the presynaptic spikes, f_F controls the degree of facilitation (with $0 \leq f_F \leq 1$), and the factor $\left(1 - P_{rel}\right)$ prevents the release probability from growing larger than 1. Note that (8) is just a modification of the saturated trace in (2) due to a nonzero 'resting' level.

In the case of depression, activity at the synapse causes P_{rel} to be decreased by $f_D P_{rel}$:

$$\frac{dP_{rel}}{dt} = \frac{P_0 - P_{rel}}{\tau_P} - f_D P_{rel}\delta\left(t - t_j^f\right), \qquad (9)$$

where f_D controls the degree of depression (with $0 \leq f_D \leq 1$), and the factor P_{rel} prevents the release probability from becoming less than 0. Note that with an equilibrium value $P_0 = 1$ (which is always possible) this equation is equivalent to the that of the simplified Tsodyks model without the inactive state (7).

4 Long-term plasticity (STDP)

Experimentally reported STDP curves vary qualitatively depending on the system and the neuron type—see Abbott and Nelson (2000) and Bi and Poo (2001) for reviews. It is therefore obvious that we cannot expect that a single STDP rule, be it defined in the framework of temporal traces outlined above or in a more biophysical framework, would hold for all experimental preparations and across all neuron and synapse types. The first spike-timing experiments were perform by Markram and Sakmann on layer 5 pyramidal neurons in neocortex (Markram et al. 1997). In the neocortex, the width of the negative window seems to vary depending on layer, and inhibitory neurons seem to have a more symmetric STDP curve. The standard STDP curve that has become an icon of theoretical research on STDP (Fig. 1 in Bi and Poo 1998) was originally found for pyramidal neurons in rat hippocampal cell culture. Inverted STDP curves have also been reported, for example in the ELL system in electric fish. This gives rise to different functional properties (Bell et al. 1997).

4.1 Pair-based STDP rules

Most models of STDP interpret the biological evidence in terms of a pair-based update rule, i.e. the change in weight of a synapse depends on the temporal difference between pairs of pre- and postsynaptic spikes:

$$\begin{aligned} \Delta w^+ &= F_+(w) \cdot \exp\left(-\left|\Delta t\right|/\tau_+\right) && \text{if } \Delta t > 0 \\ \Delta w^- &= -F_-(w) \cdot \exp\left(-\left|\Delta t\right|/\tau_-\right) && \text{if } \Delta t \leq 0, \end{aligned} \qquad (10)$$

where $\Delta t = t_i^f - t_j^f$ is the temporal difference between the post- and the presynaptic spikes, and $F_{\pm}(w)$ describes the dependence of the update on the current weight of the synapse. A pair-based model is fully specified by defining: (i) the form of $F_{\pm}(w)$; (ii) which pairs are taken into consideration to perform an update. In order to incorporate STDP into a neuronal network simulation, it is also necessary to specify how the synaptic delay is partitioned into axonal and dendritic contributions.
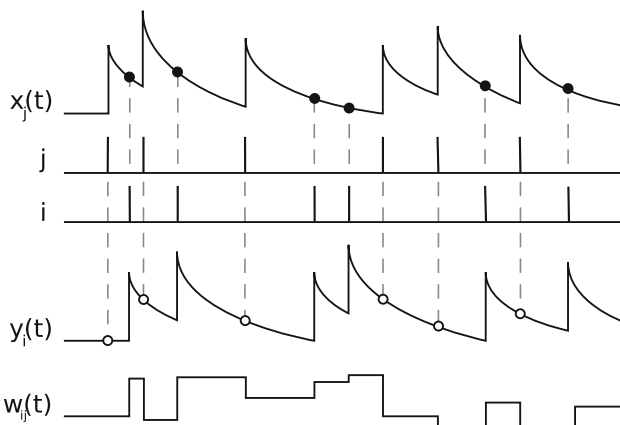
Fig. 3 Implementation of pair-based plasticity by local variables. The spikes of presynaptic neuron j leave a trace $x_j(t)$ and the spikes of the postsynaptic neuron i leave a trace $y_i(t)$. The update of the weight w_{ij} at the moment of a postsynaptic spike is proportional to the momentary value of the trace $x_j(t)$ (*filled circles*). This gives the amount of potentiation due to pre-before-post pairings. Analogously, the update of w_{ij} on the occurrence of a presynaptic spike is proportional to the momentary value of the trace $y_i(t)$ (*unfilled circles*), which gives the amount of depression due to post-before-pre pairings

A pair-based update rule can be easily implemented with two local variables: one for a low-pass filtered version of the presynaptic spike train and one for the postsynaptic spike train. The concept is illustrated in Fig. 3. Let us consider the synapse between neuron j and neuron i. Suppose that each spike from presynaptic neuron j contributes to a trace x_j at the synapse:

$$\frac{\mathrm{d}x_j}{\mathrm{d}t} = -\frac{x_j}{\tau_x} + \sum_{t_j^f} \delta\left(t - t_j^f\right), \tag{11}$$

where t_j^f denotes the firing times of the presynaptic neuron. In other words, the variable is increased by an amount of one at the moment of a presynaptic spike and decreases exponentially with time constant τ_x afterwards; see the discussion of traces in Sect. 2.3. Similarly, each spike from postsynaptic neuron i contributes to a trace y_i:

$$\frac{\mathrm{d}y_i}{\mathrm{d}t} = -\frac{y_i}{\tau_y} + \sum_{t_i^f} \delta\left(t - t_i^f\right), \tag{12}$$

where t_i^f denotes the firing times of the postsynaptic neuron. On the occurrence of a presynaptic spike, a decrease of the weight is induced proportional to the momentary value of the postsynaptic trace y_i. Likewise, on the occurrence of a postsynaptic spike a potentiation of the weight is induced
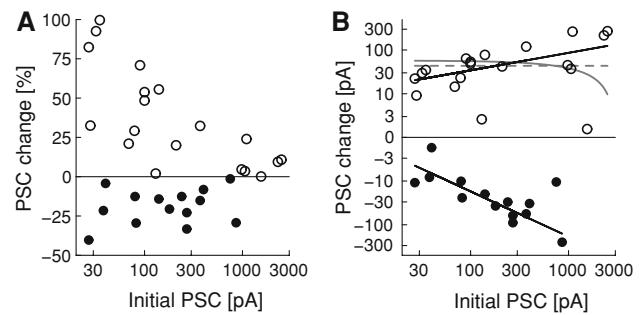
Fig. 4 a Percentage change in peak synaptic amplitude after performing the Bi and Poo (1998) spike pairing protocol as a function of initial synaptic amplitude. Data from Bi and Poo (1998). **b** As in **a**, but the absolute weight change is plotted rather than the percentage weight change, and a double logarithmic representation is used. Fitted additive potentiation update, *gray dashed curve*; fitted multiplicative potentiation update, *solid gray curve*; fitted power law update (slopes 0.4 and −1), *black curves*. Figure adapted from Morrison et al. (2007)

proportional to the trace x_j left by previous presynaptic spikes:

$$\Delta w_{ij}^+\left(t_i^f\right) = F_+\left(w_{ij}\right) x_j\left(t_i^f\right) \tag{13}$$

$$\Delta w_{ij}^-\left(t_j^f\right) = -F_-\left(w_{ij}\right) y_i\left(t_j^f\right), \tag{14}$$

or alternatively:

$$\frac{\mathrm{d}w_{ij}}{\mathrm{d}t} = -F_-\left(w_{ij}\right) y_i(t) \delta\left(t - t_j^f\right)$$
$$+ F_+\left(w_{ij}\right) x_j(t) \delta\left(t - t_i^f\right). \tag{15}$$

A pseudo-code algorithm along these lines for simulating arbitrary pair-based STDP update rules that is suitable for distributed computing is given in Morrison et al. (2007).

Depending on the definition of the trace dynamics (accumulating or saturating, see Sect. 2.3), different spike pairing schemes can be realized. Before we turn to the consequences of these subtle differences (Sect. 4.1.2) and the implementation of synaptic delays (Sect. 4.1.3), we now discuss the choice of the factors $F_+(w)$ and $F_-(w)$, i.e. the weight dependence of STDP.

4.1.1 Weight dependence of STDP

The clearest experimental evidence for the weight dependence of STDP can be found in Bi and Poo (1998), see Fig. 4a. Unfortunately, it is difficult to interpret this figure accurately, as the unit of the ordinate is percentage change, and thus not independent of the value on the abscissa. An additional confounding factor is that the timing interval used in the spike pairing protocol varies considerably across the data. However, even given these drawbacks, the rather flat dependence of the percentage weight change for depression ($\Delta w/w \approx$ constant) suggests a multiplicative dependence
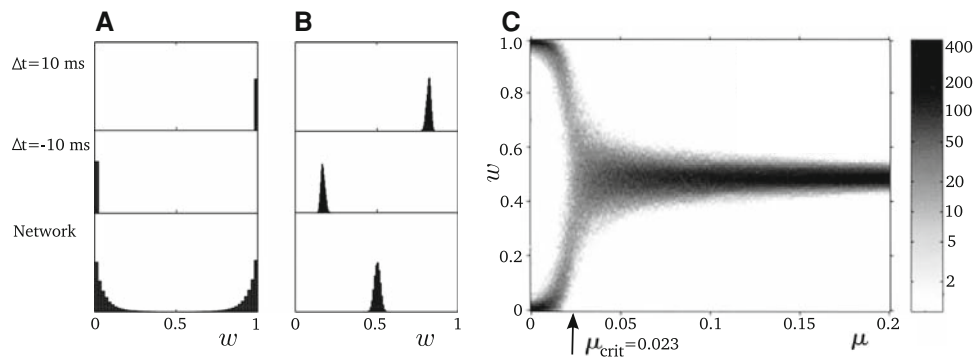
Fig. 5 **a** Histogram showing the equilibrium distribution of synaptic efficacies using the additive STDP rule with $\tau = 10$ ms, $\alpha = 1.05$, $\lambda = 0.005$. The upper two histograms show the behavior of a single synapse. In the *top panel* the presynaptic neuron is repeatedly stimulated before the postsynaptic neuron. In the middle panel the timing relation is reversed. The *bottom* histogram is the distribution of synaptic efficacies in a network of $N = 1,000$ excitatory cells with Poisson activation times and mean input rate $r_{\text{in}} = 10$ Hz. These cells converge on a single conductance-based integrate and fire cell with parameters $\tau_m = 20$ ms, $\tau_s = 5$ ms, $V_s = 5$, and $g_s = 0.01$. **b** As in **a**, but for multiplicative STDP. Adapted from Rubin et al. (2001). **c** Equilibrium weight distribution in logarithmic gray scale as a function of the weight dependence exponent μ for an integrate and fire neuron driven by 1,000 uncorrelated Poisson processes at 10 Hz. Adapted from Gütig et al. (2003)

of depression on the initial synaptic strength ($\Delta w \propto w$). For potentiation the picture is less clear.

Instead of plotting the percentage weight change, Fig. 4b shows the absolute weight change in double logarithmic representation. The exponent of the weight dependence can now be determined from the slope of a linear fit to the data, see Morrison et al. (2007) for more details. A multiplicative update rule ($F_-(w) \propto w$) is the best fit to the depression data but a poor fit to the potentiation data. The best fit to the potentiation data is a power law update ($F_+(w) \propto w^\mu$). The quality of an additive update ($F_+(w) = A_+$) fit is between the power law fit and the multiplicative fit.

4.1.1.1 Unimodal versus bimodal distributions
The choice of update rule can have a large influence on the equilibrium weight distribution in the case of uncorrelated Poissonian inputs. This was first demonstrated by Rubin et al. (2001), see Fig. 5. Here, the behavior of an additive STDP rule ($F_+(w) = \lambda$, $F_-(w) = \lambda\alpha$, where $\lambda \ll 1$ is the learning rate and α an asymmetry parameter) is compared with the behavior of a multiplicative STDP rule ($F_+(w) = \lambda(1-w)$, $F_-(w) = \lambda\alpha w$, with w in the range [0, 1]. In the lowest histograms, the equilibrium distributions are shown for a neuron receiving 1,000 uncorrelated Poissonian spike trains at 10 Hz. In the case of additive STDP, a bimodal distribution develops, whereas in the case of multiplicative STDP, the equilibrium distribution is unimodal. Experimental evidence currently suggests that a unimodal distribution of synaptic strengths is more realistic than the extreme bimodal distribution depicted in Fig. 5a, see, for example, Turrigiano et al. (1998) and Song et al. (2005). Gütig et al. (2003) extended this analysis by regarding additive and multiplicative STDP as the two extrema of a continuous spectrum of rules:

$F_-(w) = \lambda\alpha w^\mu$, $F_+(w) = \lambda(1-w)^\mu$. A choice of $\mu = 0$ results in additive STDP, a choice of $\mu = 1$ leads to multiplicative STDP, and intermediate values result in rules which have an intermediate dependence on the synaptic strength.

Gütig et al. (2003) further demonstrated that the unimodal distribution is the rule rather than the exception for update rules of this form. A bimodal distribution is only produced by rules with a very weak weight dependence (i.e. $\mu \ll 1$). Moreover, the critical value for μ at which bimodal distributions appear decreases as the the effective population size $Nr\tau$ increases, where N is the number of synapses converging onto the postsynaptic neuron, r is the rate of the input spike trains in Hz and τ is the time constant of the STDP window (assumed to be equal for potentiation and depression). Figure 5c shows the equilibrium distributions as a function of μ for $N = 1,000$, $r = 10$ Hz and $\tau = 0.02$ s. μ_{crit} is already very low for this effective population size. Because of the high connectivity of the cortex, we may expect that the effective population size in vivo would be an order of magnitude greater, and so the region of bimodal stability would be vanishingly small according to this analysis. It is worth noting that in the case that a sub-group of inputs is correlated, a bimodal distribution develops for all values of μ, whereby the synaptic weights of the correlated group become stronger than those of the uncorrelated group (data not shown—see Gütig et al. 2003). In contrast to a purely additive rule, the peaks of the distributions are not at the extrema of the permitted weight range. Moreover, the bimodal distribution does not persist if the correlations in the input are removed after learning. A unimodal distribution for uncorrelated Poissonian inputs and an ability to develop multimodal distributions in the presence of correlation is also exhibited by the additive/multiplicative update rule proposed by van Rossum et al.

(2000): $F_+(w) = \lambda$, $F_-(w) = \lambda \alpha w$; and by the power law update rule proposed by Morrison et al. (2007) and also Standage et al. (2007): $F_+(w) \propto \lambda w^\mu$, $F_-(w) \propto \lambda \alpha w$.

4.1.1.2 Fixed point analysis of STDP update rules

An insight into the similarity of behavior of all of these formulations of STDP with the exception of the additive update rule can be obtained by considering their fixed point structure. Equation (10) gives the updates of an individual synaptic weight. If the pre- and postsynaptic spike trains are stochastic, the weight updates can be described as a random walk. Using Fokker–Planck mean field theory, the average rate of change of synaptic strength corresponds to the drift of the random walk, which can be expressed in terms of the correlation between the pre- and postsynaptic spike trains (Kempter et al. 1999; Kistler and van Hemmen 2000; Kempter et al. 2001; Rubin et al. 2001; Gütig et al. 2003). Writing the presynaptic spike train as $\rho_j = \sum_{t_j^f} \delta\left(t - t_j^f\right)$ and the postsynaptic spike train as $\rho_i = \sum_{t_i^f} \delta\left(t - t_i^f\right)$, the mean rates are $\nu_{i/j} = \langle \rho_{i/j} \rangle$. Assuming stationarity, the raw cross-correlation function is given by

$$\Gamma_{ji}(\Delta t) = \langle \rho_j(t)\rho_i(t + \Delta t) \rangle_t \tag{16}$$

i.e. averaging over t while keeping the delay Δt between the two spike trains fixed. The synaptic drift is obtained by integrating the synaptic weight changes given by (10) over Δt weighted by the probability, as expressed by (16), of the temporal difference Δt occurring between a pre- and postsynaptic spike:

$$\dot{w} = -F_-(w)\int_{-\infty}^{0} d\Delta t\, K_-(\Delta t)\,\Gamma_{ji}(\Delta t)$$
$$+ F_+(w)\int_{0}^{\infty} d\Delta t\, K_+(\Delta t)\,\Gamma_{ji}(\Delta t), \tag{17}$$

where $K_\pm(\Delta t) = \exp(-|\Delta t|/\tau_\pm)$, the window function of STDP.

As we are only interested in the qualitative structure of fixed points rather than their exact location, we will simplify the analysis by assuming that the pre- and postsynaptic spike trains are independent Poisson processes with the same rate, i.e. $\langle \rho_i \rangle = \langle \rho_j \rangle = \nu$ and $\Gamma_{ji}(\Delta t) = \nu^2$. We can therefore write:

$$\dot{w} = -F_-(w)\tau_-\nu^2 + F_+(w)\,\tau_+\nu^2.$$

In general, the rate ν of a neuron is dependent on the weight of its incoming synapses and so the right side of this equation cannot be easily determined. However, we can reformulate the equation as:
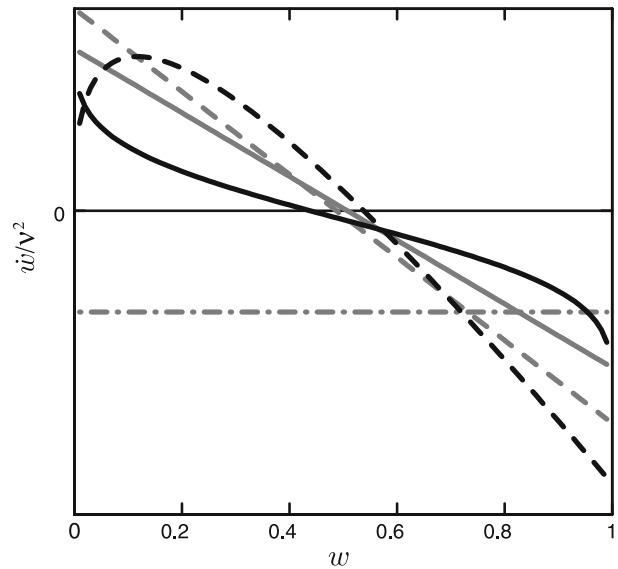
Fig. 6 The fixed point structure of STDP: The synaptic drift $\dot{w}$ scaled by the square of the pre- and postsynaptic firing rate ν is plotted as a function of the synaptic weight w for various formulations of STDP. Model parameters chosen for visual clarity as only the qualitative behavior is relevant. Additive, *gray dash-dotted curve*; multiplicative, *gray dashed curve*; van Rossum, *gray curve*; intermediate Gütig, *black curve*; power law, *dashed black curve*

$$\frac{\dot{w}}{\nu^2} = -F_-(w)\tau_- + F_+(w)\,\tau_+. \tag{18}$$

The fixed points of the synaptic dynamics are given by definition by $\dot{w} = 0$, and therefore also by $\dot{w}/\nu^2 = 0$. Figure 6 plots (18) for a range of w and a variety of STDP models. In all cases except for additive STDP the curves pass through $\dot{w}/\nu^2 = 0$ at an intermediate value of w and with a negative slope, i.e. for weights below the fixed point there is a net potentiating effect, and for weights above the fixed point there is a net depressing effect, resulting in a stable fixed point which is not at an extremum of the weight range. In the case of additive STDP there is no such fixed point, stable or otherwise.

The behavior of the additive model can be assessed more accurately by relaxing the assumption that pre- and postsynaptic spike trains can be described by independent Poisson processes. Instead, we consider a very simple neuron model in which the output spike train is generated by an inhomogeneous Poisson process with rate $\nu_i(u_i) = [\alpha u_i - \nu_0]_+$ with scaling factor α, threshold ν_0 and membrane potential $u_i(t) = \sum_j w_{ij}\epsilon\left(t - t_j^f\right)$, where $\epsilon(t)$ denotes the time course of an excitatory postsynaptic potential generated by a presynaptic spike arrival. The notation $[x]_+$ denotes a piecewise linear function: $[x]_+ = x$ for $x > 0$ and zero otherwise. In the following we assume that the argument of our piecewise linear function is positive so that we can suppress the

square brackets. Assuming once again that all input spike trains are Poisson processes with rate ν, the expected firing rate of the postsynaptic neuron is simply:

$$\nu_i = -\nu_0 + \alpha \nu \bar{\epsilon} \sum_j w_{ij},$$

where $\bar{\epsilon} = \int \epsilon(s) \mathrm{d}s$, the total area under an excitatory post-synaptic potential. The conditional rate of firing given an input spike at time t_j^f is given by

$$\nu_i(t) = -\nu_0 + \alpha \nu \bar{\epsilon} \sum_j w_{ij} + \alpha w_{ij} \epsilon \left(t - t_j^f \right),$$

thus the postsynaptic spike train is correlated with the presynaptic spike trains. This term shows up as additional spike-spike correlations in the correlation function Γ_{ji}. Hence, in addition to the terms in (18), the synaptic dynamics contains a term of the form $\alpha \nu w F_+(w) \int K_+(s) \epsilon(s) \mathrm{d}s$ that is linear rather than quadratic in the presynaptic firing rate (Kempter et al. 1999, 2001). With this additional term, (18) becomes

$$\frac{\dot{w}}{\nu} = \nu_i \left[-F_-(w)\tau_- + F_+(w)\,\tau_+ \right]$$
$$+ \alpha w F_+(w) \int K_+(s)\epsilon(s)\mathrm{d}s. \tag{19}$$

For the multiplicative models the argument hardly changes, but for the additive model it does. For $C = F_-(w)\tau_- - F_+(w)\,\tau_+ > 0$, the additive model has a fixed point which we find by setting the right-hand side of (19) to zero, i.e.

$$0 = -C\nu_i + \alpha w\, C_{ss},$$

where $C_{ss} = F_+(w) \int K_+(s)\epsilon(s)\mathrm{d}s$ denotes the contribution of the spike–spike correlations. In contrast to the curves in Fig. 6, the slope at the zero-crossing is now positive, indicating instability of the fixed point. This instability leads to the formation of a bimodal weight distribution that is typical for the additive model. Despite the instability of individual weights (which move to their upper or lower bounds), the mean firing rate of the neuron is stabilized (Kempter et al. 2001). To see this we consider the evolution of the output rate $\mathrm{d}\nu_i/\mathrm{d}t = \alpha \nu \bar{\epsilon} \sum_j \mathrm{d}w_{ij}/\mathrm{d}t$. Since $\mathrm{d}w_{ij}/\mathrm{d}t = -C\nu_i \nu + \alpha \nu w_{ij} C_{ss}$ and $\alpha \nu \bar{\epsilon} \sum_j w_{ij} = \nu_i + \nu_0$, we can write:

$$\frac{\mathrm{d}\nu_i}{\mathrm{d}t} = -\alpha \nu^2 \bar{\epsilon} N C \nu_i + \alpha \nu \left(\nu_i + \nu_0 \right) C_{ss},$$

where N is the number of synapses converging on the postsynaptic neuron. Thus we have a dynamics of the form:

$$\frac{\mathrm{d}\nu_i}{\mathrm{d}t} = -\frac{(\nu_i - \nu_{FP})}{\tau_\nu}$$

with a fixed point given by:

$$\nu_{FP} = \frac{C_{ss}\nu_0}{N C \nu \bar{\epsilon} - C_{ss}} \tag{20}$$

and a time constant $\tau_\nu = (\alpha \nu [N C \nu \bar{\epsilon} - C_{ss}])^{-1}$. Note that stabilization at a positive rate requires that $\nu_0 > 0$ and $C > 0$. The first condition states that, in the absence of any input, the neuron does not show any spontaneous activity, and this is trivially true for all standard neuron models, including the integrate-and-fire model. The latter condition is equivalent to the requirement that the integral over the STDP curve be negative: $\int \mathrm{d}s[F_+(w)K_+(s) - F_-(w)K_-(s)] = -C < 0$. Exact conditions for stabilization of output rates are given in Kempter et al. (2001). Since for constant input rates ν we have $\nu_i = \nu_0 + \alpha \nu \bar{\epsilon} \sum_j w_{ij}$, stabilization of the output rate implies normalization of the summed weights. Hence STDP can lead to a control of total presynaptic input and of the postsynaptic firing rate – a feature that is usually associated with homeostatic processes rather than Hebbian learning per se (Kempter et al. 1999, 2001; Song et al. 2000).

Note that the existence of a fixed point and its stability does not crucially depend on the presence of soft or hard bounds on the weight. Equations (18) and (19) can equate to zero for hard-bounded or or unbounded rules.

4.1.1.3 Consequences for network stability

Results on the consequences of STDP in large-scale networks are few and far between, and tend to contradict each other. Part of the reason for the lack of simulation papers on this important subject is the fact that simulating such networks consumes huge amounts of memory, is computationally expensive, and potentially requires extremely long simulation times to overcome transients in the weight dynamics which can be of the order of hundreds of seconds of biological time. A lack of theoretical papers on the subject can be explained by the complexity of the interactions between the activity dynamics of the network and the weight dynamics, although some progress is being made in this area (Burkitt et al. 2007).

It was recently shown that power law STDP is compatible with balanced random networks in the asynchronous-irregular regime (Morrison et al. 2007), resulting in a unimodal distribution of weights and no self-organization of structure. This result was verified for Gütig et al. (2003) STDP for an intermediate value of the exponent ($\mu = 0.4$). Although it has not yet been possible to perform systematic tests, it seems likely that all the formulations of STDP with the fixed point structure discussed in Sect. 4.1.1.1 would give qualitatively similar behavior. The results for additive STDP seem to be more contradictory. Izhikevich et al. (2004) reported self-organization of neuronal groups, whereas the chief feature of the networks investigated by Iglesias et al. (2005)

seems to be extensive withering of the synaptic connections. In the former case, it is the existence of many strong synapses which defines the network, in the latter, the presence of many weak ones. This discrepancy may be attributable to different choices for the effective stabilized firing rates (20) in combination with different choices of delays in the network, see Sect. 4.1.3.
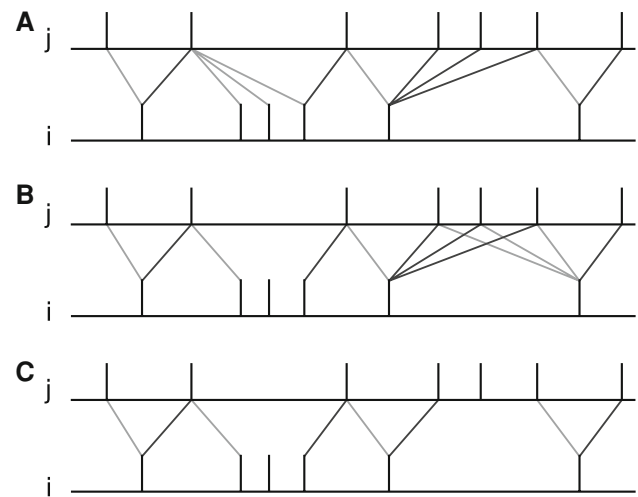
4.1.2 Spike pairing scheme

There are many possible ways to pair pre- and postsynaptic spikes to generate a weight update in an STDP model. In an all-to-all scheme, each presynaptic spike is paired with all previous postsynaptic spikes to effect depression, and each postsynaptic spike is paired with all previous presynaptic spikes to effect potentiation. This is the interpretation used for the fixed point analysis in Sect. 4.1.1.1 and can be implemented using local variables as demonstrated in Sect. 4.1. In a nearest neighbor scheme, only the closest interactions are considered. However, there are multiple possible interpretations of nearest neighbor, as can be seen in Fig. 7. Nearest neighbor schemes can also be realized in terms of appropriately chosen local variables. The symmetric nearest-neighbor scheme shown in Fig. 7a can be implemented by pre- and postsynaptic traces that reset to 1, rather than incrementing by 1 as is the case for the all-to-all scheme. In the case of the presynaptic centered interpretation depicted in Fig. 7B, the postsynaptic trace resets to 1 as in the previous example, but the presynaptic trace must be implemented with a slightly more complicated dynamics:

$$\frac{\mathrm{d}x_j}{\mathrm{d}t} = -\frac{x_j}{\tau_x} + \sum_{t_j^f}\left(1 - x_j^-\right)\delta\left(t - t_j^f\right) - \sum_{t_i^f}x_j^-\delta\left(t - t_i^f\right),$$

where t_j^f and t_i^f denote the firing times of the pre- and postsynaptic neurons respectively, and x_j^- gives the value of x_j just before the update. In other words, the trace is reset to 1 on the occurrence of a presynaptic spike and reset to 0 on the occurrence of a postsynaptic spike. Similarly, the reduced symmetric interpretation shown in Fig. 7c can be implemented by pre- and postsynaptic 'doubly resetting' traces of this form.

It is sometimes assumed that the scheme used makes no difference, as the ISI of cortical network models is typically an order of magnitude larger than the time constant of the STDP window. However, this is not generally true (Kempter et al. 2001; Izhikevich and Desai 2003; Morrison et al. 2007). For a review of a wide variety of schemes and their consequences, particularly with respect to selectivity of higher-frequency inputs, see Burkitt et al. (2004). Experimental results on this issue suggest limited interaction between pairs of spikes. Sjostrom et al. (2001) found that their


Fig. 7 Examples of nearest neighbor spike pairing schemes for a presynaptic neuron j and a postsynaptic neuron i. In each case, the *dark gray* indicate which pairings contribute toward depression of a synapse, and *light gray* indicate which pairings contribute toward potentiation. **a** Symmetric interpretation: each presynaptic spike is paired with the last postsynaptic spike, and each postsynaptic spike is paired with the last presynaptic spike (Morrison et al. 2007). **b** Presynaptic centered interpretation: each presynaptic spike is paired with the last postsynaptic spike and the next postsynaptic spike (Izhikevich and Desai 2003; Burkitt et al. 2004: Model II). **c** Reduced symmetric interpretation: as in **c** but only for immediate pairings (Burkitt et al. 2004: Model IV, also implemented in hardware by Schemmel et al. 2006)

data was best fit by a nearest neighbor interaction similar to Fig. 7c but giving precedence to LTP, i.e. a postsynaptic spike can only contribute to a post-before-pre pairing if it has not already contributed to a pre-before-post pairing. However, this result may also be due to the limitations of pair-based STDP models to explain the experimentally observed frequency dependence, see Sect. 4.2. More recently, Froemke et al. (2006) demonstrated that the amount of LTD was not dependent on the number of presynaptic spikes following a postsynaptic spike, suggesting nearest-neighbor interactions for depression as in Fig. 7c. However, the amount of LTP was negatively correlated with the number of presynaptic spikes preceding a postsynaptic spike. This suggests that multiple spike pairings contribute to LTP, but not in the linear fashion of the all-to-all scheme, which would predict a positive correlation between the number of spikes and the amount of LTP. Again, these results are good evidence for the limitations of pair-based STDP rules.

4.1.3 Synaptic delays

Up until now we have referred to Δt as the temporal difference between a pre- and a postsynaptic spike, i.e. $\Delta t = t_i^f - t_j^f$. However, many classic STDP experiments are expressed in terms of the temporal difference between the start of the EPSP and the postsynaptic spike (Markram et al. 1997;
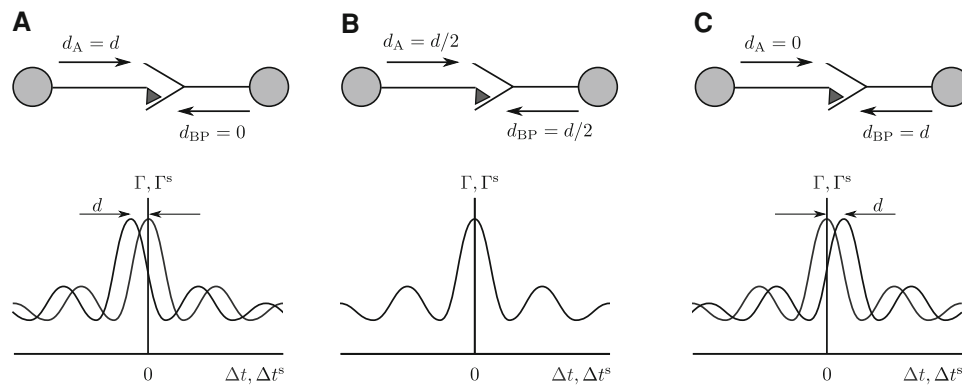
Fig. 8 Different partitions of synaptic delays and the resulting shift of the raw cross-correlation function as perceived at the synapse (*black curves*) with respect to the raw cross-correlation function as perceived at the soma (*gray curves*). The cross-correlation functions shown are purely illustrative and do not result from a specific network model. **a** All synaptic delay is axonal, backpropagation delay is 0. The synaptic raw cross-correlation function is shifted to the left by d. **b** The axonal delay is the same as the backpropagation delay. The synaptic raw cross-correlation function is identical to the somatic raw cross-correlation function. **c** All synaptic delay is dendritic, axonal delay is 0. The synaptic raw cross-correlation function is shifted to the right by d. See Senn et al. (2002) for a more detailed treatment of the partitioning of synaptic delays

Bi and Poo 1998). In fact, when a presynaptic spike is generated at t_j^f, it must first travel down the axon before arriving at the synapse, thus arriving at $t_j^s = t_j^f + d_A$, where d_A is the axonal propagation delay. Similarly, a postsynaptic spike at t_i^f must backpropagate through the dendrite before arriving at the synapse at $t_i^s = t_i^f + d_{BP}$, where d_{BP} is the backpropagation delay. Consequently, the relevant temporal difference for STDP update rules is $\Delta t^s = t_i^s - t_j^s$ as initially suggested by Gerstner et al. (1993) and Debanne et al. (1998). Senn et al. (2002) showed that under fairly general conditions, STDP may cause adaptation in the presynaptic and postsynaptic delays in order to optimize the effect of the presynaptic spike on the postsynaptic neuron. In order to calculate the synaptic drift as in (17), we therefore need to integrate the synaptic weight changes over Δt^s, weighted by the raw cross-correlation function *at the synapse*. With $\Gamma_{ji}(\Delta t) = \Gamma_{ji}(\Delta t^s + (d_A - d_{BP}))$, we reformulate (17) as:

$$
\dot{w} = -F_-(w) \int_{-\infty}^{0} d\Delta t^s K_-\left(\Delta t^s\right) \Gamma_{ji}\left(\Delta t^s + (d_A - d_{BP})\right)
$$
$$
+ F_+(w) \int_{0}^{\infty} d\Delta t^s K_+\left(\Delta t^s\right) \Gamma_{ji}\left(\Delta t^s + (d_A - d_{BP})\right).
$$

$$(21)$$

In the case of independent Poisson processes as in Sect. 4.1.1.1, the shift of the raw cross-correlation function by $(d_A - d_{BP})$ has no effect, as $\Gamma_{ji}(\Delta t)$ is constant. Generally, however, this is not the case. For example, networks of neurons, both in experiment and simulation, typically exhibit oscillations with a period several times larger than the

synaptic delay, even when individual spike trains are irregular (see Kriener et al. 2008, for discussion). If the axonal delay is the same as the backpropagation delay, i.e. $d_A = d_{BP} = d/2$, where d is the total transmission delay of the spike, the raw cross-correlation function at the synapse is the same as the raw cross-correlation at the soma:

$$
\Gamma_{ji}^s\left(\Delta t^s\right) = \Gamma_{ji}\left(\Delta t^s + (d_A - d_{BP})\right) = \Gamma_{ji}\left(\Delta t\right).
$$

This situation is depicted in Fig. 8b. Let w_0 be the synaptic weight for which the synaptic drift given in (21) is 0, i.e. the fixed point of the synaptic dynamics for the cross-correlation shown. If the axonal delay is larger than the backpropagation delay, this results in a shift of the raw cross-correlation function to the left. This is shown in Fig. 8a for the extreme case of $d_A = d$, $d_{BP} = 0$, resulting in a net shift of d. This increases the value of the first integral in (21) and decreases the second integral, such that $\dot{w} < 0$ at w_0. Conversely, if the axonal delay is smaller than the backpropagation delay, the raw cross-correlation function is shifted to the right (Fig. 8c, for the extreme case of $d_A = 0$, $d_{BP} = d$). This decreases the value of the first integral in (21) and increases the second integral, such that $\dot{w} > 0$ at w_0. Therefore, a given network dynamics may cause systematic depression, systematic potentiation or no systematic change at all to the synaptic weights, depending on the partition of the synaptic delay into axonal and dendritic contributions. Systematic synaptic weight changes can in turn result in qualitatively different network behavior. For example, in Morrison et al. (2007) small systematic biases in the synaptic weight dynamics were applied to a network with an equilibrium characterized by a unimodal weight distribution and medium rate (< 10 Hz) asynchronous irregular activity dynamics. Here, a

small systematic depression led to a lower weight, lower rate equilibrium also in the asynchronous irregular regime, whereas a systematic potentiation led to a sudden transition out of the asynchronous irregular regime: the activity was characterized by strongly patterned high-rate peaks of activity interspersed with silence, and the unimodal weight distribution splintered into several peaks.

4.2 Beyond pair effects

There is considerable evidence that the pair-based rules discussed above cannot give a full account of STDP. Specifically, they reproduce neither the dependence of plasticity on the repetition frequency of pairs of spikes in an experimental protocol, nor the results of recent triplet and quadruplet experiments.

STDP experiments are usually carried out with about 60 pairs of spikes. The temporal distance of the spikes in the pair is of the order of a few to tens of milliseconds, whereas the temporal distance between the pairs is of the order of hundreds of milliseconds to seconds. In the case of a facilitation protocol (i.e. pre-before-post), standard pair-based STDP models predict that if the repetition frequency is increased, the strength of the depressing interaction (i.e. post-before-pre) becomes greater, leading to less net potentiation. This prediction is independent of whether the spike pairing scheme is all-to-all or nearest neighbor (see Sect. 4.1.2). However, experiments show that increasing the repetition frequency leads to an increase in potentiation (Sjostrom et al. 2001). Other recent experiments employed multiple-spike protocols, such as repeated presentations of symmetric triplets of the form pre-post-pre and post-pre-post (Bi and Wang 2002; Froemke and Dan 2002; Wang et al. 2005; Froemke et al. 2006). Standard pair-based models predict that the two sequences should give essentially the same results, as they each contain one pre-post pair and one post-pre pair. Experimentally, quite different results are observed.

Here we review two examples of simple models which account for these experimental findings. For other models which also reproduce frequency dependence or multiple-spike protocol results, see Abarbanel et al. (2002), Senn (2002) and Appleby and Elliott (2005).

4.2.1 Triplet model

One simple approach to modeling STDP which addresses these issues is the triplet rule developed by Pfister and Gerstner (2006). This model is based on sets of three spikes (one presynaptic and two postsynaptic). As in the case of pair-based rules, the triplet rule can be easily implemented with local variables as follows. Similarly to pair-based rules, each spike from presynaptic neuron j contributes to a trace x_j at the synapse:
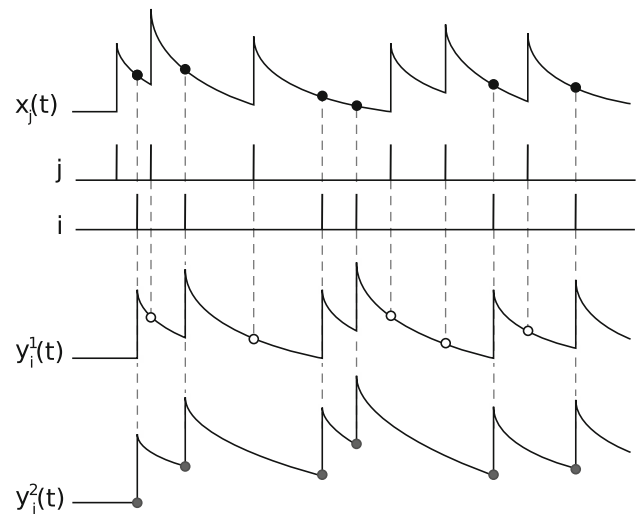
Fig. 9 Implementation of the triplet rule by local variables. The spikes of presynaptic neuron j contribute to a trace $x_j(t)$, the spikes of postsynaptic neuron i contribute to a fast trace $y_i^1(t)$ and a slow trace $y_i^2(t)$. The update of the weight w_{ij} at the moment of a presynaptic spike is proportional to the momentary value of the fast trace $y_i^1(t)$ (*unfilled circles*), as in the pair-based model (see Fig. 3). The update of the weight w_{ij} at the moment of a postsynaptic spike is proportional to the momentary value of the trace $x_j(t)$ (*black filled circles*) and the value of the slow trace $y_i^2(t)$ just before the spike (*gray filled circles*)

$$\frac{dx_j}{dt} = -\frac{x_j}{\tau_x} + \sum_{t_j^f} \delta\left(t - t_j^f\right),$$

where t_j^f denotes the firing times of the presynaptic neuron. Unlike pair-based rules, each spike from postsynaptic neuron i contributes to a fast trace y_i^1 and a slow trace y_i^2 at the synapse:

$$\frac{dy_i^1}{dt} = -\frac{y_i^1}{\tau_1} + \delta\left(t - t_i^f\right)$$
$$\frac{dy_i^2}{dt} = -\frac{y_i^2}{\tau_2} + \delta\left(t - t_i^f\right),$$

where $\tau_1 < \tau_2$, see Fig. 9. LTD is induced as in the standard STDP pair model given in (13), i.e. the weight change is proportional to the value of the fast postsynaptic trace y_i^1 evaluated at the moment of a presynaptic spike. The new feature of the rule is that LTP is induced by a triplet effect: the weight change is proportional to the value of the presynaptic trace x_j evaluated at the moment of a postsynaptic spike and also to the slow postsynaptic trace y_i^2 remaining from previous postsynaptic spikes:

$$\Delta w_{ij}^+\left(t_i^f\right) = F_+\left(w_{ij}\right) x_j\left(t_i^f\right) y_i^2\left(t_i^{f-}\right)$$

where t_i^{f-} indicates that the function y_i^2 is to be evaluated before it is incremented due to the postsynaptic spike at t_i^f.
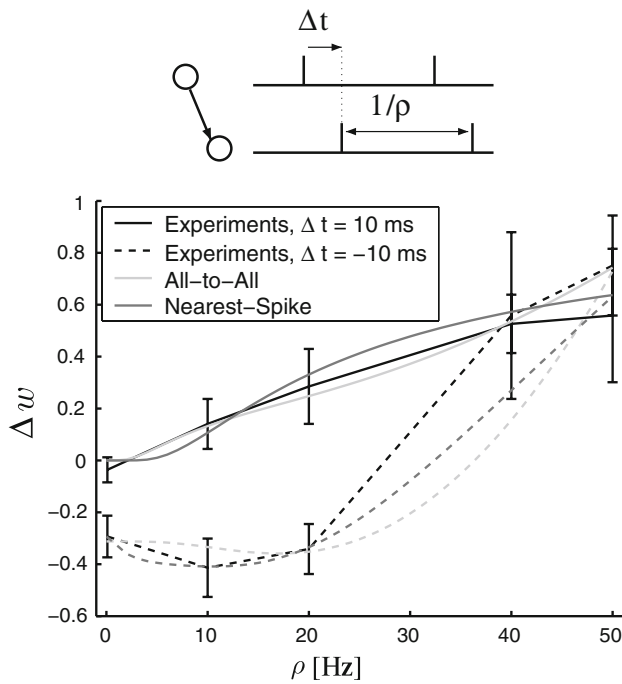
Fig. 10 The triplet rule reproduces the finding that increased frequency of pair repetition leads to increased potentiation in visual cortex pyramidal neurons. Data from Sjostrom et al. (2001), figure adapted from Pfister and Gerstner (2006)

Analogously to pair-based models, the triplet rule can also be implemented with nearest-neighbor rather than all-to-all spike pairings by an appropriate choice of trace dynamics, see Sect. 4.1.2.

The triplet rule reproduces experimental data from visual cortical slices (Sjostrom et al. 2001) that increasing the repetition frequency in the STDP pairing protocol increases net potentiation (Fig. 10). It also gives a good fit to experiments based on triplet protocols in hippocampal culture (Wang et al. 2005). The main functional advantage of such a triplet learning rule is that it can be mapped to a Bienenstock–Cooper–Munro learning rule (Bienenstock et al. 1982): if we assume that the pre- and postsynaptic spike trains are governed by Poisson statistics, the triplet rule exhibits depression for low postsynaptic firing rates and potentiation for high postsynaptic firing rates. If we further assume that the triplet term in the learning rule depends on the mean postsynaptic frequency, a sliding threshold between potentiation and depression can be defined. In this way, the learning rule matches the requirements of the BCM theory and inherits the properties of the BCM learning rule such as input selectivity. From BCM properties, we can immediately conclude that the model should be useful for receptive field development. Note that earlier efforts to show that STDP maps to the BCM model (Izhikevich and Desai 2003; Senn et al. 2000) demonstrated neither an exact mapping nor a sliding threshold. The

exact relationship between the above triplet model and other models is discussed in Pfister and Gerstner (2006).

4.2.2 Suppression model

An alternative model to address the inability of standard pair-based models to account for data obtained from triplet and quadruplet spike protocols was developed by Froemke and Dan (2002). They observed that in triplet protocols of the form pre-post-pre, as long as the intervals between the spikes were reasonably short (< 15 ms), the timing of the pre–post pair was a better predictor for the change in the synaptic strength than either the timing of the post–pre pair or of both timings taken together. Similarly, in post–pre–post protocols, the timing of the first post-pre pairing was the best predictor for the change of synaptic strength. On the basis of this observation, they proposed a model in which the synaptic weight change is not just dependent on the timing of a spike pair, but also on the *efficacy* of the spikes. Each spike of presynaptic neuron j sets the presynaptic spike efficacy ϵ_j to 0 whereafter it recovers exponentially to 1 with a time constant τ_j. The efficacy of the nth presynaptic spike is given by:

$$\epsilon_j^n = 1 - \exp\left(-\left(t_j^n - t_j^{n-1}\right)/\tau_j\right),$$

where t_j^n denotes the nth spike of neuron j. In other words, the efficacy of a spike is suppressed by the proximity of a previous spike. Similarly, the postsynaptic spike efficacy is reset to 0 by each spike of postsynaptic neuron i, recovering exponentially to 1 with time constant τ_i. The model can be implemented with local variables as follows. Each presynaptic spike contributes to an efficacy trace $\epsilon_j(t)$ with dynamics:

$$\frac{d\epsilon_j}{dt} = -\frac{\epsilon_j - 1}{\tau_j} - \sum_{t_j^f} \epsilon_j^- \delta\left(t - t_j^f\right),$$

where ϵ_j^- denotes the value of ϵ_j just before the update. The standard presynaptic trace x_j given in (11) is adapted to take the spike efficacy into account:

$$\frac{dx_j}{dt} = -\frac{x_j}{\tau_x} + \sum_{t_j^f} \epsilon_j^- \delta\left(t - t_j^f\right),$$

i.e. each presynaptic spike increments x_j by the value of the spike efficacy before the update. Similarly, each postsynaptic spike contributes to an efficacy trace $\epsilon_i(t)$ with dynamics:

$$\frac{d\epsilon_i}{dt} = -\frac{\epsilon_i - 1}{\tau_i} - \sum_{t_i^f} \epsilon_i^- \delta\left(t - t_i^f\right),$$

and a postsynaptic trace y_i with increments weighted by the postsynaptic spike efficacy:

$$\frac{dy_i}{dt} = -\frac{y_i}{\tau_y} + \sum_{t_i^f} \epsilon_i^- \delta\left(t - t_i^f\right).$$

The weight updates on the occurrence of a post- or presynaptic spike are therefore given by:

$$\Delta w_{ij}^+\left(t_i^f\right) = F_+\left(w_{ij}\right) x_j\left(t_i^f\right) \epsilon_i^-\left(t_i^f\right)$$
$$\Delta w_{ij}^-\left(t_j^f\right) = -F_-\left(w_{ij}\right) y_i\left(t_j^f\right) \epsilon_j^-\left(t_j^f\right).$$

This model gives a good fit to triplet and quadruplet protocols in visual cortex slice, and also gives a much better prediction for synaptic modification due to natural spike trains (Froemke and Dan 2002). However, it does not predict the increase of LTP with the repetition frequency observed by Sjostrom et al. (2001). A revised version of the model (Froemke et al. 2006) also accounts for the switch of LTD to LTP at high frequencies by modifying the efficacy functions.

4.3 Voltage dependence

Traditional LTP/LTD experiments employ the following induction paradigm: the postsynaptic neuron is held at a fixed depolarization while one or several presynaptic neurons are activated. Often a presynaptic pathway is stimulated extracellularly, so that several presynaptic neurons are activated. Depending on the level of the postsynaptic membrane potential, the activated synapses increase their efficacy while other non-activated synapses do not change their weight (Artola et al. 1990; Artola and Singer 1993). More recently, depolarization has also been combined with STDP experiments. In particular, Sjostrom et al. (2004) showed a dependence of synaptic weight changes on the synaptic membrane potential just before a postsynaptic spike.

There is an ongoing discussion whether the voltage dependence is more fundamental than the dependence on postsynaptic spiking. Indeed, voltage dependence alone can generate STDP-like behavior (Brader et al. 2007), as the membrane potential behaves in a characteristic way in the vicinity of a spike (high shortly before a spike, and low shortly after). Alternatively, a dependence on the slope of the postsynaptic membrane potential has also been shown to reproduce the characteristic STDP weight change curve (Saudargiene et al. 2003). The voltage effects caused by backpropagating spikes is implicitly contained in the mechanistic formulation of STDP models outlined above. In particular, the fast postsynaptic trace y_1 in the above triplet model could be seen as an approximation of a backpropagating action potential. However, the converse is not true: a pure STDP rule does not automatically generate a voltage dependence. Moreover, synaptic effects caused by subthreshold

depolarization in the absence of postsynaptic firing cannot be modeled by standard STDP or triplet models.

4.4 Induction versus maintenance

We stress that all the above models concern induction of potentiation and depression, but not their maintenance. The induction of LTP may take only a few seconds: for example, stimulation with 50 pairs of pre- and postsynaptic spikes given at 20 Hz takes less than 3 s. However, afterwards the synapse takes 60 min or more to consolidate these changes, and this process may also be interrupted (Frey and Morris 1997). During this time synapses are 'tagged', that is, they are ready for consolidation. Consolidation is thought to rely on a different molecular mechanism than that of induction. Simply speaking, gene transcription is necessary to trigger the building of new proteins that increase the synaptic efficacy.

4.4.1 Functional consequences

Long-term stability of synapses is necessary to retain memories that have been learned, despite ongoing activity of presynaptic neurons. A simple possibility used in many models is that plasticity is simply switched off once the neuron has learned what it should. This approach makes sense in the context of reward-based learning: the learning rate goes to zero once the actual reward equals the expected reward and learning stops automatically (see Sect. 5.2). It also makes sense in the framework of supervised learning (see Sect. 5.1). Learning is normally driven by the difference between desired output and actual output. However, in the context of unsupervised learning it is inconsistent to switch off the dynamics. Nevertheless, receptive field properties should be retained for a fairly long time even if the stimulation characteristic changes.

4.4.2 Bistability model

A simple model of maintenance has been proposed by Fusi et al. (2000). The basis of the model is a hidden variable that has an unstable fixed point (threshold). If the variable has a value above threshold it converges towards 1; otherwise towards 0. To stay within the framework of the previous sections, let us suppose that the weight w is calculated by one of the STDP or short-term plasticity models. Maintenance is implemented by adding on top of the STDP dynamics a slow bistable dynamics (Gerstner and Kistler 2002):

$$\frac{dw}{dt} = \alpha(w) = -w\,(1-w)(w_\theta - w)/\tau_\alpha,$$

where τ_α is a time constant of consolidation in the range of several minutes of biological time. The result is that in the absence of any stimulation, individual synapses evolve

towards binary values of 0 or 1 which are intrinsically stable fixed points of the slow dynamics. As a result, rather strong stimuli are necessary to perturb the synaptic dynamics.

4.4.3 Biological evidence

Whether single synapses themselves are binary or continuous is a matter of intense debate. Some experiments have suggested that synapses are binary (Petersen et al. 1998; O'Connor et al. 2005). However, this would seem to result in a bistable distribution of weights which is at odds with the unimodal distribution reported by other studies (Turrigiano et al. 1998; Sjostrom et al. 2001; Song et al. 2005), and with the finding that the magnitude of LTP/LTD increases with the number of spike pairs in a protocol until saturation is reached (Froemke et al. 2006).

Some possibilities to reconcile these findings include: (i) since pairs of neurons form several contacts with each other, it is likely that in standard plasticity experiments several synapses are measured at the same time; (ii) LTP and STDP results are typically reported as pooled experiments over several pairs of neurons. Under the assumption that the upper bound is not the same for all synapses, a broad distribution could result; (iii) both unimodal distribution and bimodal distributions could be stable. Untrained neurons would show a unimodal distribution whereas neurons that have learned to respond to a specific pattern would develop a bimodal distribution of synaptic weights (Toyoizumi et al. 2007); (iv) all synapses are binary, but the efficacy of the 'strong' state is subject to short-term plasticity and homeostasis; (v) some synapses are binary and some are not. Potentially a combination of several of these possibilities must be considered in order to explain the experimental findings.

5 Supervised and reinforcement learning

All the models considered in Sect. 4 are unsupervised 'Hebbian' rules: changes are triggered as a result of combined action of pre- and postsynaptic neurons. The postsynaptic neuron itself is driven by its input arising from presynaptic neurons. There is no notion of whether or not the postsynaptic output is 'good' or 'useful'. If, however, the local variables are combined with global teacher or reinforcement signals, completely different learning paradigms are possible.

5.1 Supervised learning

Supervised plasticity has been demonstrated experimentally by Fregnac and Schulz (2006): the behavior of a (cortical) neuron can be changed by pairing some class of stimuli with an (artificial) increase of neural activity while pairing another class of stimuli with a decrease of responsiveness.

Theoretical studies have demonstrated that a teacher-forced STDP approach can be used to learn precise spike times (Legenstein et al. 2005; Pfister et al. 2006). In a natural situation, this would mean that a few strong neural inputs can drive the neuron and therefore drive learning of other inputs. If these strong inputs are controlled in a task-specific way, they act as a teacher for the postsynaptic neuron. For a practical realization of this idea see Brader et al. (2007).

5.2 Reinforcement learning

If neuronal activity leads to actions, feedback may arise from the environment in forms of reward (a piece of pizza) or punishment (burnt fingers). It is thought that success of an action is signaled by neuromodulators—a top candidate is dopamine (Schultz et al. 1997). Dopamine signals are closely related to a quantity in reinforcement learning known as δ, that can be interpreted as the difference between the received reward and the expected reward. Here 'reward' means current or future rewards that can be reliably predicted. In reinforcement learning, the difference between actual and expected rewards plays an important role for the update of weights in Q-learning, SARSA, and related variants of temporal difference learning (Sutton and Barto 1998).

Under a suitable interpretation of the role of pre- and postsynaptic neurons, the weight update rules can be derived from an optimality framework (Pfister et al. 2006). The learning rule can be interpreted as a Hebbian learning based on joined pre- and postsynaptic activity, but conditioned on the presence of a global reward signal. Variants of such reinforcement rules for spiking neurons have been developed (Seung 2003; Pfister et al. 2006; Izhikevich 2007; Florian 2007).

6 Discussion

Pair-based STDP models can be decomposed into three aspects: weight dependence, spike-pairing scheme and delay partition (Sect. 4.1). We have shown that all of these aspects can have significant consequences for the behavior of the model system under investigation. However, in many cases there is not enough experimental data to settle these questions definitively. Therefore, choices for each aspect should be made consciously and take into consideration the relevant available experimental findings. Moreover, these choices should be explicitly documented and critically addressed: it should be clear to what extent results depend on the specific choices.

In particular, the choice of STDP weight dependence is critical. The available evidence suggests that both potentiation and depression are dependent on the weight. Whereas it is useful to start with very simplified models to gain insight, we now know that STDP models which assume some weight

dependence produce qualitatively different behavior from the additive model. Moreover, weight dependent rules are no harder to implement computationally than additive rules. In the absence of fresh experimental evidence supporting an additive rule, weight dependent rules should therefore be considered as the standard.

Pair-based models of STDP have their limitations. They give incorrect predictions for many experiments such as triplet and quadruplet protocols and cannot account for synaptic modification due to natural spike trains or pairing protocols at different frequencies. Models of STDP that are beyond the pair-based framework (Sect. 4.2) can account for these findings at the cost of only a small number of additional variables, and so should attract increasing theoretical interest.

In this manuscript, we have considered models in which synaptic modifications depend only on spike timing. However, this ignores many aspects of synaptic plasticity which may prove to be of great importance to the functioning of the brain, and will therefore have to be taken into consideration in future phenomenological modeling. Most STDP models assume that the absolute synaptic strength is modified (but see Senn 2002). However, it may turn out that a formulation in terms of the release probability is a more accurate description, thus allowing a unified view of short-term and long-term plasticity. Additionally, STDP has been shown to be sensitive to a number of factors beyond spike timing, for example active dendritic properties and the location of the synapse on the dendrite - see Kampa et al. (2007) for a review. There is also substantial evidence that inhibition is an important physiological feature fine-tuning induction and maintenance of LTP/LTD. Inhibition gates induction of LTP/LTD as a function of physiological conditions and physiologically-induced changes in the activity of networks (Larson and Lynch 1986; Pacelli et al. 1989; Radpour and Thomson 1991; Steele and Mauk 1999; Nishiyama et al. 2000; Togashi et al. 2003). Here, the main challenge is to derive appropriate phenomenological models from experiments and detailed biophysical models. Finally, although some progress has been made in investigating the interactions of STDP with other plasticity mechanisms such as homeostasis and heterosynaptic spread of LTP/LTD (van Rossum et al. 2000; Toyoizumi et al. 2005, 2007; Triesch 2007), this complex topic remains largely unexplored. In this area, the main challenge is to perform analytical and simulation studies which can identify and characterize their composite effects, and investigate their functional consequences.

References

Abarbanel H, Huerta R, Rabinovich M (2002) Dynamical model of long-term synaptic plasticity. Proc Natl Acad Sci USA 99(15): 10132–10137

Abbott LF, Nelson SB (2000) Synaptic plasticity: taming the beast. Nat Neurosci 3(Suppl):1178–1183

Abbott LF, Varela JA, Sen K, Nelson SB (1997) Synaptic depression and cortical gain control. Science 275:220–223

Appleby P, Elliott T (2005) Synaptic and temporal ensemble interpretation of spike-timing-dependent plasticity. Neural Comput 17(11): 2316–2336

Artola A, Bröcher S, 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

Artola A, Singer W (1993) Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. Trends Neurosci 16(11):480–487

Badoual M, Zou Q, Davison AP, Rudolph M, Bal T, Fregnac Y, Destexhe A (2006) Biophysical and phenomenological models of multiple spike interactions in spike-timing dependent plasticity. Int J Neural Systems 16:79–97

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

Bi G-q, Poo M (2001) Synaptic modification by correlated activity: Hebb's postulate revisited. Annu Rev Neurosci 24:139–166

Bi G-q, Wang H (2002) Temporal asymmetry in spike timing-dependent synaptic plasticity. Physiol Behav 77:551–555

Bi G-q, Poo M-m (1998) Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci 18:10464–10472

Bienenstock EL, Cooper LN, Munro PW (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. J Neurosci 2(1):32–48

Billings G, van Rossum M (2008) Memory retention and spike timing dependent plasticity (preprint)

Bliss TVP, Collingridge GL (1993) A synaptic model of memory: long-term potentiation in the hippocampus. Nature 361:31–39

Bliss TVP, Lomo T (1973) Long-lasting potentiation of synaptic transmission in the dendate area of anaesthetized rabbit following stimulation of the perforant path. J Physiol 232:331–356

Brader JM, Senn W, Fusi S (2007) Learning real world stimuli in a neural network with spike-driven synaptic dynamics. Neural Comput 19(11):2881–2912

Burkitt AN, Gilson M, van Hemmen JL (2007) Spike-timing-dependent plasticity for neurons with recurrent connections. Biol Cybern 96(5):533–546

Burkitt AN, Meffin H, Grayden DB (2004) Spike-timing-dependent plasticity: the relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. Neural Comput 16:885–940

Cooper L, Intrator N, Blais B, Shouval HZ (2004) Theory of cortical plasticity. World Scientific, Singapore

Dayan P, Abbott LF (2001) Theoretical neuroscience. MIT Press, Cambridge

Debanne D, Gähwiler BH, Thompson SM (1998) Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. J Physiol (Lond) 507:237–247

Dudek SM, Bear MF (1992) Homosynaptic long-term depression in area CA1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. Proc Natl Acad Sci USA 89:4363–4367

Florian RV (2007) Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. Neural Comput 19:1468–1502

Fregnac Y, Schulz DE (2006) Activity-dependent regulation of receptive field properties of cat area 17 by supervised Hebbian learning. J Neurobio 41:69–82

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

Froemke R, Tsay I, Raad M, Long J, Dan Y (2006) Contribution of individual spikes in burst-induced long-term synaptic modification. J Neurophysiol 95:1620–1629

Froemke RC, Dan Y (2002) Spike-timing-dependent synaptic modification induced by natural spike trains. Nature 416(6879):433–438

Fusi S, Annunziato M, Badoni D, Salamon A, Amit DJ (2000) Spike-driven synaptic plasticity: Theory, simulation, VLSI implementation. Neural Comput 12(10):2227–2258

Fusi S, Drew PJ, Abbott LF (2005) Cascade models of synaptically stored memories. Neuron 45(4):599–611

Gerstner W, Kempter R, van Hemmen JL, Wagner H (1996) A neuronal learning rule for sub-millisecond temporal coding. Nature 383:76–78

Gerstner W, Kistler W (2002) Spiking neuron models: single neurons, populations, plasticity. Cambridge University Press, Cambridge

Gerstner W, Ritz R, van Hemmen JL (1993) Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. Biol Cybern 69(5–6):503–515

Gewaltig M-O, Diesmann M (2007) NEST (neural simulation tool). Scholarpedia 2(4):1430

Graupner M, Brunel N (2007) STDP in a bistable synapse model based on CaMKII and associated signaling pathways. Public Library Sci Comput Biol 3(11):e221

Grossberg S (1987) The adaptive brain I. Elsevier, Amsterdam

Gupta A, Wang Y, Markram H (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. Science 287:273–278

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

Gütig R, Aharonov R, Rotter S, Sompolinsky H (2003) Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. J Neurosci 23(9):3697–3714

Hebb DO (1949) The organization of behavior: a neuropsychological theory. Wiley, New York

Iglesias J, Eriksson J, Grize F, Tomassini M, Villa A (2005) Dynamics of pruning in simulated large-scale spiking neural networks. Biosystems 79:11–20

Izhikevich EM (2007) Solving the distal reward problem through linkage of STDP and dopamine signaling. Cereb Cortex 17(10):2443–2452

Izhikevich EM, Desai NS (2003) Relating STDP to BCM. Neural Comput 15:1511–1523

Izhikevich EM, Gally JA, Edelman GM (2004) Spike-timing dynamics of neuronal groups. Cereb Cortex 14:933–944

James W (1890) Psychology (briefer course). Holt, New York

Kampa BM, Letzkus JJ, Stuart GJ (2007) Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. Trends Neurosci 30(9):456–463

Kandel ER, Schwartz JH, Jessel TM (2000) Principles of neural science, 4th edn. McGraw-Hill, New York. ISBN 978-0838577011

Kelso SR, Ganong AH, Brown TH (1986) Hebbian synapses in hippocampus. Proc Natl Acad Sci USA 83:5326–5330

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

Kempter R, Gerstner W, van Hemmen JL (2001) Intrinsic stabilization of output rates by spike-based Hebbian learning. Neural Comput 12:2709–2742

Kistler WM, van Hemmen JL (2000) Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic action potentials. Neural Comput 12:385–405

Kriener B, Tetzlaff T, Aertsen A, Diesmann M, Rotter S (2008) Correlations and population dynamics in recurrent cortical networks. Neural Comput (in press)

Larson J, Lynch G (1986) Induction of synaptic potentiation in hippocampus by patterned stimulation involves two events. Science 232:985–988

Legenstein R, Naeger C, Maass W (2005) What can a neuron learn with spike-timing-dependent plasticity? Neural Comput 17(11):2337–2382

Levy WB, Steward D (1983) Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. Neuroscience 8:791–797

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

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

Lu J, Li C, Zhao JP, Poo M-m, Zhang X (2007) Spike-timing-dependent plasticity of neocortical excitatory synapses on inhibitory interneurons depends on target cell type. J Neurosci 27:9711–9720

Malenka RC, Kauer J, Zucker R, Nicoll RA (1988) Postsynaptic calcium is sufficient for potentiation of hippocampal synaptic transmission. Science 242:81–84

Malenka RC, Nicoll RA (1993) NMDA-receptor-dependent plasticity: multiple forms and mechanisms. Trends Neurosci 16:480–487

Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275:213–215

Markram H, Sakmann B (1995) Action potentials propagating back into dendrites trigger changes in efficacy of single-axon synapses between layer V pyramidal neurons. Soc Neurosci Abstr 21(3):2007

Markram H, Wang Y, Tsodyks M (1998) Differential signaling via the same axon of neocortical pyramidal neurons. Proc Natl Acad Sci USA 95(9):5323–5328

Miller K, Keller JB, Stryker MP (1989) Ocular dominance column development: analysis and simulation. Science 245:605–615

Morrison A, Aertsen A, Diesmann M (2007) Spike-timing dependent plasticity in balanced random networks. Neural Comput 19:1437–1467

Morrison A, Straube S, Plesser HE, Diesmann M (2007) Exact subthreshold integration with continuous spike times in discrete time neural network simulations. Neural Comput 19(1):47–79

Ngezahayo A, Schachner M, Artola A (2000) Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. J Neurosci 20(7):2451–2458

Nishiyama M, Hong K, Mikoshiba K, Poo M, Kato K (2000) Calcium stores regulate the polarity and input specificity of synaptic modification. Nature 408(6812):584–588

O'Connor D, Wittenberg G, Wang S-H (2005) Graded bidirectional synaptic plasticity is composed of switch-like unitary events. Proc Natl Acad Sci USA 102:9679–9684

Oja E (1982) A simplified neuron model as a principal component analyzer. J Math Biol 15:267–273

Pacelli GJ, Sue W, Keslo SR (1989) Activity-induced depression of synaptic inhibition during LTP-inducing patterned stimulation. Brain Res 486:26–32

Petersen C, Malenka R, Nicoll R, Hopfield J (1998) All-or-none potentiation of CA3-CA1 synapses. Proc Natl Acad Sci USA 95:4732–4737

Pfister J-P, Gerstner W (2006) Triplets of spikes in a model of spike timing-dependent plasticity. J Neurosci 26:9673–9682

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

Radpour S, Thomson AM (1991) Coactivation of local circuit NMDA receptor mediated EPSPs induces lasting enhancement of minimal schaffer collateral epsps in slices of rat hippocampus. Eur J Neurosci 3:602–613

Roberts PD (1999) Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. J Comput Neurosci 7:235–246

Rotter S, Diesmann M (1999) Exact digital simulation of time-invariant linear systems with applications to neuronal modeling. Biol Cybern 81(5/6):381–402

Rubin J, Lee D, Sompolinsky H (2001) Equilibrium properties of temporally asymmetric Hebbian plasticity. Phys Rev Lett 86:364–367

Rubin JE, Gerkin RC, Bi G-q, Chow CC (2005) Calcium time course as a signal for spike-timing-dependent plasticity. J Neurophysiol 93:2600–2613

Saudargiene A, Porr B, Wörgötter F (2003) How the shape of pre- and postsynaptic signals can influence STDP: a biophysical model. Neural Comput 16:595–626

Schemmel J, Gruebl A, Meier K, Mueller E (2006) Implementing synaptic plasticity in a VLSI spiking neural network model. In: Proceedings of the 2006 international joint conference on neural networks. IEEE Press, pp 1–6

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593–1599

Senn W (2002) Beyond spike timing: the role of nonlinear plasticity and unreliable synapses. Biol Cybern 87:344–355

Senn W, Markram H, Tsodyks M (2000) An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. Neural Comput 13:35–67

Senn W, Schneider M, Ruf B (2002) Activity-dependent development of axonal and dendritic delays, or, why synaptic transmission should be unreliable. Neural Comput 14(3):583–619

Seung HS (2003) Learning spiking neural networks by reinforcement of stochastic synaptic transmission. Neuron 40:1063–1073

Shouval HZ, Bear MF, Cooper LN (2002) A unified model of NMDA receptor dependent bidirectional synaptic plasticity. Proc Natl Acad Sci USA 99:10831–10836

Sjostrom P, Turrigiano G, Nelson S (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 32:1149–1164

Sjostrom PJ, Turrigiano GG, Nelson SB (2004) Endocannabinoid-dependent neocortical layer-5 LTD in the absence of postsynaptic spiking. J Neurophysiol 92(6):3338–3343

Song S, Miller KD, Abbott LF (2000) Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat Neurosci 3(9):919–926

Song S, Per S, Reigl M, Nelson S, Chklovskii D (2005) Highly nonrandom features of synaptic connectivity in local cortical circuits. Public Library Sci Biol 3(3):0507–0519

Standage D, Jalil S, Trappenberg T (2007) Computational consequences of experimentally derived spike-time and weight dependent plasticity rules. Biol Cybern 96(6):615–623

Steele PM, Mauk MD (1999) Inhibitory control of LTP and LTD: stability of synapse strength. J Neurophysiol 81:1559–1566

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Adaptive Computation and Machine Learning. The MIT Press

Thomson AM, Deuchars J, West DC (1993) Large, deep layer pyramid-pyramid single axon EPSPs in slices of rat motor cortex display paired pulse and freuquency-dependent depression, mediated presynaptically and self-facilitation mediated postsynaptically. J Neurophysiol 70(6):2354–2369

Thomson AM, Lamy C (2007) Functional maps of neocortical local circuitry. Front Neurosci 1:19–42

Togashi K, Kitajima T, Aihara T, Hong K, Poo M, Nishiyama M (2003) Gating of activity-dependent long-term depression by GABAergic activity in the hippocampus. In: SocNeurosciAbstr, pp 123.4

Toyoizumi T, Pfister J, Aihara K, Gerstner W (2007) Optimality model of unsupervised spike-timing-dependent plasticity: synaptic memory and weight distribution. Neural Comput 19(3):639–671

Toyoizumi T, Pfister JP, Aihara K, Gerstner W (2005) Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission. Proc Natl Acad Sci USA 102(14):5239–5244

Triesch J (2007) Synergies between intrinsic and synaptic plasticity mechanisms. Neural Comput 19(4):885–909

Tsodyks M, Pawelzik K, Markram H (1998) Neural networks with dynamic synapses. Neural Comput 10:821–835

Tsodyks M, Uziel A, Markram H (2000) Synchrony generation in recurrent networks with frequency-dependent synapses. J Neurosci 20, RC1 (1–5)

Tsodyks MV, Markram H (1997) The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. Proc Natl Acad Sci USA 94:719–723

Turrigiano G, Abbott LF, Marder E (1994) Activity-dependent changes in the intrinsic properties of pyramidal neurons. Science 264:974–977

Turrigiano GG, Leslie KR, Desai NS, Rutherford LC, Nelson SB (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. Nature 391:892–896

Turrigiano GG, Nelson SB (2004) Homeostatic plasticity in the developing nervous system. Nat Rev Neurosci 5:97–107

Rossum MCWvan , Bi G-q, Turrigiano GG (2000) Stable Hebbian learning from spike timing-dependent plasticity. J Neurosci 20(23):8812–8821

Wang H-X, Gerkin RC, Nauen DW, Bi G-q (2005) Coactivation and timing-dependent integration of synaptic potentiation and depression. Nat Neurosci 8(2):187–193

Zou Q, Destexhe A (2007) Kinetic models of spike-timing dependent plasticity and their functional consequences in detecting correlations. Biol Cybern 97(1):81–97