Development/Plasticity/Repair

Stable Learning in Stochastic Network States

Sami El Boustani,* Pierre Yger,* Yves Frégnac, and Alain Destexhe

Unité de Neurosciences, Information et Complexité, CNRS, Unité Propre de Recherche 3293, 91198 Gif-sur-Yvette, France

The mammalian cerebral cortex is characterized *in vivo* by irregular spontaneous activity, but how this ongoing dynamics affects signal processing and learning remains unknown. The associative plasticity rules demonstrated *in vitro*, mostly in silent networks, are based on the detection of correlations between presynaptic and postsynaptic activity and hence are sensitive to spontaneous activity and spurious correlations. Therefore, they cannot operate in realistic network states. Here, we present a new class of spike-timing-dependent plasticity learning rules with local floating plasticity thresholds, the slow dynamics of which account for metaplasticity. This novel algorithm is shown to both correctly predict homeostasis in synaptic weights and solve the problem of asymptotic stable learning in noisy states. It is shown to naturally encompass many other known types of learning rule, unifying them into a single coherent framework. The mixed presynaptic and postsynaptic dependency of the floating plasticity threshold is justified by a cascade of known molecular pathways, which leads to experimentally testable predictions.

Introduction

Various forms of synaptic plasticity have been discovered and characterized in the mammalian brain experimentally as well as theoretically. One type of associative plasticity is based on the relative spike timing between presynaptic and postsynaptic cells (Levy and Steward, 1983; Bell et al., 1997; Markram et al., 1997; Bi and Poo, 1998). This so-called spike-timing-dependent plasticity (STDP) has been shown to produce an asymmetric and reversible form of synaptic potentiation and depression depending on the temporal order of presynaptic and postsynaptic cell spikes. Models have been proposed to capture such mechanisms and have been very successful in accounting for the emergence of receptive fields, hippocampal place cells, synaptic competition, and stability (Gerstner et al., 1996; Song et al., 2000; van Rossum et al., 2000; Song and Abbott, 2001; Gütig et al., 2003; Yu et al., 2008). However, previous models suffer from an important limitation: their sensitivity to spontaneous activity. Because cerebral cortex is characterized by sustained and irregular spontaneous activity (Softky and Koch, 1993; Shadlen and Newsome, 1998), classical STDP models cannot be used to learn in such stochastic-like activity states, leading to either a complete loss of memory or a shift of the network dynamics toward a pathological synchronous state (Morrison et al., 2007). Because of this "catastrophic-forgetting" effect, memory retention in *in vivo*-like network states seems inconsistent with STDP.

We investigate a natural way to circumvent this problem through metaplasticity. Since early work on long-term potentiation (LTP) and depression (LTD), the need for higher-order regulation of plasticity rules has been evident. This process, termed metaplasticity (Deisseroth et al., 1995; Abraham and Bear, 1996), has been shown to take place through various molecular pathways and to be critical for the stabilization of learned statistics (Abraham, 2008). One major role of metaplasticity is to promote in a conditional way (depending on the previous history of synapses) persistent changes in synaptic weights during learning. A learning rule with activity-dependent induction threshold, known as the Bienenstock, Cooper, and Munro (BCM) theory, was introduced previously (Bienenstock et al., 1982). The BCM model is considered as a precursor form of metaplasticity and was shown to be consistent with STDP (Senn et al., 2001; Izhikevich and Desai, 2003; Pfister and Gerstner, 2006; Clopath et al., 2010). However, in these models, the activity-dependent threshold was introduced artificially or its postsynaptic dependency could not account for priming experiments reported in the hippocampus (Huang et al., 1992; Christie and Abraham, 1992; Mockett et al., 2002).

Thus, to our knowledge, no model has yet been proposed to account comprehensively for the multiple scales on which plasticity may occur and, in particular, the influence of previous presynaptic activity patterns. We introduce here an STDP model endowed with fast and slow induction processes that can, respectively, account for STDP, short-term competitive interactions (Sjöström et al., 2001; Wang et al., 2005) as well as long-term plasticity regulation as observed in hippocampal priming experiments. This metaplastic model of STDP, or mSTDP, provides a possible solution to the catastrophic-forgetting problem.

Materials and Methods

The mSTDP model

Biological basis. The mSTDP model is based on a number of experimental observations *in vitro*. First, important nonlinear interactions in STDP were reported when spike pairings occur closely in time or if triplet protocols are considered (Markram et al., 1997; Sjöström et al., 2001; Wang et al., 2005). Second, priming experiments were conducted in classical LTP/LTD experiments by stimulating the synaptic pathway to be conditioned before the

LTP/LTD induction protocol with a high-frequency tetanus or a low-frequency stimulation (Christie and Abraham, 1992; Huang et al., 1992; Wang et al., 1998; Mockett et al., 2002). These studies demonstrated that subthreshold preactivation of LTP or LTD mechanisms could impact their reactivation. Based on these results, it appears that plastic processes occurring at the same synapse through the same molecular pathways exhibit regulation on three distinct timescales that account, respectively, for the temporal and order specificity to induce synaptic weight changes (STDP rule), synaptic competition between paired and unpaired synapses (short-term interactions), and metaplasticity (long-term interactions). We propose here a unified phenomenological model in which the functional diversity of the synaptic effects naturally emerges across these timescales. This mSTDP model is based on putative biomolecular interactions between kinases and phosphatases acting on the synaptic weights and is inspired by experimental *in vitro* data on the effect of priming stimulations on LTP and LTD induction in the hippocampus.

The mSTDP algorithm. To implement this as a learning rule, LTP and LTD are treated separately such that the final synaptic weight change results from a linear summation of the scaling contribution produced by these two processes. Both LTP and LTD will be described through the action of two variables: a fast variable responsible for the direct action of these plasticity rules and a slow variable that will act as an induction threshold for authorizing the expression of plasticity. Both variables are computed on eligibility traces resulting from a triplet rule (Pfister and Gerstner, 2006).

For LTP, every presynaptic spike followed by a postsynaptic spike will increase the corresponding eligibility trace by an amount governed by the STDP exponential curve. This eligibility trace then decays with a slower time constant accounting for the triplet interaction. Both the fast and slow actions of LTP are then computed directly with this eligibility trace. The fast action will give rise to an instantaneous synaptic weight increase at the time of a postsynaptic spike, whereas the slow action that is acting at the same instant is a nonlinear function of the eligibility trace averaged over a long timescale and across synapses. The latter will act as a threshold for plasticity induction so that we only consider the positive amount of synaptic change resulting from the difference between fast and slow variables. This will be computed by considering a rectified function of their difference. The mechanism is identical for LTD except that the slow variable will be driven by the LTP eligibility trace accounting for the concurrent changes in LTP and LTD induction thresholds (Mockett et al., 2002). In the special case in which the slow variables are taken as constants, this model formally reduces to a learning rule similar to a model introduced previously by Senn et al. (2001) in which LTP and LTD are both subject to fixed induction thresholds. In this case, the model exhibits plasticity induction thresholds reminiscent of the Artola, Bröcher, and Singer (ABS) learning rule (Artola et al., 1990; Artola and Singer, 1993). In that study, the authors reported that the depolarization level of the postsynaptic neuron defines a plasticity threshold for LTD and a second threshold for LTP at higher depolarization levels. When the postsynaptic membrane potential was clamped at values under these thresholds, LTD and LTP were both absent. In the study by Senn et al. (2001), these thresholds are defined at the suprathreshold level but obey the same induction crossing order for LTD and LTP for increasing postsynaptic spiking activity.

For the LTP mechanism, and for a given synaptic connection, our model posits that every presynaptic spike elicits instantaneous incrementation of a variable r_{LTP}, followed by exponential decay with a time constant τ_{LTP}. This variable is used to detect pairing activity within the LTP side of the STDP curve:

$$\frac{dr_{LTP}}{dt} = -\frac{r_{LTP}}{\tau_{LTP}} + \sum_{t_{pre}} \delta(t - t_{pre}). \quad (1)$$

Once a postsynaptic spike is triggered, the current value of r_{LTP} is used to determine the instantaneous increment of the LTP eligibility trace δw_{LTP}, which decays with a time constant T_{LTP}, slower than the STDP time constant (τ_{LTP}):

$$\frac{d\delta w_{LTP}}{dt} = -\frac{\delta w_{LTP}}{T_{LTP}} + \sum_{t_{post}} r_{LTP}(t)\delta(t - t_{post}). \quad (2)$$

Note that, once Equation 1 is solved and used in Equation 2, the dynamics of δw_{LTP} depends on the correlation between presynaptic and postsynaptic spiking activity through the product of these point processes. The LTP eligibility trace δw_{LTP} is used to define two different actions on the synaptic weight. The instantaneous action of LTP takes place at the postsynaptic spike time and results in an increase given by the value of δw_{LTP} right after the spike. We chose a linear relationship for the sake of simplicity, but a monotonic nonlinear function involving more detailed biophysical processes could be used as well.

The second component introduced in our model describes the downregulation on the amount of LTP with a slow time constant and thus monitors the past history of a nonlinear function of δw_{LTP}. We hypothesize that this LTP downregulation term can be modeled as an exponential operator on δw_{LTP} averaged over a long time interval T and over a spatially distributed presynaptic ensemble, allowing some form of contextual regulation: $\delta w_{LTP}^{th}(t) = \left\langle \frac{\alpha_{LTP}}{T} \int_{-\infty}^{t} d\tau e^{(\tau-t)/T} e^{\beta \delta w_{LTP}(\tau)} \right\rangle_{pre}$. The coefficient α_{LTP} defines the impact of LTP downregulation relative to the potentiation term, and β determines the rate of modulation of the LTP downregulation. The average over a presynaptic ensemble determines the nature (including heterosynaptic effects) and the spatial extent of competition between presynaptic neurons. Unless stated otherwise, we will consider for simplicity an average over all presynaptic neurons such that the memory context is the same for all synapses, whether they are stimulated or not, and all synapses share the same plasticity thresholds (the slow variables δw_{LTP}^{th}). The slow variable is acting on δw_{LTP} right after each postsynaptic spike so that the complete synaptic increment at time t is given by the integration over times preceding t of the rectified difference between these two terms, such as

$$w_{LTP}(t) = \lim_{\epsilon \to 0^+} \lambda \int_{-\infty}^{t} ds \sum_{t_{post}} \delta(s - t_{post} - \epsilon) \left[\delta w_{LTP}(s) \right.$$

$$\left. - \left\langle \frac{\alpha_{LTP}}{T} \int_{-\infty}^{s} d\tau e^{(\tau-s)/T} e^{\beta \delta w_{LTP}(\tau)} \right\rangle_{pre} \right]_+$$

$$= \lim_{\epsilon \to 0^+} \lambda \int_{-\infty}^{t} ds \sum_{t_{post}} \delta(s - t_{post} - \epsilon)[\delta w_{LTP}(s) - \delta w_{LTP}^{th}(s)]_+, \quad (3)$$

where λ is the learning magnitude relative to the synaptic weight. In this equation and in the following, the product of two identical point processes for the postsynaptic or presynaptic activity is performed with an infinitesimal difference ϵ such that the value used for the weight update is taken immediately after the spike.

As mentioned above, the slow variable δw_{LTP}^{th} is used as an induction threshold by considering the net effect of LTP as the rectified difference between the fast and slow variables at each postsynaptic spike. In the absence of synaptic activity, the slow variable is equal to α_{LTP}, which is also the lower bound for this variable. As long as fluctuations of the fast variable can overcome the impact of the slow variable, LTP can occur. This induction threshold changes according to the past history of the synapses and can eventually prevent any additional potentiation.

For LTD, each postsynaptic spike followed by a presynaptic spike elicits behavior governed by a set of equations similar to Equations 1 and 2:

$$\frac{dr_{LTD}}{dt} = -\frac{r_{LTD}}{\tau_{LTD}} + \alpha \sum_{t_{post}} \delta(t - t_{post}). \quad (4)$$

$$\frac{d\delta w_{LTD}}{dt} = -\frac{\delta w_{LTD}}{T_{LTD}} + \sum_{t_{pre}} r_{LTD}(t)\delta(t - t_{pre}), \quad (5)$$

where α accounts for the asymmetric impact of the LTD and LTP sides of STDP. The LTD eligibility trace δw_{LTD} is hypothesized to have an instantaneous effect on AMPARs that will act to reduce the net synaptic weight of the synapse by δw_{LTD} right after the presynaptic neuron has fired. Moreover, the downregulation on the LTP variable is assumed to lower the threshold for LTD. This will be accounted for by introducing an additional term that is inversely proportional to the LTP downregulation term averaged over the past history and the same presynaptic ensemble (as shown by the sign change in the exponential term and its dependency on LTP): $\delta w_{LTD}^{th}(t) = \left\langle \dfrac{\alpha_{LTD}}{T} \int_{-\infty}^{t} d\tau e^{(\tau-t)/T} e^{-\beta \delta w_{LTP}(\tau)} \right\rangle_{pre}$. The coefficient α_{LTD} is the magnitude of facilitation of LTD that acts right after each presynaptic spike. Therefore, the net depressive effect on the synaptic weight at time t is given by the integration over times preceding t of the rectified difference between these two terms:

$$
w_{LTD}(t) = \lim_{\epsilon \to 0+} \lambda \int_{-\infty}^{t} ds \sum_{t_{pre}} \delta(s - t_{pre} - \epsilon) \left[\delta w_{LTD}(s) \right.
$$

$$
\left. - \left\langle \frac{\alpha_{LTD}}{T} \int_{-\infty}^{s} d\tau e^{(\tau-s)/T} e^{-\beta \delta w_{LTP}(\tau)} \right\rangle_{pre} \right]_{+}
$$

$$
= \lim_{\epsilon \to 0+} \lambda \int_{-\infty}^{t} ds \sum_{t_{pre}} \delta(s - t_{pre} - \epsilon) [\delta w_{LTD}(s) - \delta w_{LTD}^{th}(s)]_{+}.
\tag{6}
$$

The rectified difference creates plasticity induction thresholds for LTP and LTD that are essential for stabilization of individual synaptic weights. In the absence of synaptic activity, the resting value for the slow variable is α_{LTD}, which is also the upper bound for this variable. Similar models could be obtained with other positive supralinear functions. Altogether, the synaptic weight fluctuations are dictated by the difference between LTP and LTD:

$$
w(t) = w_{LTP}(t) - w_{LTD}(t).
\tag{7}
$$

We will assume lower bounds ($w_{min} = 0$ nS), and, in the case of Figure 9, we will introduce realistic upper hard bounds to avoid the local divergence of one or very few synaptic weights. This latter situation seems unrealistic because it would require all the impact of the distributed synaptic competition to be restricted to the reinforcement of a single synapse while forcing all other synapses to their lower bound value. Note that both the mSTDP algorithm and the BCM rule provide asymptotic stability of the mean of the synaptic weight distribution, as will be shown below. The further addition of hard bounds in the mSTDP model forbids the divergence of the weights of individual synapses because of the stochastic nature of competition in this latter model (which does not exist in rate-based models such as BCM).

A weight-dependent version of mSTDP could be obtained by adding this dependency in Equation 7, as described by van Rossum et al. (2000), Gütig et al. (2003), and Morrison et al. (2007), so that each LTP and LTD synaptic weight update is modulated by the current weight value.

From BCM plasticity to stochastic mSTDP

mSTDP without slow variables. Here, we formally describe the relationship between mSTDP and the BCM rule (Bienenstock et al., 1982) for a simple model in which two neurons are connected through a plastic synapse and follow independent Poisson processes. We will first consider the case in which the slow variables δw_{LTP}^{th} and δw_{LTD}^{th} have been set to a fixed value of 0 so that the rectification can be neglected. In this setting, the model reduces to a triplet rule defined by the fast variables δw_{LTP} and δw_{LTD} alone in Equations 3 and 6. The

computation of the mean synaptic variation as a function of presynaptic and postsynaptic rate is straightforward and similar to (Pfister and Gerstner, 2006)

$$
\frac{d\langle w(t) \rangle}{dt} = \rho_{pre}(\lambda \tau_{LTP} T_{LTP} \rho_{post}^2 + \lambda \rho_{post}(\tau_{LTP} - \alpha \tau_{LTD} T_{LTD} \rho_{pre}
$$

$$
- \alpha \tau_{LTD})), \tag{8}
$$

$$
= \rho_{pre} \Phi(\rho_{post}, \rho_{pre}), \tag{9}
$$

where ρ_{pre} and ρ_{post} are, respectively, the presynaptic and postsynaptic firing rates, and $\Phi(\rho_{post}, \rho_{pre})$ is a quadratic function of ρ_{post} that crosses the abscissa at 0 and another point given by

$$
\theta_{LTD}^0 = \frac{\alpha \tau_{LTD} T_{LTD}}{\tau_{LTP} T_{LTP}} \rho_{pre} + \frac{\alpha \tau_{LTD} - \tau_{LTP}}{\tau_{LTP} T_{LTP}}, \tag{10}
$$

which is positive if

$$
\frac{1}{T_{LTD}\rho_{pre} + 1} < \alpha \frac{\tau_{LTD}}{\tau_{LTP}}, \tag{11}
$$

a condition that is true for any ρ_{pre} if $\tau_{LTP} < \alpha \tau_{LTD}$. Otherwise, the existence of a second positive crossing point will depend on the presynaptic firing rate. Note that the dependence on the presynaptic firing rate was not present in the description of the threshold in the original BCM rule (Bienenstock et al., 1982). However, in the case in which T_{LTD} is very small or 0, the floating threshold dependence on presynaptic firing rate disappears in our model. This scenario is formally equivalent to the triplet rule found by Pfister and Gerstner (2006) after fitting their model to *in vitro* data (Sjöström et al., 2001; Wang et al., 2005).

mSTDP with fast and slow variables. One of the main ingredients of the BCM theory in providing metaplasticity is that the threshold between LTP and LTD for various postsynaptic firing rates can shift as a supralinear function of the past activity of the postsynaptic neuron. We show here that the slow variables linked to the LTP downregulation fulfill similarly this role. However, they do so without the constraint of a supralinear power relationship between the LTD–LTP transition threshold and the mean postsynaptic activity introduced in an ad hoc manner in the BCM model. The variation of the slow components in LTP and LTD can be considered constant during measurement of the BCM curve because the corresponding timescales are several orders of magnitude apart. Since we can compute at least numerically the expected asymptotic values of $\overline{\delta w_{LTP}^{th}}$ and $\overline{\delta w_{LTD}^{th}}$ for a given presynaptic and postsynaptic firing rate, we can derive the relationship between the BCM threshold and the postsynaptic firing averaged over the past history. We can thus write the mean synaptic weight derivative as

$$
\frac{d\langle w(t) \rangle}{dt} \simeq \lambda \left\langle \sum_{t_{post}} \delta(t - t_{post})[\delta w_{LTP}(t) - \overline{\delta w_{LTP}^{th}}]_{+} \right\rangle
$$

$$
- \lambda \left\langle \sum_{t_{pre}} \delta(t - t_{pre})[\delta w_{LTD}(t) - \overline{\delta w_{LTD}^{th}}]_{+} \right\rangle, \tag{12}
$$

where the ε term is omitted in the notation for the sake of simplification. The first averaged term in Equation 12 can be computed through the joint probability distribution of the postsynaptic spike times $\{t_{post}\}$ and the stochastic process $\delta w_{LTP}(t)$: $P(\delta w_{LTP}, t_{post})$. The average can then be performed on the rectified function using this joint distribution. Moreover, by using the conditional probability definition and the Poisson stationarity, we can write $P(\delta w_{LTP}, t_{post}) = P(\delta w_{LTP}|t_{post})P(t_{post}) = P(\delta w_{LTP}|t_{post})\rho_{ost}$. A similar expression can be obtained for LTD such that we can write Equation 12 as

Table 1. LTP and LTD asymptotic behaviors

	LTP	LTD
$\langle\delta w\rangle \ll \overline{\delta w^{\text{th}}}$	0	0
$\langle\delta w\rangle \sim \overline{\delta w^{\text{th}}}$	$\lambda\rho_{\text{post}}\left(\dfrac{1}{2}\langle\delta w_{\text{LTP}}\rangle_{t_{\text{post}}} - \dfrac{1}{2}\overline{\delta w_{\text{LTP}}^{\text{th}}} + \dfrac{\sigma(\delta w_{\text{LTP}})_{t_{\text{post}}}}{\sqrt{2\pi}}\right)$	$-\lambda\rho_{\text{pre}}\left(\dfrac{1}{2}\langle\delta w_{\text{LTD}}\rangle_{t_{\text{pre}}} - \dfrac{1}{2}\overline{\delta w_{\text{LTD}}^{\text{th}}} + \dfrac{\sigma(\delta w_{\text{LTD}})_{t_{\text{pre}}}}{\sqrt{2\pi}}\right)$
$\langle\delta w\rangle \gg \overline{\delta w^{\text{th}}}$	$\lambda\rho_{\text{post}}\left(\langle\delta w_{\text{LTP}}\rangle_{t_{\text{post}}} - \overline{\delta w_{\text{LTP}}^{\text{th}}}\right)$	$-\lambda\rho_{\text{pre}}\left(\langle\delta w_{\text{LTD}}\rangle_{t_{\text{pre}}} - \overline{\delta w_{\text{LTD}}^{\text{th}}}\right)$

The middle and last columns contain the asymptotic expression of LTP and LTD terms for the three possible regimes corresponding to the cases in which the fast variables are well below $\langle\delta w\rangle \ll \overline{\delta w^{\text{th}}}$, close to $\langle\delta w\rangle \sim \overline{\delta w^{\text{th}}}$, or well above $\langle\delta w\rangle \gg \overline{\delta w^{\text{th}}}$ the slow variables. Depending on the past activity of the synapse and on the parameter set, any combination of these LTP/LTD terms can occur and describe the current plasticity pattern.

$$\frac{d\langle w(t)\rangle}{dt} = \lambda\int_{\overline{\delta w_{\text{LTP}}^{\text{th}}}}^{\infty} d\delta w_{\text{LTP}}\rho_{\text{post}}P(\delta w_{\text{LTP}}|t_{\text{post}})(\delta w_{\text{LTP}} - \overline{\delta w_{\text{LTP}}^{\text{th}}})$$

$$- \lambda\int_{\overline{\delta w_{\text{LTD}}^{\text{th}}}}^{\infty} d\delta w_{\text{LTD}}\rho_{\text{pre}}P(\delta w_{\text{LTD}}|t_{\text{pre}})(\delta w_{\text{LTD}} - \overline{\delta w_{\text{LTD}}^{\text{th}}}). \quad (13)$$

We will consider three asymptotic cases for LTP and LTD to study the effect of the slow variables. In the linear regime ($\langle\delta w\rangle \gg \overline{\delta w^{\text{th}}}$), the threshold–linear function is approximated by a linear function so that a simple average can be performed in Equation 13. In the threshold regime ($\langle\delta w\rangle \ll \overline{\delta w^{\text{th}}}$), the linear–threshold function is almost always equal to 0 so that we can neglect the corresponding integrals in Equation 13. Finally, we will also consider the case in which the distribution is close to the threshold ($\langle\delta w\rangle \sim \overline{\delta w^{\text{th}}}$). In Equation 13, the term for LTP can be written as follows (an identical computation can be done for LTD):

$$\frac{d\langle w_{\text{LTP}}(t)\rangle}{dt} = \lambda\rho_{\text{post}}(\langle\delta w_{\text{LTP}}\rangle_{t_{\text{post}}} - \overline{\delta w_{\text{LTP}}^{\text{th}}} - F_{\text{LTP}}(\overline{\delta w_{\text{LTP}}^{\text{th}}})), \quad (14)$$

where

$$F_{\text{LTP}}(\overline{\delta w_{\text{LTP}}^{\text{th}}}) = \int_{-\infty}^{\overline{\delta w_{\text{LTP}}^{\text{th}}}} d\delta w_{\text{LTP}}P(\delta w_{\text{LTP}}|t_{\text{post}})(\delta w_{\text{LTP}} - \overline{\delta w_{\text{LTP}}^{\text{th}}}),$$

$$(15)$$

is a function of the LTP threshold and $\langle\ldots\rangle_{t_{\text{post}}}$ is a conditional average over the postsynaptic spike times. In this special case when the mean of the fast variable is assumed to be close to the threshold, $F_{\text{LTP}}(\overline{\delta w_{\text{LTP}}^{\text{th}}})$ can be expanded around the mean value of the distribution $P(\delta w_{\text{LTP}})$. We will assume that the scale of synaptic weight changes λ is very small compared with the postsynaptic and presynaptic mean firing rates so that the distributions $P(\delta w_{\text{LTP}}|t_{\text{post}})$ and $P(\delta w_{\text{LTD}}|t_{\text{pre}})$ in Equation 13 can be approximated by Gaussian probability distributions (diffusion approximation). The general form of such an expansion for a Gaussian distribution $P(x)$ of mean value m and variance σ^2 is given by

$$F(c) = \int_{-\infty}^{c} dxP(x)(x - c) = -\frac{\sigma}{\sqrt{2\pi}} - \frac{1}{2}(c - m)$$

$$+ \mathbb{O}((c - m)^2). \quad (16)$$

The LTP contribution can therefore be written as

$$\frac{d\langle w_{\text{LTP}}(t)\rangle}{dt} \simeq \lambda\rho_{\text{post}}\left(\frac{1}{2}\langle\delta w_{\text{LTP}}\rangle_{t_{\text{post}}} - \frac{1}{2}\overline{\delta w_{\text{LTP}}^{\text{th}}} + \frac{\sigma(\delta w_{\text{LTP}})_{t_{\text{post}}}}{\sqrt{2\pi}}\right).$$

$$(17)$$

From the definition of the fast variable in Equation 3, it can be shown that the last term grows as ρ_{post}^2 for Poisson processes (the variance of the corresponding compound shot noise process grows as ρ_{post}^4). Therefore, the overall form of the BCM curve is unchanged. Note, however, that this equation is only valid in the limit of the diffusion approximation. If the thresholds are set to 0, the mean should also be

close to 0 and the diffusion approximation is not valid anymore. These different asymptotic cases are summarized in Table 1. In particular, $\rho_{\text{post}}\langle\delta w_{\text{LTP}}\rangle_{t_{\text{post}}} = \tau_{\text{LTP}}T_{\text{LTP}}\rho_{\text{pre}}\rho_{\text{post}}^2 + \tau_{\text{LTP}}\rho_{\text{pre}}\rho_{\text{post}}$ and $\rho_{\text{pre}}\langle\delta w_{\text{LTD}}\rangle_{t_{\text{pre}}} = \alpha\tau_{\text{LTD}}T_{\text{LTD}}\rho_{\text{post}}\rho_{\text{pre}}^2 + \alpha\tau_{\text{LTD}}\rho_{\text{pre}}\rho_{\text{post}}$ such that $\lambda(\rho_{\text{post}}\langle\delta w_{\text{LTP}}\rangle_{t_{\text{post}}} - \rho_{\text{pre}}\langle\delta w_{\text{LTD}}\rangle_{t_{\text{pre}}}) = \rho_{\text{pre}}\Phi(\rho_{\text{post}}, \rho_{\text{pre}})$. This last relation corresponds to the curve of Equation 9 found in the case in which the slow variables are neglected. In the situation in which slow variables are taken into account, more complex behavior can be expected from the synaptic weights depending on the definition of the presynaptic ensemble sharing the same slow variables, and this can be explored by introducing a new function $\Phi_1(\rho_{\text{post}}, \rho_{\text{pre}})$. This function is derived from the reference BCM curve (case without slow variables in our model) with a shift in the floating threshold proportional to $\dfrac{\overline{\delta w_{\text{LTP}}^{\text{th}}}}{\rho_{\text{pre}}}$, where $\overline{\delta w_{\text{LTP}}^{\text{th}}}$ is an exponential function of the presynaptic activity averaged over time and over the presynaptic ensemble. This specific presynaptic dependence will be responsible for the various forms of synaptic competition observed in our simulations. In particular, when all synapses share the same slow variables (identical $\overline{\delta w_{\text{LTP}}^{\text{th}}}$ for all synapses), the ones that are targeted by presynaptic neurons with higher firing rates tend to express a lower LTD–LTP transition threshold. Conversely, when thresholds are assumed to be independent across synapses, a high presynaptic firing rate results in higher floating threshold because of the exponential term. This will be illustrated in the following simulations.

Because the LTP and LTD thresholds depend on the past history of the synapse and on the coefficients α_{LTP} and α_{LTD}, any combination of the terms in Table 1 can occur during the learning process. However, from the definition of these thresholds as exponential functions of the postsynaptic and presynaptic activities (Eqs. 3, 6), it can be shown that the BCM sliding threshold always increases with the postsynaptic and presynaptic firing rate past histories and that $\rho_{\text{post}} = 0$ is always a crossing point. In particular, because of the induction thresholds, there is a postsynaptic domain including 0 (null activity level) where no plasticity can take place, which corresponds to rare coincidences between presynaptic and postsynaptic spiking activities: the "zero domain." Note that this plateau domain differs from the origin singularity at 0 in the BCM model, because it is not limited to the 0 point but encompasses a postsynaptic activity domain that is similar to the dead zone assumed in the ABS model. Beyond this stable domain, three regimes can be outlined for increasing past synaptic activity. In the pure LTP regime, the BCM curve is entirely positive and only potentiation can occur. This scenario can happen when the past synaptic activity was very low and the corresponding slow variables could not shift the curve downward. For higher past synaptic activity, the slow variables $\delta w_{\text{LTP}}^{\text{th}}$ and $\delta w_{\text{LTD}}^{\text{th}}$ have reached values high enough to define a non-null BCM-like floating threshold. In this particular regime, LTD or LTP will be observed depending on the current level of postsynaptic firing rate. Finally, in the pure LTD regime, the LTP threshold has completely overtaken the fast variable so that the BCM sliding threshold has increased to infinity and only depression can occur. In this last situation, the zero domain will shrink to the single point at 0 because of the decreasing LTD threshold. Altogether, these terms produce, for increasing past postsynaptic activity, the desired supralinear shift of the BCM sliding threshold toward higher values. This behavior is further illustrated in Figure 4.

Neuron model

We consider leaky conductance-based integrate-and-fire neurons, with membrane time constant $\tau_m = 20$ ms, and resting membrane potential $V_{rest} = -70$ mV. When the membrane potential V reaches the spiking threshold $V_{thresh} = -54$ mV, a spike is generated and the membrane potential is held at the resting potential for a refractory period of duration $\tau_{ref} = 5$ ms. Synaptic connections are modeled as conductance changes:

$$\tau_m \frac{dV(t)}{dt} = (V_{rest} - V(t)) + g_{exc}(t)(E_{exc} - V(t))$$
$$+ g_{inh}(t)(E_{inh} - V(t)), \quad (18)$$

where the reversal potentials are $E_{exc} = 0$ mV and $E_{inh} = -70$ mV. The synaptic activation is modeled as an instantaneous conductance increase followed by exponential decay:

$$\tau_{exc} \frac{dg_{exc}(t)}{dt} = -g_{exc}(t) + S_{exc}(t), \quad (19)$$

$$\tau_{inh} \frac{dg_{inh}(t)}{dt} = -g_{inh}(t) + S_{inh}(t), \quad (20)$$

with time constants $\tau_{exc} = 5$ ms and $\tau_{inh} = 5$ ms. $S_{exc/inh}(t)$ are the synaptic spike trains —point processes—coming from the excitatory and inhibitory populations, respectively.

Constraining the model from in vitro *data*

The data of Sjöström et al. (2001) and Wang et al. (2005) shown in Figure 2 were used to constrain the model with a least-squares fit, similar to the one described by Pfister and Gerstner (2006). The time constants for the STDP time windows were fixed to $\tau_{LTP} = 20$ ms and $\tau_{LTD} = 25$ ms, and parameters found by the model are $T_{LTP} = 845$ ms, $T_{LTD} = 995$ ms, and $\alpha = 0.46$. In all the other simulations, the two parameters T_{LTP} and T_{LTD} were set to 1 s.

Priming simulations

For the LTP/LTD priming protocols (see Fig. 3) corresponding to Huang et al. (1992) and Mockett et al. (2002), respectively, we used a two-neuron model. Synaptic weights are chosen such that weak presynaptic tetanic bursts are able to elicit some spikes in the postsynaptic neuron ($w_{init} = 50$ nS and $w_{max} = 100$ nS). The strong values of those weights could be reduced if more than two neurons were considered. Parameters for the plasticity rule were $\beta = 0.15$, $\alpha_{LTP} = 15$, and $\alpha_{LTD} = 2$. For LTP priming, the presynaptic neuron received 10 weak tetanic bursts (20 spikes/s for 250 ms each), with an interburst interval of 1 s. Then, 10 s later, it received a strong burst at 100 spikes/s for 200 ms. The learning rate was fixed to $\lambda = 5 \times 10^{-3}$ and $T = 50$ s. For LTD priming, the presynaptic neuron receives a low-frequency stimulation at 15 spikes/s for 40 s. Then, 40 s later, the exact same stimulation occurs. The learning rate was fixed to $\lambda = 0.8 \times 10^{-4}$, and $T = 100s$.

Feedforward convergence simulations

In feedforward convergence simulations (see Figs. 4–9), the network was composed of 1000 excitatory and 250 inhibitory neurons connected to a postsynaptic neuron. Only excitatory connections are subject to the mSTDP learning rule, and parameters for the plasticity were $\beta = 0.15$, $\alpha_{LTP} = 2.5$, $\alpha_{LTD} = 2.3$, and $T = 5$ s. The initial weight for the presynaptic population was $w_{init} = 0.3$ nS. The ratio between excitatory and inhibitory weights is such that $w_{inh} = 33 \times w_{exc}$. Every neuron in the presynaptic population had a constant firing rate of 8 spikes/s, and the learning rate was fixed to $\lambda = 2.5 \times 10^{-6}$. For Figure 9, presynaptic neurons are stimulated with wrapped Gaussian profiles of rates $v_i = 1 + 50 \times \exp(-(i - \mu)^2/2\sigma^2)$ spikes/s, the center μ being shifted randomly every 50 ms over all possible positions $i \in \{1, ..., 1000\}$ and with $\sigma = 80$. Here an upper bound for the synaptic weights was set to $w_{max} = 0.5$ nS, and the learning rate was divided by 2 to avoid strong fluctuations near the boundaries.

Network model

The network was composed of $N_{exc} = 63 \times 63$ excitatory and $N_{inh} = 31 \times 31$ inhibitory neurons, arranged on a two-dimensional grid of artificial size 1 mm². Because of the rather small neuron density, the metric unit used to describe the network can be considered arbitrary, but we will assume, for the sake of comparability between stimulation protocols and for establishing realistic propagation delays, that 1 mm corresponds to ~1° of visual angle in cat primary visual cortex V1. The grid has periodic boundary conditions to avoid any border effects. Each neuron was sparsely connected with the rest of the network with a connection probability that depended on the distance Δ_{ij} between neurons according to a Gaussian profile $p_{ij} = e^{-\frac{\Delta_{ij}^2}{2\sigma^2}}$ with a spatial spread of $\sigma = 0.1$ mm. For each neuron, $K = \varepsilon(N_{exc} + N_{inh})$ incoming connections were drawn by randomly picking other neurons in the network that will or will not create a projection according to a rejection method based on the Gaussian profile. The connection density was $\varepsilon = 0.1$, and self-connections were discarded.

The network was set to an asynchronous irregular (AI) state (Brunel, 2000) with a population rate of ~7 spikes/s and with the mean coefficient of variation of the interspike interval (ISI CV) of 1. Initial synaptic weights were drawn from Gaussian distributions with mean values: $\delta g_{exc} = 5$ nS and $\delta g_{inh} = 50$ nS. The SDs of the Gaussian distributions are one-third of their mean values. Every neuron received an additional Poisson input at 500 spikes/s in Figure 10 and 300 spikes/s in Figure 11. To reproduce lateral propagation along horizontal connections, as supported by *in vivo* intracellular evidence (Bringuier et al., 1999) and voltage-sensitive dye imaging (Benucci et al., 2007), synaptic delays are considered as being linearly dependent on distance, i.e., if Δ_{ij} is the distance between two neurons, we have $d_{ij} = d_{syn} + \frac{\Delta_{ij}}{v}$, where d_{ij} is the delay between the two neurons, $d_{syn} = 0.2$ ms is the minimal delay attributable to synaptic transmission, and v is the velocity of dendritic conduction. In all simulations, $v = 0.2$ mm/ms, in agreement with biological measurements (Bringuier et al., 1999), i.e., delays in the network are in the range $\left[0.2 \text{ ms}, 0.2 + \frac{\sqrt{2}}{2v} \simeq 3.7 \text{ ms} \right]$, where $\sqrt{2}/2$ mm is the longest possible connection in the network.

In all simulations, only recurrent excitatory connections are subject to plasticity, and the learning scale is set to $\lambda = 1.5 \times 10^{-6}$. For Figure 11, an external layer of stimulating neurons driven by Poisson processes is added on top of the recurrent network. All external neurons project locally to all neurons of the recurrent network within a radius of 0.3 mm, with synaptic weights $g_{exc} = 5$ nS. In this external layer, moving bars were modeled as linear wave fronts with a profile described by a half-Gaussian function. The bars were 0.2 mm large and 0.1 mm long, with a peak rate amplitude of 4 spikes/s. This pattern of activity originated every 100 ms at random places and propagated perpendicularly to the wave front in a random direction ϕ, with a constant speed $s = 10$ mm/s (that would correspond to 10°/s in the visual field according to the chosen metric to which cells in V1, and mostly secondary visual cortex V2, will respond). In the unbiased case, values of ϕ are drawn from a uniform distribution ranging from 0 to 2π, whereas in the biased case, they are drawn from a wrapped Gaussian distribution centered on 0 with variance $\sigma^2 = \pi/4$.

For the quantification of the preferred direction, we considered, for each neuron, the centered map of its afferent presynaptic weights. In this neuron-centered reference frame, an incoming input originates from a position identified by its polar coordinates (r, θ), such that the map value at (r, θ) corresponds to the weight w of the projection. Angular positions of all the afferent connections were then gathered and weighted by their corresponding synaptic weights. The peak value of this weighted histogram is used to define the structural direction preference of the neuron, i.e., the angular direction along which the composite synaptic contribution of all synaptic weights recruited by the bar is the strongest.

Simulator

All simulations were performed using a modified version of the NEST 2.0 simulator (Gewaltig and Diesmann, 2007), allowing updates of the synaptic weights at each time step of the simulation, controlled with the

PyNN interface (Davison et al., 2008). The integration time step of our simulations was 0.1 ms.

Results

Biophysical correlates of the mSTDP model

The working assumption of our mSTDP model is based on electrophysiological *in vitro* results reported previously (Wang et al., 2005), showing evidence that LTP and LTD are elicited by calcium influx through distinct dedicated channels, NMDA and L-type, respectively. For LTP, and for every pre–post pairing, calcium enters the cell through NMDA channels and activates second-messenger-dependent molecular processes leading to potentiation (presumably through calmodulin binding and subsequent activation of kinases; see Lisman, 1989, 1994). For LTP, we consider that the STDP window is defined by an exponential curve that dictates the instantaneous amount of Ca^{2+}/calmodulin that will be formed at the time of the postsynaptic spike. In a recent work by Faas et al. (2011), the authors have shown that the binding of calmodulin with calcium takes place over a very short timescale, in accordance with our phenomenological description of this process as an instantaneous rise of the corresponding variable at the time of postsynaptic spikes. This variable will be labeled δw_{LTP}. The action of this variable is twofold. On a very short timescale, the corresponding molecule will activate a kinase K that will phosphorylate AMPA receptors, thus increasing the excitatory synaptic conductance, acting as a phenomenological substrate for LTP, as in previous biophysical models (Senn et al., 2001; Shouval et al., 2002; Badoual et al., 2006; Pfister and Gerstner, 2006; Graupner and Brunel, 2007; Clopath et al., 2010). On a slower timescale, it will also activate another pathway, resulting in a nonlinear downregulation of the kinase activation. This additional hypothetical slow molecular pathway is critical in this model and has not been taken into account in previous models of calcium-dependent kinase kinetics that usually rely on positive regulation of the kinase (Lisman, 1989, 1994; Okamoto and Ichikawa, 1994; Lisman et al., 2002; Graupner and Brunel, 2007).

In our phenomenological model, these fast and slow actions were implemented through two different operators. The rapid Ca^{2+}/calmodulin action is modeled as an instantaneous update of synaptic weight given by the value of δw_{LTP} at the time of the postsynaptic spike. The long-lasting LTP downregulation acts as an induction threshold during these postsynaptic events and will be modeled as an average of $\exp(\beta \delta w_{LTP})$ over the past history and over a presynaptic ensemble. This function accounts for the strong nonlinear increase in the LTP threshold when weak tetanic priming stimulations are used before a strong tetanic induction in hippocampal slices (Huang et al., 1992). Unless stated otherwise, the presynaptic ensemble will be averaged over all previously active synapses projecting on the postsynaptic neuron such that they all share the same slow variables. Note also that the fast variable does not need to exhibit an instantaneous rise but can also have a finite rise time as long as the timescale is significantly faster than that of the slow variables. The net effect on the synaptic weight is proportional to the kinase concentration, thus resulting in the rectified difference between these two terms (see Materials and Methods). This LTP mechanism is depicted in Figure 1*A* and obeys Equation 3 for instantaneous weight changes induced by LTP: w_{LTP}.

Note that other types of ensemble average can be considered for the slow variables. For instance, the sharing of the slow variables by synapses can be done on a local basis (restricted to a specific domain of the neurite) and not on a global basis as described above. This could correspond to local groups of synapses located in the same dendritic branch and sharing some of the proteins through local diffusion. This will result in different learning rules, as described in later sections.

Regarding LTD, it has been shown that priming stimulations can facilitate LTD induction with low-frequency stimulation (Christie and Abraham, 1992; Wang et al., 1998; Mockett et al., 2002). Although the direct action of LTD is assumed to be conveyed by calcium influx through dedicated L-type channels (Graef et al., 1999; Wang et al., 2005), the priming effect has been shown to occur through NMDA channels and to covary with the metaplasticity observed for LTP (Mockett et al., 2002). Indeed, LTD facilitation and LTP suppression are both elicited in priming experiments that activate NMDA receptors. We assumed that the direct action of "pre after post" pairing results in the binding of an independent Ca^{2+}/calmodulin (variable δw_{LTD}) to a phosphatase that will be responsible for rapid AMPA receptor dephosphorylation and endocytosis. Moreover, we hypothesized that this process can be facilitated by the Ca^{2+}/calmodulin-dependent LTP downregulation generated through NMDA channels that will make more Ca^{2+}/calmodulin molecules available for the phosphatase. This is modeled as a buffering term acting at the presynaptic spike times and that is inversely proportional to the slow LTP action $\exp(\beta \delta w_{LTP})$, averaged over the past history and over the same presynaptic ensemble. This LTD mechanism is depicted in Figure 1*B* and obeys Equation 6 for instantaneous weight changes induced by LTD: w_{LTD}. The time evolution of the synaptic weight is given by the time integral over the synaptic changes (Eq. 7).

Thus, one of the main consequences of these mechanisms is that the LTP or LTD changes to be expressed will have to cross a plasticity induction threshold. This by itself is similar to a previously proposed plasticity model (Senn et al., 2001), except that the threshold here is moving as a function of the past activity attributable to the dual timescale interactions that we propose to underlie metaplasticity. Changes in thresholds in this model result from short-term upregulations and downregulations of purely presynaptic vesicle release probability at individual synapses through nonlinear messenger interactions. In contrast, the mSTDP model offers the possibility of studying several scenarios, ranging from learning rules with a single shared postsynaptic floating threshold (as in BCM) to multiple and purely presynaptic plasticity induction thresholds (as in the study by Senn et al., 2001).

The mSTDP model can be more readily compared with the BCM hypothesis, because the two learning rules produce a floating plasticity threshold driven, respectively, by presynaptic and postsynaptic activity (mSTDP) or shared by all the synapses (BCM) and evolving on a very slow timescale. Note that these two features are absent in the model of Senn et al. (2001), in which the sliding feature of the threshold arises from kinetic equations regulating the vesicle release probability. However, this adaptive threshold is driven by the same time constant causing the synaptic weight changes. Therefore, this model does not account for the multiple plasticity timescales necessary to reproduce *in vitro* priming experiments, neither does this rule incorporate plasticity induction thresholds shared among all synapses critical to ensure synaptic competition in the BCM framework. More recently, models based on triplet interactions were introduced that reproduce some features of the BCM rule (Pfister and Gerstner, 2006; Clopath et al., 2010). However, the slow sliding threshold in these models was introduced as an ad hoc homeostatic process governing the amplitude of depression (Clopath et al., 2010) or potentiation (Pfister and Gerstner, 2006), without a structural relationship with the plasticity algorithm. The induction thresholds for LTP and LTD are therefore kept fixed in these models. We will show further in the text that our biochemical scenario can
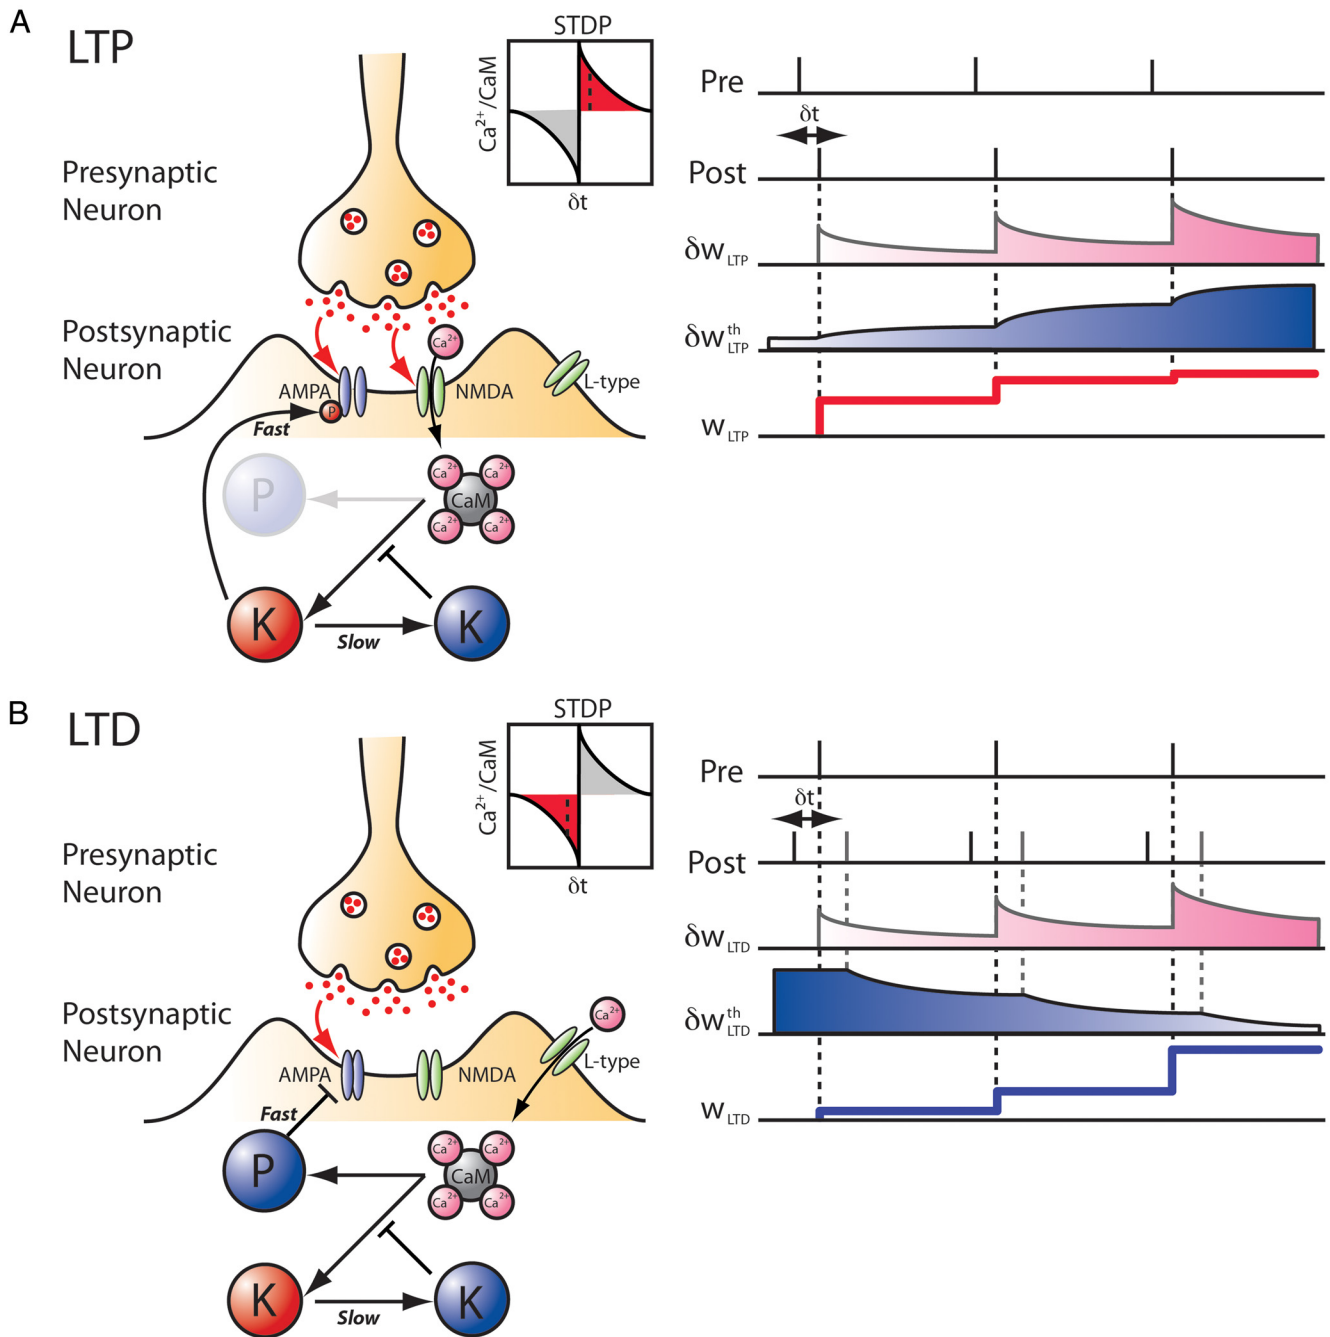
Figure 1. Schematic description of the model of mSTDP. ***A***, Mechanism of LTP with a fast potentiation action and slower downregulation of the instantaneous amount of potentiation at the postsynaptic spike times. The plasticity rule used for LTP is inspired by putative molecular pathways depicted on the left. A pairing protocol between presynaptic and postsynaptic spike trains is illustrated on the right with the time evolution of each variable of the model. The slowly increasing threshold (δw^{th}_{LTP}, blue curve) progressively reduces the effect of fast potentiation (δw_{LTP}, pink curve) on the synaptic weight (w_{LTP}, thick red lines) at subsequent postsynaptic spike times. The pairing time difference δt is relative to the presynaptic spike time ($\delta t = t_{post} - t_{pre}$). ***B***, Mechanism of LTD with a fast depression action and a facilitation produced by LTP downregulation at the presynaptic spike times. The plasticity rule used for LTD is inspired by putative molecular pathways depicted on the left. On the right, a triplet protocol illustrates how the progressive lowering of the LTD threshold (δw^{th}_{LTD}, blue curve) produced by repeated pre–post pairings can increase the functional impact of fast depression (δw_{LTD}, pink curve) produced by previous post–pre pairings, resulting in larger depression of the synaptic weight (w_{LTD}, thick blue lines) at the presynaptic spike times.

account for the classic "floating plasticity threshold" found in previous experimental and theoretical BCM-like studies of synaptic plasticity, but in this case, the "floating feature" is an emergent property of the model.

Single-cell STDP experiments

Before considering the mSTDP with the fast and slow variable actions, we first compared our model with plasticity protocols

occurring at a fast timescale so that the slow variables can be considered constant. In particular, to determine whether this phenomenological model is consistent with STDP experiments, we show that its predictions reproduce the plasticity effects reported for first- and second-order interspike interactions. In Figure 2 *A*, the biexponential STDP curve was obtained using a pairing protocol at 1 Hz. Triplet interactions were obtained in Figure 2, *B* and *C*, for a given set of parameters chosen to
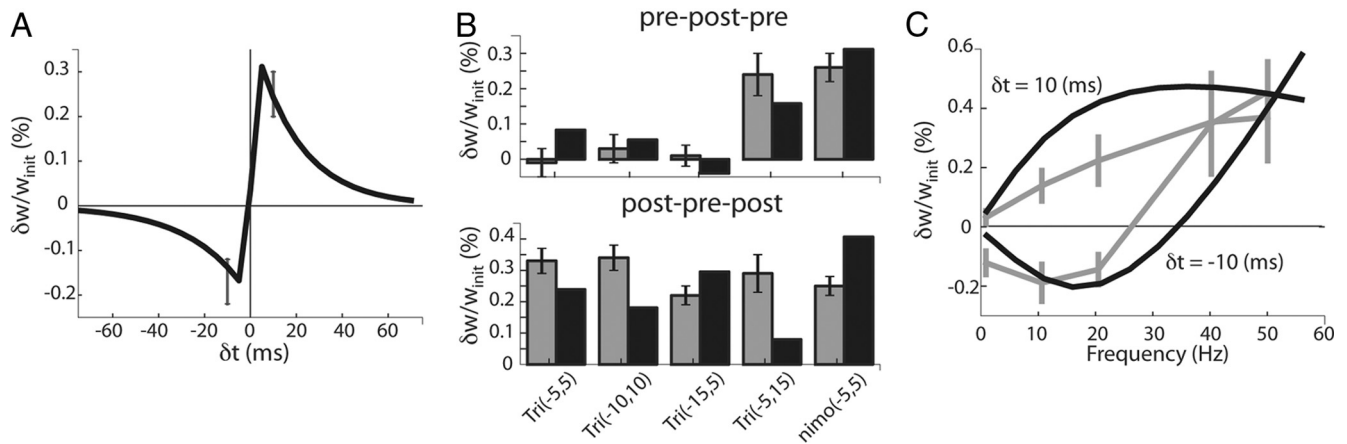
Figure 2. Reproduction of STDP experiments *in vitro*. Data shown in gray have been used to fit the mSTDP model parameters for STDP and short-term interactions (with the slow variables set to 0). ***A***, Typical biexponential STDP curve, linking the relative amplitude of the synaptic change (positive for LTP, negative for LTD) with the $t_{post} - t_{pre}$ interval and obtained with a simulated pairing protocol. Gray error bars in the graph are taken from Wang et al. (2005) recorded in hippocampal neurons for $\delta t = \pm 10$ ms. ***B***, Relative synaptic changes induced by triplet protocols (denoted Tri) compared with experimental data obtained by Wang et al. (2005) (top, pre–post–pre; bottom, post–pre–post). Gray bars correspond to the original recordings for each protocol in the abscissa (with time intervals between each spike). Black bars correspond to the prediction of the model. The last set of columns represents the condition in which L-type calcium channels are blocked with nimo, a specific L-type calcium channel antagonist. ***C***, STDP dependence on the pairing frequency for $\delta t = \pm 10$ ms. In gray are original data obtained for visual cortex neurons by Sjöström et al. (2001) for this protocol.

qualitatively match the data reported by Sjöström et al. (2001) and Wang et al. (2005) and with the slow variables set to 0 for the sake of simplicity. Note that these results are conserved even in the presence of slow variables (data not shown). Previous studied plasticity models have also succeeded in reproducing these data (Pfister and Gerstner, 2006; Clopath et al., 2010) or other similar datasets (Senn et al., 2001) with generalized short-term interactions. In the study by Senn et al. (2001), the short-term interaction is produced by regulation of presynaptic vesicle release probability. Here, the short-term interaction is inherited from the assumption that the Ca^{2+}/calmodulin-bound concentration decays slowly compared with the STDP time constant. Figure 2*B* displays the synaptic changes for different patterns of presynaptic and postsynaptic spike triplets (Wang et al., 2005). In pre–post–pre firing patterns, for instance, the synapse displays virtually no plasticity, whereas significantly larger potentiation is obtained for post–pre–post patterns that would produce comparable synaptic changes if short-term interactions were neglected. Moreover, when simulating L-type channel blockade ("nimo" application condition) by forcing δw_{LTD} to 0, we obtained a comparable potentiation between pre-post-pre and post-pre-post patterns, a prediction in accordance with experimental data. In Figure 2*C*, plasticity curves are plotted for fixed pairing intervals ($\delta t = \pm 10$ ms) occurring at various frequencies. In particular, for post–pre pairings, the synapse first experiences depression followed by potentiation at higher frequencies, as reported by Sjöström et al. (2001). The main features of the model responsible for this frequency-dependent effect are the extra time constants added by the binding of Ca^{2+} with independent calmodulin molecules (variables $\delta w_{LTP/LTD}$). These time constants, similar to those of the triplet model by Pfister and Gerstner (2006), are able of capturing not only pairwise interactions within pre–post pairings but also higher-order interactions. Note that these high-order plasticity rules were also shown to naturally emerge from the dynamics of kinase/phosphatase interaction with calcium and glutamate in a previous biophysical model (Badoual et al., 2006).
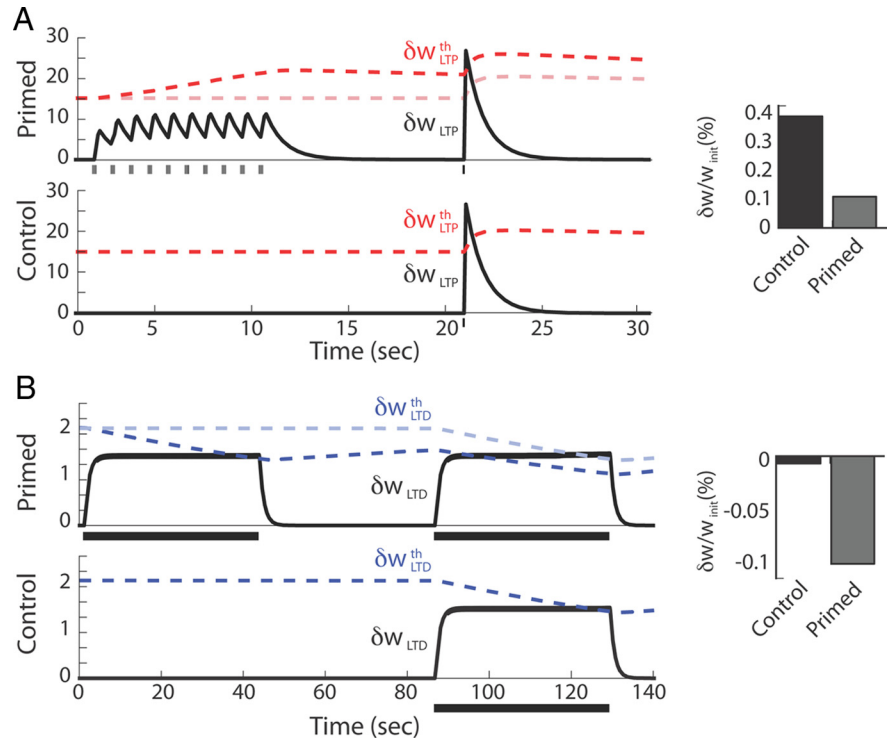
Priming experiments

We next considered experimental data involving timescales in which the action of the mSTDP slow variables is critical. To confirm that the mSTDP model faithfully reproduces the dynamics of metaplasticity as reported *in vitro*, we replicated priming experiments that were shown to either inhibit LTP or facilitate LTD. In Figure 3*A*, the protocol used by Huang et al. (1992) was reproduced in a two-neuron model. When weak tetanus priming stimulations were used before subsequent strong tetanus stimulation, the resulting synaptic weight change was considerably lower than the change elicited by the strong tetanus stimulation alone. Preactivation of LTP increases its induction threshold and then decreases the amount of LTP obtained during the strong tetanus stimulation. Conversely, LTD induction by low-frequency stimulation was significantly facilitated when a previous low-frequency priming stimulation was used (Fig. 3*B*). This result is in accordance with LTD facilitation reported by Christie and Abraham (1992), Wang et al. (1998), and Mockett et al. (2002). Low-frequency stimulation of the LTD pathway decreases the induction threshold for LTD and therefore increases the amount of depression compared with a control situation without priming. Moreover, in the study by Mockett et al. (2002), it was shown that this facilitation occurs through NMDA channels and is concurrent with LTP downregulation, which is in accordance with our definition of the slow variables both depending on δw_{LTP}. Note that the results shown in Figure 3 were obtained with a set of parameters different from that used for the rest of the study to mimic the protocols used *in vitro*. To avoid excessively long simulations, time constants smaller than those observed experimentally were used in the rest of the study. Nevertheless, similar results on priming experiments can be obtained with our choice of parameters if the timescale and burst amplitude of each protocol are properly scaled.

Relation with BCM

The dependence of LTP and LTD induction thresholds on the past synaptic activity in mSTDP introduces new interesting functional properties that can ensure stable learning as well as synap-

tic competition. We thus further studied the relationship between mSTDP and the original BCM plasticity rule first introduced to endow plasticity with stable convergence through synaptic competition (Bienenstock et al., 1982). Previous theoretical studies have succeeded in reproducing some aspects of the BCM rule by considering nonlinear interactions between presynaptic and postsynaptic spike patterns in STDP (Senn et al., 2001; Izhikevich and Desai, 2003; Burkitt et al., 2004; Pfister and Gerstner, 2006; Clopath et al., 2010). This correspondence was achieved either by assuming first-neighbor interactions between spikes (Izhikevich and Desai, 2003; Burkitt et al., 2004) or short-term interactions involving slower variables (Senn et al., 2001; Pfister and Gerstner, 2006; Clopath et al., 2010). However, these models were unable to account for one salient feature of the BCM rule: the floating plasticity threshold shared among all synapses integrating the impact of mean postsynaptic activity changes on slow timescales, which was either included as an ad hoc mechanism or was absent (Bush et al., 2010). The sliding plasticity threshold was originally assumed to be a supralinear function of the past postsynaptic firing rate history and was introduced to ensure nontrivial convergence and to avoid weight divergence in the plasticity algorithm (Bienenstock et al., 1982). In Figure 4A, we first

Figure 3. Reproduction of priming experiments using the mSTDP rule. A, Simulation of an LTP priming experiment inspired by Huang et al. (1992), in which weak tetanus stimulation can elicit LTP suppression. The top shows the time evolution of fast and slow LTP variables with a priming stimulation, and the bottom shows the control condition without priming (with δw^{th}_{LTP} also illustrated in light red in the top). The postsynaptic spike times corresponding to the weight updates are depicted below each graph. The histograms on the right are the resulting synaptic weight changes relative to their initial value for the two conditions. B, Same as A but with a protocol inspired by Mockett et al. (2002) in which previous low-frequency stimulation can elicit LTD facilitation for subsequent low-frequency stimulations. Slow variables are drawn in blue. The presynaptic spike times corresponding to the weight updates are depicted below each graph.

plotted the amplitude and sign of plasticity for various presynaptic and postsynaptic firing rates in a model consisting of a pair of neurons linked by an ideal synapse and firing as Poisson processes. Figure 4B shows two BCM curves that correspond to selected lines in the diagram of Figure 4A. These curves were computed in the absence of slow variables and are compared with theoretical predictions (Eq. 9). The depression domain and thus the BCM threshold increases with presynaptic firing rate as expected from the model definition (see Materials and Methods). This dependence is absent from the original BCM rule and creates differential plasticity changes between synapses originating from neurons with different firing rates as reported previously (Burkitt et al., 2004). This form of synaptic competition usually requires fine tuning of the model parameters and produces unstable distributions of synaptic weights because no homeostatic regulation is present in the absence of slow variables to ensure a normalization of the overall synaptic weight impact on the postsynaptic neuron. Moreover, in the parameter regime of the triplet rule used in Figure 4A, regions of anti-Hebbian learning exist in which an increase in the presynaptic firing rate results in a decrease of the corresponding synaptic weight. When the value of T_{LTD} is decreased, the presynaptic dependence progressively disappears until a fixed BCM threshold is observed independently of any change in the presynaptic rate when $T_{LTD} = 0$. In Figure 4C, we illustrate several plasticity rules for different values of the triplet parameters T_{LTD} and α. In particular, for $T_{LTD} = 0$ and $\alpha = 0.46$, only LTP can occur. The choice of parameters in the last panel is formally similar to that found by Pfister and

Gerstner (2006) when fitting their triplet model to the data of Sjöström et al. (2001) and Wang et al. (2005); in their studies, as well as in this particular parametric case of our mSTDP model, the BCM threshold is constant and shared by all synapses.

We then studied the role of the slow variables δw^{th}_{LTP} and δw^{th}_{LTD} in modulating the LTD–LTP transition plasticity BCM threshold with the past postsynaptic firing rate history, as originally used by Bienenstock and colleagues to simulate the effect of rearing conditions (dark rearing vs normal visual experience) on visual cortical plasticity. Measuring this relationship directly from the model requires the decoupling of slow and fast variables such that the past history of the postsynaptic activity is not affected during the measurement of the "instantaneous" BCM threshold. Therefore, we assumed that the slow time constant T is slow enough compared with other time constants to be considered infinite during the BCM curve measurement. Moreover, we also assumed that the presynaptic average is carried over all synapses. In this asymptotic regime ($T \to \infty$), all synapses share common LTP and LTD thresholds described by the coefficients $\overline{\delta w^{th}_{LTP}}$ and $\overline{\delta w^{th}_{LTD}}$ of Equations 3 and 6 that act as constant thresholds defined by the past activity regime. The new BCM plasticity curves and thresholds can thus be computed for any pair of fixed LTP and LTD thresholds. Figure 4D shows the plasticity thresholds obtained for different values of these coefficients. Note that the existence of an LTD induction threshold was absent from the original BCM rule (Bienenstock et al., 1982) but was introduced at the subthreshold level in the ABS rule (Artola and Singer, 1993) and at the suprathreshold level by Senn et al. (2001). However, in
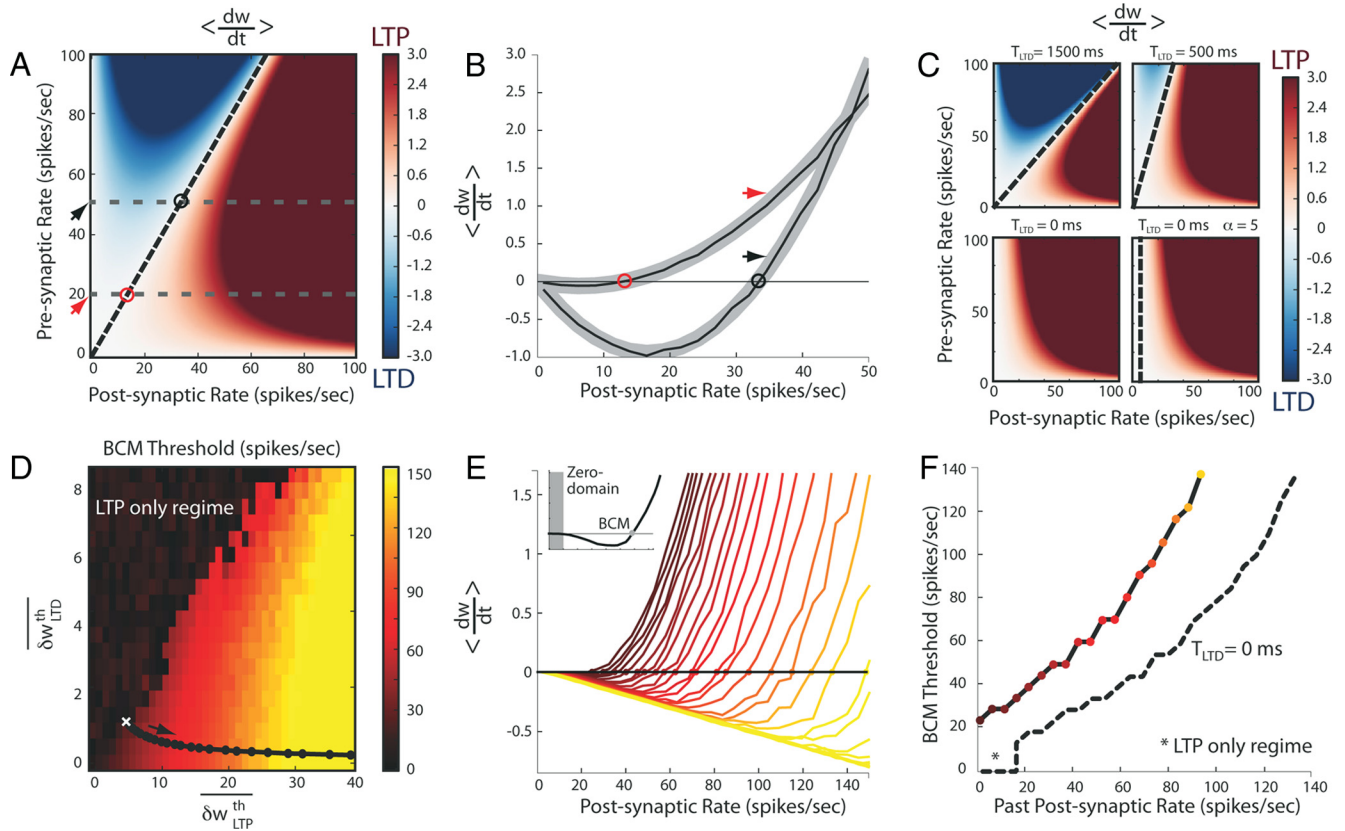
Figure 4. Relationship between mSTDP and the BCM plasticity rule. **A**, Mean synaptic change $\langle dw/dt \rangle$ (LTP in red, LTD in blue) as a function of presynaptic and postsynaptic firing rates in a two-neuron model with Poisson-process firing statistics and in the absence of slow variables. The oblique dotted line represents the crossing point of the plasticity curve between LTD and LTP. **B**, Two BCM-like plasticity curves, linking synaptic change with mean postsynaptic activity, correspond to the gray dashed lines in **A**, respectively, for a presynaptic firing rate of 20 spikes/s (red arrow) and 50 spikes/s (black arrow), with the LTD–LTP transition threshold marked by an open circle. Theoretical predictions are shown in bold gray. **C**, Same as in **A** for different T_{LTD} and α values. Note that α has been set to 0.46 in the first three panels and $\alpha = 5.00$ in the last panel. **D**, Values of the LTD–LTP transition threshold (BCM) are plotted as a function of asymptotic (constant) slow variables $\overline{\delta w_{LTP}^{th}}$ and $\overline{\delta w_{LTD}^{th}}$, corresponding to an infinite T. The black curve represents a trajectory obtained for a presynaptic firing rate of 10 spikes/s and different postsynaptic firing rates when $T = 10$ s and for $(\alpha_{LTP}, \alpha_{LTD}) = (5, 1.15)$. **E**, BCM-like plasticity curves obtained along the black trajectory with the same color code as in **D**. In the inset, a representative curve depicting the singularity around the floating threshold (BCM) and the zero domain of stability at the origin. **F**, Supralinear dependence of the floating plasticity threshold on the past postsynaptic firing rate history corresponding to the curves shown in **E**. In dashed black, the same curve when $T_{LTD} = 0$ ms. The input rate-dependent region in which only LTP can occur is indicated by an asterisk.

the ABS model, LTP and LTD thresholds were considered as fixed values. In the model by Senn et al., the effective floating threshold was not shared among synapses, contrary to the mSTDP rule. The slow variables are nonlinear functions of the postsynaptic activity history through the eligibility traces δw_{LTP} and δw_{LTD} in Equations 3 and 6 that are functions of the past postsynaptic activity for Poisson processes. For a given parameter set $(\alpha_{LTP}, \alpha_{LTD})$ and for a finite time constant $T = 10$ s and presynaptic rate of 10 spikes/s, successive stationary postsynaptic firing rates resulted in effective stationary variables δw_{LTP}^{th} and δw_{LTD}^{th} that are plotted as trajectories in Figure 4D. The white cross corresponds to a quiescent past postsynaptic activity, and the arrow indicates the floating plasticity thresholds for higher past activities. Figure 4E shows the plasticity curves as a function of the past postsynaptic firing rate (assumed constant during each protocol) for the trajectories drawn in Figure 4D. As described in Materials and Methods, two curve-axis crossing points can be highlighted: the first point at the origin of the postsynaptic activity axis, and the second point at the transition between pure LTD and LTP domains (BCM-like threshold). The corresponding two regimes are illustrated in Figure 4E (inset). The stable zero domain around the origin is usually large enough to operate as a "dead zone" and supports a static integrative mode with fixed synaptic weights. In this domain, no synaptic weight change will occur as long as the

level of presynaptic and postsynaptic coincidence remains small. This domain can be shrunk after high past synaptic activity (LTD-only regime). However, the second point, the LTD–LTP transition BCM threshold, increases with increasing past synaptic activity as described in Table 1. Figure 4F shows the BCM thresholds for increasing postsynaptic past activity corresponding to the different curves in Figure 4E. This relation is supralinear as required by the formal demonstration of stability conditions during competitive BCM learning (Bienenstock et al., 1982). The increase rate of this function as a function of the past synaptic activity is driven by the parameter β that describes the rate of LTP downregulation (Eq. 3). It is worth noting that the floating BCM threshold is a singular operating point that can only guarantee the convergence of the mean synaptic weight. The stochastic nature of neuronal firing creates a residual competition that slowly diffuses changes in the synaptic weights around the mean self-stabilized operating point. Thus, only the zero domain guarantees both individual and mean synaptic stability, allowing the cell to operate reliably in a passive (non-adaptive) integrative mode.

Stable competition of mSTDP during ongoing activity

To show that the mSTDP model produces stored weight patterns that are robust to ongoing activity, we investigated several net-

work models, namely feedforward and recurrent network models. In the feedforward model, we considered a presynaptic population of 1000 excitatory and 250 inhibitory neurons projecting to a postsynaptic neuron, in which only the excitatory synapses were subject to plasticity. To mimic the stochastic ongoing activity observed *in vivo*, each presynaptic neuron followed a Poisson process with a mean rate of 8 spikes/s. The parameters α_{LTP} and α_{LTD} were chosen such that no synaptic changes were observed in this regime, considered as the normal operating point of the network. One important issue, common to all models of synaptic plasticity, concerns the operating point of the synapses. In the ABS and Senn et al. (2001) models, the implicit reference is the quiescent network state: as soon as the postsynaptic activity increases, the lower plasticity threshold crossed by synapses is that of LTD induction. In contrast, the BCM and mSTDP models were proposed to account for the *in vivo* case. In this situation, the presence of ongoing activity differs from the absence of firing at rest *in vitro*. For BCM, the bending of the covariance rule, so that it converges to the origin, the "zero-activity" point was proposed only to ensure mathematical stability. Note that a pure covariance product rule (Sejnowski, 1977) would lead to instability and change in the sign of the synaptic weight (if allowed) near the origin. In our model, however, the "fixed" regime below LTP and LTD thresholds guarantees that individual synaptic weights remain fixed during ongoing activity. This stable information transmission domain may be thought to correspond to the basal level of network spiking correlation during spontaneous activity, which has been reported to be low *in vivo* (Renart et al., 2010) or during sparse evoked activity during natural scene viewing (Vinje and Gallant, 2000). However, both zero-domain and "crossing" (floating threshold of transition between LTD and LTP) points can be used to stabilize the mean synaptic weight, although the singular nature of the floating threshold still allows for a residual competition. This singularity has a possible functional drawback because competition of infinite duration will induce a progressive increase in the variance of the weight distribution without affecting its mean (first-order stability remains valid), as will be shown in the following. This second-order instability can be seen as a correlate of synaptic competition but in the long term results in a synaptic weight distribution in which a single weight overtakes all the others. Alternatively, our model could have been parameterized in such a way as to create a stability domain around the LTD–LTP crossing point, similar to the zero domain at the origin. This second stability domain would stabilize the full synaptic weight distribution at the operating point between LTP and LTD. Such a second stability domain was proposed on the basis of the experimental validation of the BCM rule (Frégnac et al., 1988), as an operating domain matching the dynamical range of evoked sensory and ongoing activity (Frégnac, 2002). These differences in the working state of the synapse will be the focus of future developments of the mSTDP model. In the following, we considered different heterogeneous presynaptic firing patterns to study the evolution of synaptic weights produced by mSTDP.

We first studied the case in which the presynaptic ensemble average of slow variables is made over all synapses, thus defining common LTP and LTD thresholds shared among all synapses as required by the BCM rule. We separately considered the effect of heterogeneous presynaptic firing rate (detailed in Figs. 5, 6) and synchrony (detailed in Fig. 7) within the presynaptic neuron population (see illustration in Fig. 5A). The nonlinear interspike interactions responsible for the BCM-like features of the plasticity curve can create competition between groups of neurons with different firing rates further stabilized by the slow variables. We

thus increased the mean firing rate of half the presynaptic excitatory neurons and, as expected, observed a rapid separation of their synaptic weights from those of neurons with unchanged rate (Fig. 5B). In particular, because of the strong impact of the LTP fast variables compared with LTD ($\alpha = 0.46$) and our choice of presynaptic rates (8 and 16 Hz), the synaptic weights of population 2 were increased when the mean firing rate of population 2 got higher. This change was accompanied by a simultaneous decrease in the synaptic weights fed by population 1, whose firing rate was kept unchanged. This effect is expected from the BCM-like competition feature of the plasticity algorithm, which tends to keep the average synaptic weight and thus the postsynaptic firing rate constant. The floating plasticity threshold dependence on the past postsynaptic activity is illustrated for both input populations in Figure 5C, with lower floating thresholds for synapses fed by population 2 as expected from the BCM threshold presynaptic dependence discussed in Materials and Methods. In these conditions, population 1 experiences depression and population 2 potentiation.

Eventually, when all synaptic weights of population 1 reached their lower bound at 0, the average synaptic weight corresponding to population 2 became stabilized so that the postsynaptic firing rate reached a new stationary regime. If the stimulation pattern was maintained, the variance of the global synaptic weight distribution would slowly increase because of the residual competition. To avoid weight divergence of individual synapses (stability is guaranteed by the mSTDP model for the mean but not for the variance of the synaptic weight distribution), this plasticity phase was then followed by a phase in which presynaptic firing rates were relaxed to their original values, resulting in an inactivation of LTP and LTD. Variables δw_{LTP} and δw_{LTD} for both populations returned under their respective thresholds $\delta w_{\mathrm{LTP}}^{\mathrm{th}}$ and $\delta w_{\mathrm{LTD}}^{\mathrm{th}}$ corresponding to the zero domain of stability (Fig. 5D). The resulting postsynaptic firing rate was comparable with its original value, and the synaptic weights remained constant in this new ongoing regime (Fig. 5B). Therefore, spatially heterogeneous distribution of presynaptic firing rates elicits, at the level of the common postsynaptic target cell, synaptic weight scaling of a competitive nature. Ensuring that the fast variables return under their respective thresholds thus results in a segregation and stabilization of individual synaptic weights. This result does not depend on the initial weights as can be concluded from similar simulations in which the initial values were drawn from a Gaussian distribution (data not shown).

However, as discussed in Figure 4A–C in the case of a single presynaptic source, the presynaptic dependence of the LTD–LTP transition threshold can invert the outcome of competition between the input populations, when presynaptic firing rates are chosen so as to produce LTD instead of LTP in the absence of slow variables. If, for instance, the presynaptic firing rates are chosen to be 20 spikes/s for population 1 (instead of 8 spikes/s in Fig. 5) and 50 spikes/s for population 2 (instead of 16 spikes/s in Fig. 5), a depression is predicted by Figure 4A for both synaptic sets, the larger the higher the presynaptic rate. Consequently, the synaptic weights corresponding to population 2 will become smaller than for population 1 (Fig. 6A). Indeed, in the singular BCM regime (around the LTD–LTP threshold), the floating plasticity threshold of population 1 is smaller than that of population 2, thus inverting the pattern of competition (Fig. 6B). Note that this presynaptic dependence is lost when T_{LTD} is very small or equal to 0. In Figure 6D–F, the same simulation was run with $T_{\mathrm{LTD}} = 50$ ms and resulted in synaptic weight distributions comparable with those in Figure 5, with stronger weights for popula-
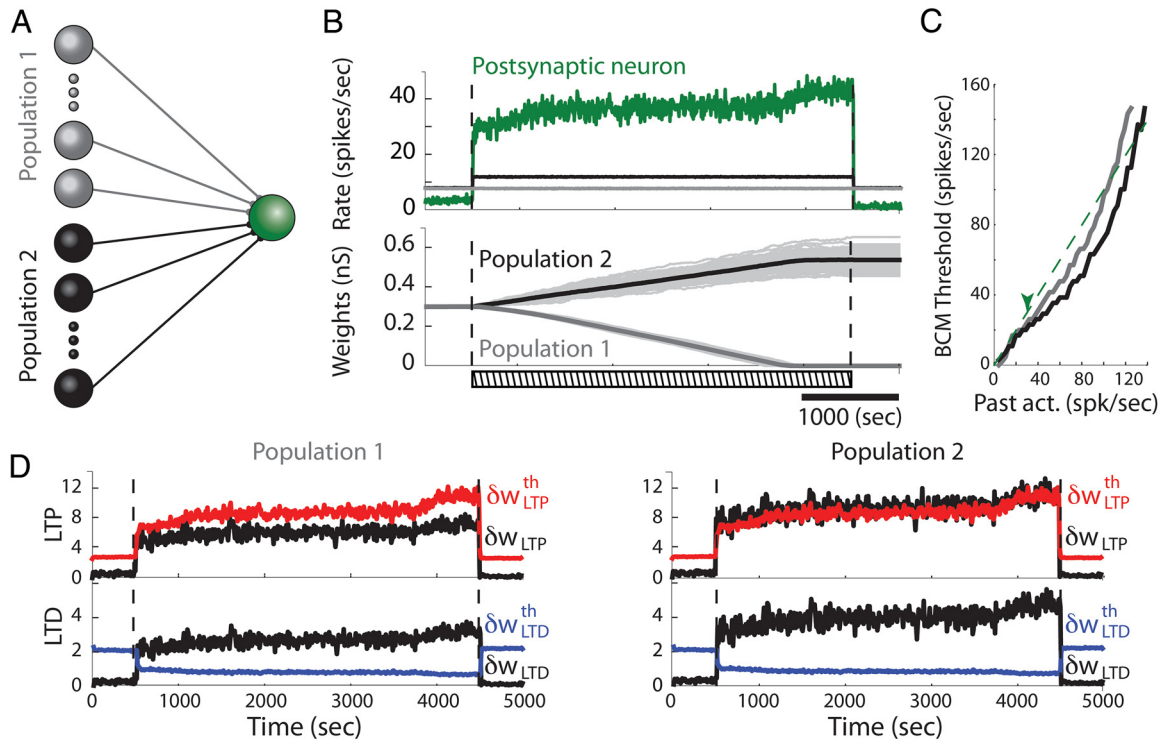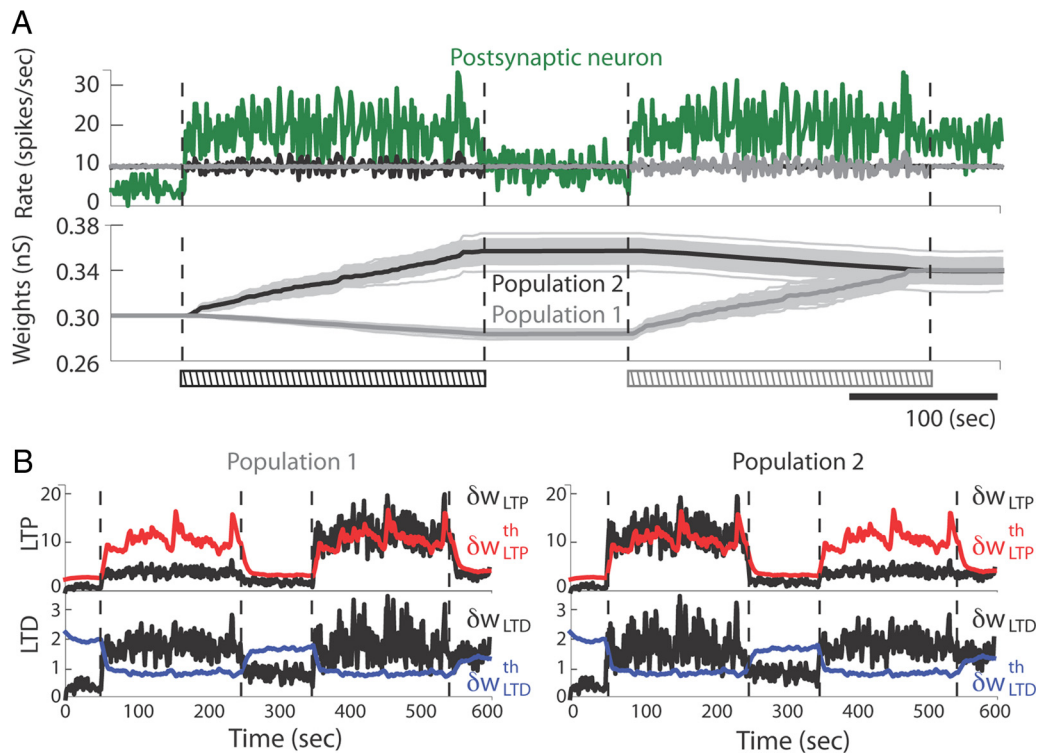
Figure 5. Firing rate-based competition produced by mSTDP. ***A***, Illustration of the feedforward model (only excitatory cells are shown). ***B***, Time evolution of firing rates (top) and synaptic weights (bottom) in a feedforward model with presynaptic Poisson processes. In this protocol, the postsynaptic neuron receives Poisson input from two presynaptic populations. After 500 s, presynaptic neurons in population 2 double their firing rates from 8 to 16 spikes/s (black curve), whereas the other population remains unchanged (gray curve). The firing rate of the postsynaptic neuron is shown in green. The bottom plot represents individual synaptic weights in light gray, as well as the mean weight of population 1 in dark gray and population 2 in black. After 4000 s, the firing rate of population 2 is relaxed to its initial value. The dashed bar indicates the onset, offset, and duration of the stimulation. ***C***, Floating plasticity threshold dependence on the past postsynaptic firing rate for both populations. The green arrow indicates the stable postsynaptic firing rate during the competitive learning protocol, before population 1 weights reach their lower bound at 0. The dashed green line indicates identity. ***D***, Time evolution of the fast variables δw_{LTP} and δw_{LTD} (black curves), as well as the corresponding slow variables δw_{LTP}^{th} (red) and δw_{LTD}^{th} (blue).
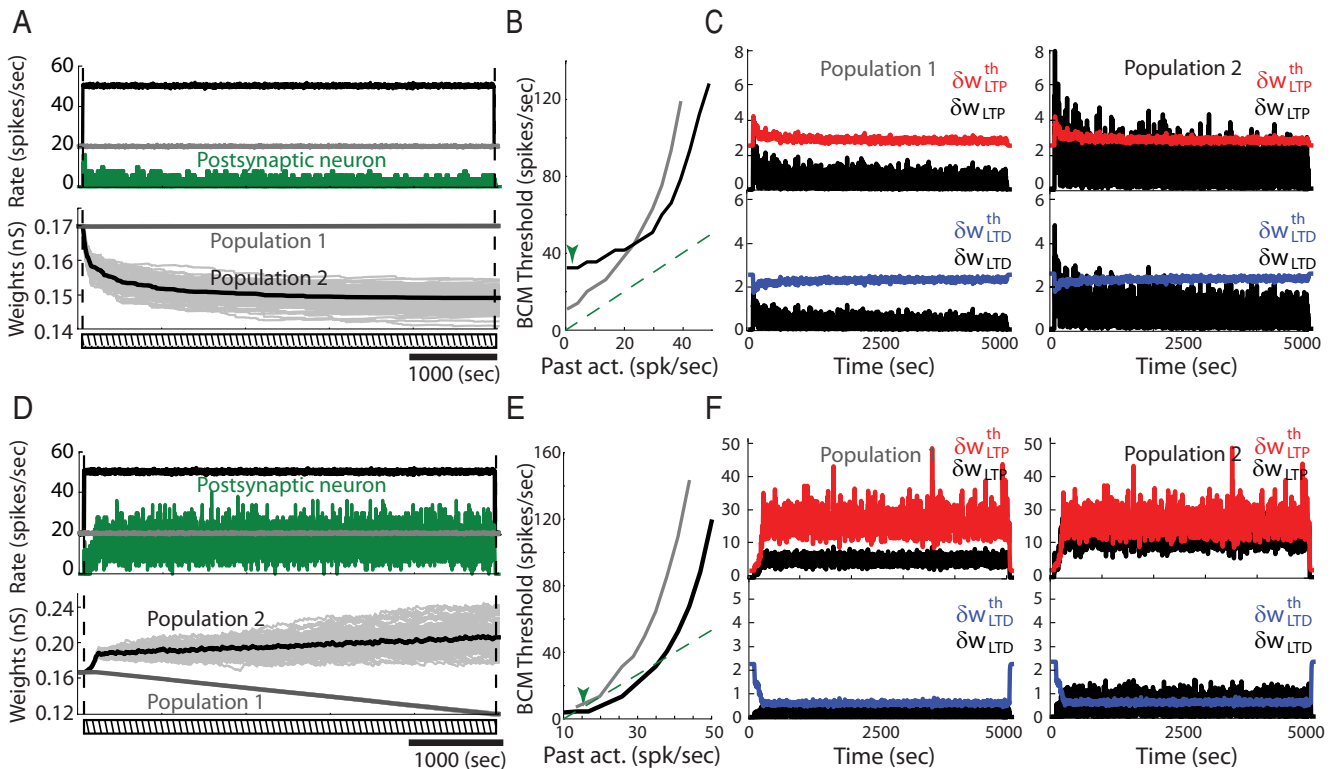
Figure 6. Effect of the time constant T_{LTD} on the presynaptic dependence of the floating plasticity threshold. ***A–C***, Same as in ***B–D*** in Figure 5 but with a mean firing rate of 20 spikes/s for population 1 and 50 spikes/s for population 2. Here, as in Figure 5, $T_{LTD} = 1000$ ms. ***D–F***, Same as in ***A–C*** but with $T_{LTD} = 50$ ms.

Figure 7. Synchrony-based competition produced by mSTDP. ***A***, Time evolution of firing rates (top) and synaptic weights (bottom) in a feedforward model with presynaptic Poisson processes. In this protocol, population 2 neurons are synchronous Poisson processes (correlation coefficient of 0.25) between 50 and 250 s (black crosshatched bar), and population 1 neurons are synchronous Poisson processes between 350 and 550 s (gray crosshatched bar). Outside these synchrony periods, the synaptic inputs are uncorrelated. ***B***, Time evolution of the fast variables δw_{LTP} and δw_{LTD} (black curves), as well as the corresponding slow variables δw_{LTP}^{th} and δw_{LTD}^{th}.

tion 2. This can be seen in particular in the dependence of the floating plasticity threshold on the past postsynaptic activity (Fig. 6*E*), which shows smaller values for population 2, corresponding to potentiation in the singular BCM regime.

We next considered the case of correlated Poisson processes of variable synchrony with a fixed mean firing rate. As reported by Song and Abbott (2001), if half of the presynaptic neurons are driven by synchronous inputs while the other half are randomly activated, heterosynaptic competition occurs between the synchronized and unsynchronized sets of synapses and produces two separate synaptic weight distributions corresponding to each input population. In Figure 7*A*, during a first phase, half the presynaptic neurons (population 2) were driven by correlated Poisson processes (correlation coefficient of 0.25 for 200 s). The resulting synaptic weight distribution after the rapid learning phase is separated into two distinct sharp unimodal distributions. If the synchronous firing of population 2 were to last longer, the synaptic weights of population 1 would eventually reach 0 and the average synaptic weight of population 2 would become stable as in Figure 5 (data not shown). Instead, this phase was then followed by a stationary phase of 100 s, during which the spiking activity of both populations was decorrelated. Fast variables were brought back below their induction thresholds (Fig. 7*B*), resulting in a stabilization of individual synaptic weights. To show that these changes are not irreversible and that homogeneous distributions can be recovered with appropriate synaptic inputs, we then inverted the two population spiking patterns for the same stimulation duration. The weight distribution self-organized into a unimodal distribution in which the synaptic weights of the two populations were indistinguishable. The combined effect of these alternated patterns of correlated stimulations induced a global

potentiation of the mean synaptic weight, each set of synapses benefiting more from LTP than LTD, as can be seen in the time evolution of all variables (Fig. 7*B*). From this result, we concluded that presynaptic populations can be segregated based on both first-order (mean rate) and second-order (synchrony) input statistics, thus resulting in efficient synaptic competition mechanisms that rely on different aspects of the mSTDP rule.

mSTDP with independent slow variables
We also considered the case in which synapses do not share a common plasticity threshold. Interestingly, new forms of synaptic competition can arise in this situation. We first replicated the stimulation protocol of Figure 5 where two presynaptic populations fire with different mean rates. In Figure 8*A*, the increase in postsynaptic firing rate elicits a fast and strong change in population 2 synaptic weights. However, instead of being potentiated as in Figure 5, these weights are strongly depressed. Indeed, the increase in slow variables (Fig. 8*C*) produced by the strong presynaptic and postsynaptic firing rate only affects the synapses of population 2, resulting in a shift toward depression for the corresponding weights until the fast variables return below their respective thresholds. Eventually, individual synaptic weights settle at constant values and the postsynaptic firing becomes stationary. Although this pattern of competition resembles that observed in Figure 6*A–C*, this effect is not the result of the presynaptic dependence of the LTD–LTP transition threshold. Indeed, the protocol is identical to that illustrated in Figure 5, and the singular LTD–LTP transition BCM point still keeps the same form of dependence on the past synaptic activity (Fig. 8*B*) for small T_{LTD}. The homeostatic regulation acts in this latter case on each synapse independently and therefore produces a strong depression to
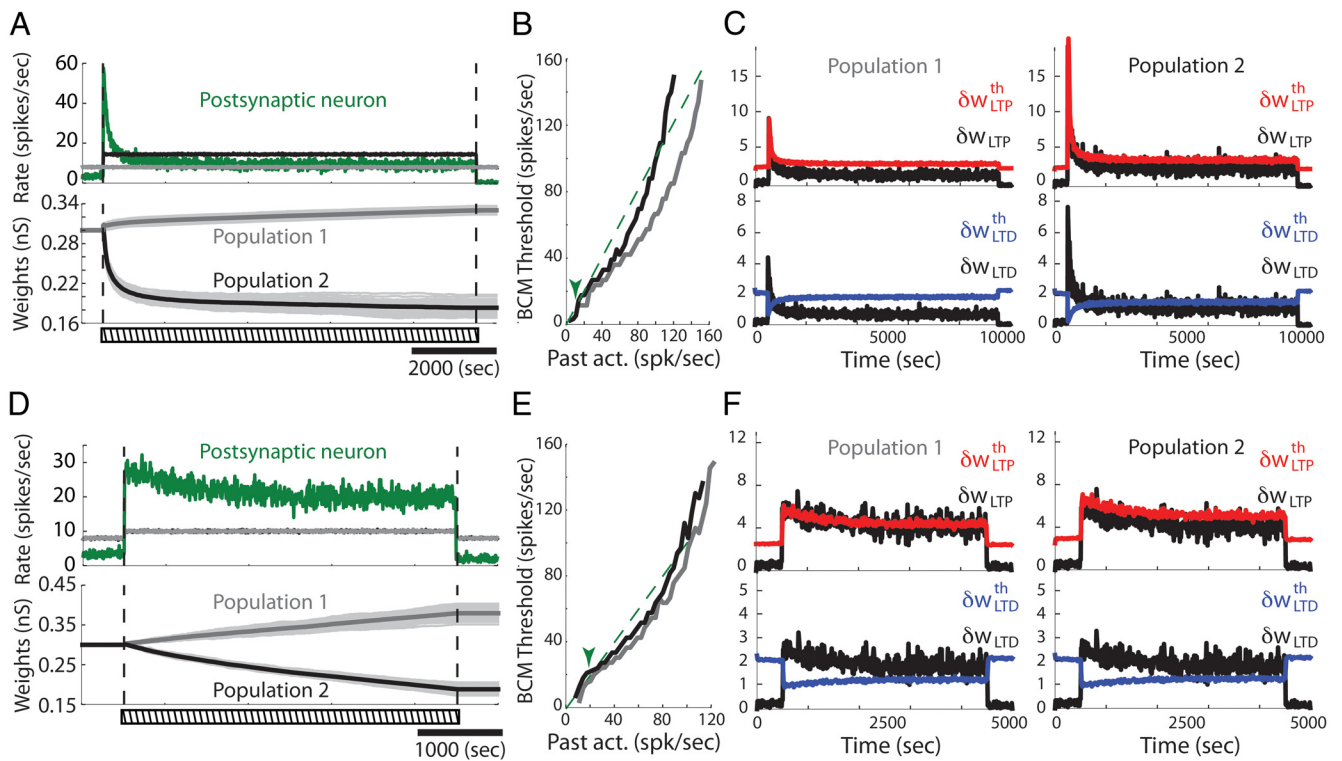
Figure 8. Synaptic weight dynamics for mSTDP rules with independent slow variables. *A*, The time evolution of the firing rates of the presynaptic and postsynaptic neurons is shown in the top and the synaptic weight distributions of populations 1 and 2 in the bottom. Population 2 neurons increase their firing rate during the period delimited by the dashed black bar. *B*, Floating plasticity threshold dependence on the past postsynaptic firing rate for both populations. *C*, Time evolution of the fast and slow variables of both populations 1 and 2. The color code is identical to that of Figures 5–7. *D–F*, Same as in *A–C* but for a uniform increase in firing rate and for a bimodal distribution of coefficients α_{LTP} with means of 2.3 and 2.7, corresponding to populations 1 and 2, respectively.

compensate for the increase in presynaptic firing rate and keep the postsynaptic mean rate constant. Therefore, different forms of learning rules can arise based on the spatial extent of synaptic interactions within the same neurite.

This synaptic weight segregation can arise from functional competition but can also result from structural specificity. If the coefficients α_{LTP} are drawn from a bimodal distribution, two distinct populations of synaptic weights emerge for a fixed and homogeneous presynaptic firing rate (Fig. 8*D–F*). These populations correspond respectively to each mode of the bimodal distribution. Population 1, having smaller LTP induction thresholds, experiences an increase in the synaptic weights, whereas population 2 synaptic weights are dominated by depression. If synapses interact only locally, synaptic weight distributions can be spontaneously segregated as a result of structural differences that may be related, for instance, to different dendritic branching morphologies.

Emergence of receptive fields with mSTDP

To illustrate the emergence of receptive fields in a simplified retinotopic feedforward model, we next considered an additional protocol in which retinotopically organized presynaptic neurons were driven by Poisson processes with mean rate following spatial Gaussian profiles (Song and Abbott, 2001; Pfister and Gerstner, 2006). Every 50 ms, the center of the Gaussian stimulation profile was randomly drawn among the presynaptic neurons with periodic conditions. After an initial phase of ∼500 s, the postsynaptic firing rate reached a stationary regime (Fig. 9*A*) and a sharp and localized receptive field emerged in the presynaptic population (Fig. 9*B,C*) ∼500 s later. After ∼2000 s, the synaptic weights reached their upper ($w_{\text{max}} = 0.5$ nS) or lower ($w_{\text{max}} = 0$ nS)

bound, and the receptive field stayed stable for the rest of the simulation. If the synaptic weight is unbounded, the excessive competition between synapses could shrink the receptive field to a single winner-take-all synapse. To avoid this anomalous trivial weight distribution, some additional constraint must be used, such as imposing weight dependency or hard bounds, as in Figure 9.

Apart from the general issue of local versus global (mean) stability that we already discussed at length, these simulations demonstrate that mSTDP is well suited for the creation of stable functional maps during development. In this model in which synapses can be ordered with a spatial arrangement based on the stimulation patterns, we studied the effect of spatial spread of LTP and LTD thresholds (slow variables). Instead of averaging these slow variables over all synapses at each time step, we assumed that the focally induced slow variable changes $\delta w_{\text{LTP}}^{\text{th}}$ and $\delta w_{\text{LTD}}^{\text{th}}$ are convolved with a sharp Gaussian function every millisecond to mimic a spatiotemporal diffusion. The corresponding diffusion process was characterized by a diffusion coefficient D given by the ratio between the Gaussian variance and the elementary time step ($D = 10^3$ units^2/s, where the spatial unit is the elementary distance between two neighboring neurons). The simulation shows that, in this scenario, the postsynaptic neuron settles in a regime in which a bimodal receptive field emerges (Fig. 9*D–F*). Although this bimodal receptive field can spontaneously emerge at any position from one simulation to another, the two modes are always equidistant in both directions because of the periodic conditions. For a finite diffusion coefficient, local competition can produce more complex structures than global competition mediated by threshold values shared by all synapses. We hypothesize that this local competition, caused by the depen-
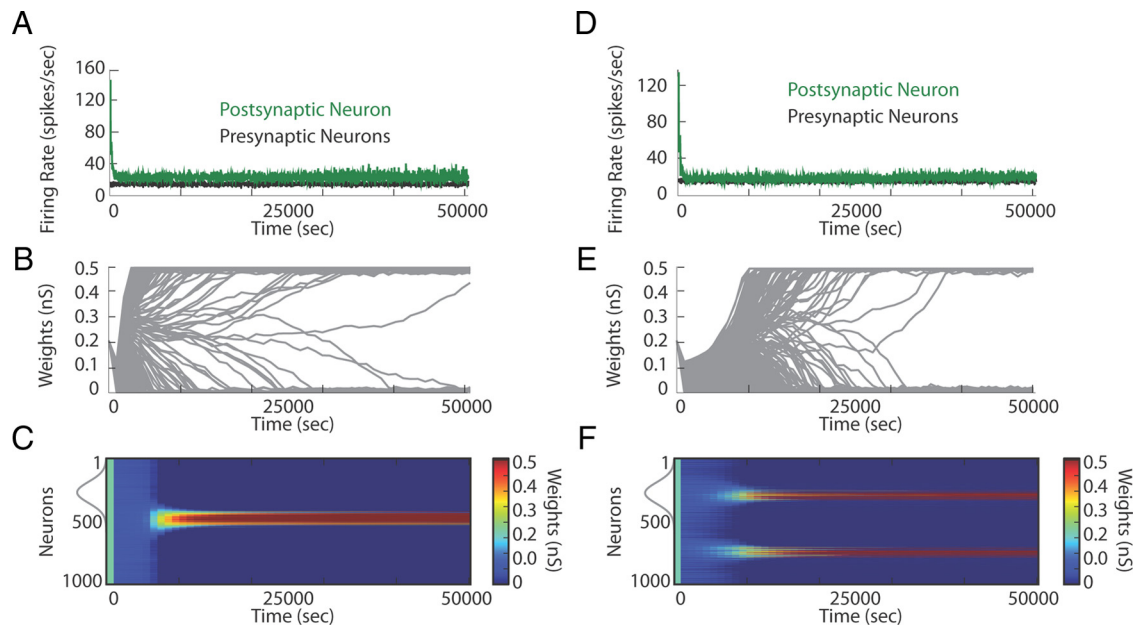
Figure 9. Emergence of receptive fields with random and sparse activation of a spatially Gaussian distributed stimulation. **A**, Time evolution of the mean presynaptic (black) and postsynaptic (green) neuron firing rates. **B**, Time evolution of individual synaptic weights with hard upper bounds. These simulations were done by averaging slow threshold variables over all synapses. **C**, Emergence of a unimodal receptive field (one-dimensional mapping on the neuronal sheet). Synaptic weights are color coded and illustrated over the whole stimulation period. The Gaussian profile of the stimulation is shown in gray. **D–F**, Same as in **A–C** but with finite slow variables diffusion coefficient ($D = 10^3$ units ^{2}/s).

dence of slow variables on the presynaptic firing activity, could occur independently in each dendritic compartment in biological neurons. For very small diffusion coefficients, slow variables are not shared among synapses so that no competition can take place between nearby synapses and no spatial structure emerges after stimulation. From a functional point of view, local competition increases the sensitivity of the postsynaptic neuron to the presynaptic ensemble by creating multiple segregated receptive fields in the stimulus space, possibly resulting in synapse clustering on distinct dendritic regions of the same neuron (Jia et al., 2010).

Stable learning in a recurrent network with stochastic-like regimes

We next considered the difficult problem of synaptic plasticity in recurrent networks displaying *in vivo*-like spontaneous activity. In this case, synaptic weights directly change the network dynamics and vice versa (Lubenov and Siapas, 2008). For appropriate parameters, balanced networks of integrate-and-fire neurons can display asynchronous irregular (AI) states (Brunel, 2000; Vogels and Abbott, 2005; Kumar et al., 2008; El Boustani and Destexhe, 2009; Marre et al., 2009) that produce spike discharge patterns similar to those reported in awake animals (Steriade, 2001; El Boustani et al., 2007). Recurrent neuronal networks have been studied with weight-independent STDP rules (Izhikevich et al., 2004; Morrison et al., 2007; Kang et al., 2008; Lubenov and Siapas, 2008; Gilson et al., 2010). It has been reported that these rules usually disrupt the ongoing regime by strengthening coactive neurons, thus producing strong synchronous volleys of firing (Morrison et al., 2007).

It was also shown that the addition of homeostatic mechanisms and axonal delays can stabilize the dynamics and partially abolish the pathological regimes by creating multiple subpopulations of synchronized neurons (Izhikevich et al., 2004) or by setting the network dynamics into a mixed state with episodic synchronous firing (Lubenov and Siapas, 2008). However, in the

latter state, the network loses the memory-retention property usually expressed by weight-independent STDP rules. If endowed with a weight-dependent STDP rule, such networks can sustain AI regimes with stable synaptic weight distributions (Morrison et al., 2007). However, individual synaptic weights are subject to strong fluctuations as a result of the ongoing spiking activity. Consequently, changes in the weight distribution are always transient because synaptic weights relax to their stable stationary distribution. There is therefore a fundamental tradeoff (catastrophic forgetting) with classical STDP rules between maintaining the stability of network dynamics or stabilizing the traces of the learning process (memory retention). We compared a weight-dependent STDP rule (van Rossum et al., 2000) and the mSTDP rule in a topological network with lateral propagation at finite speed, displaying an AI regime with a stable synaptic weight distribution for weight-dependent STDP rules. For 1000 s, these networks were subject to their spontaneous activity and synaptic weights converged to their equilibrium distribution. Then a circular area of the network with radius $r = 0.2$ mm (network side length $L = 1$ mm) was stimulated with an additional external input of 250 spikes/s while mean firing rates and mean afferent synaptic weights coming from the stimulated and unstimulated regions were monitored (Fig. 10 A). In these simulations, plasticity was authorized only within recurrent connections and externally driven synapses were kept fixed. When the network was stimulated, the mean intracortical synaptic weights projecting to (afferent) or sent by (efferent) the cortical neurons in the stimulated region were strongly affected by the increase in firing rate (Fig. 10 $B.2$, mSTDP). Only synaptic weights involving a direct lateral connection with the stimulated region were affected ("lateral receptive field").

The resulting stable pattern of synaptic weights displayed two distinct but complementary regions inscribed within the part of the receptive field corresponding to the trained lateral connectivity in which mean afferent weights per neuron had been either
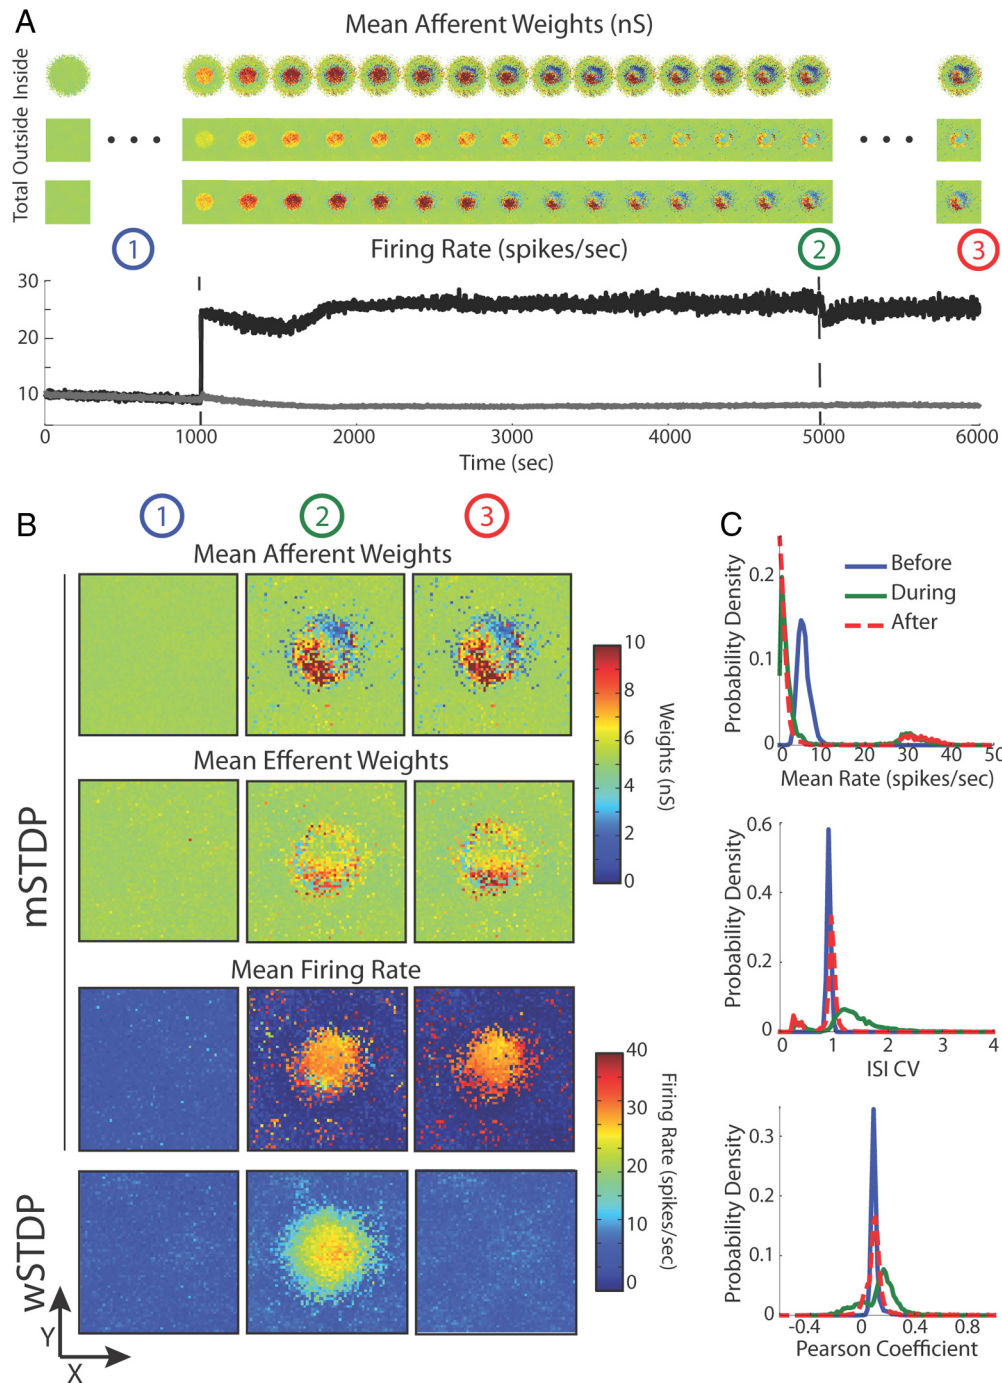
Figure 10. Stable learning in a topographically organized recurrent two-dimensional neuronal network with *in vivo*-like spontaneous activity. **A**, Time evolution of the unstimulated population (gray curve) and stimulated population by an external drive within a circular area (black curve). The mean afferent synaptic weights for each neuron in the network are displayed as snapshots in the course of time above the graph. The mean afferent weights fed by neurons inside the stimulated region, outside the stimulated region, and from all neurons are plotted in the first, second, and third rows, respectively. **B**, From left to right, The columns labeled by circled numbers correspond, respectively, to the stable spontaneous regime (**1**), the end of the stimulation epoch (**2**), and the relaxation ongoing regime 1000 s after the stimulation (**3**). Network snapshots of the mean afferent and efferent synaptic weight per neuron are plotted in the two first rows for mSTDP. In the last two rows, the mean firing rates for each neuron in the network are plotted for the mSTDP and the weight-dependent STDP rules. Circled numbers at the top of each column refer to the same phases described in **A** and **B**. **C**, Network firing statistics before, during, and after stimulation: top, firing rate distribution; middle, ISI CV; bottom, mean Pearson's coefficient computed over 2000 distinct neuron pairs.

potentiated or depressed (Fig. 10*B*, mSTDP). The inhomogeneous changes within the lateral receptive field are presumably caused by the local connection distribution for each neuron following a Gaussian profile with SD of 0.1 mm and the propagation delays allowing a strong local competition in the stimulated region. The resulting firing rate at the end of the stimulation was

considerably increased inside the spatial domain defined by the lateral receptive field, whereas the AI state of the network (Pearson's coefficients and ISI CV) observed before the stimulation was conserved after the stimulation has been removed (Fig. 10*C*). Therefore, mSTDP allowed the network to remain in its ongoing stochastic-like regime while displaying long-lasting memory re-

tention. These modifications did not last in the classical STDP model (Fig. 10*B.3*, STDP) because of the spontaneous network activity. We also tested this paradigm for a weight-independent STDP rule. Although the network synaptic weights kept track of the stimulation, the dynamics settled in a pathological regime of synfire explosions (Morrison et al., 2007) (data not shown). This catastrophic-forgetting tradeoff was avoided by the mSTDP model, which induced stable synaptic changes in a stable network state through the existence of a floating plasticity threshold (Fig. 10*B.3*, mSTDP). In this case, the system operates in a stable domain of activity (see zero domain above).

Finally, we considered a stimulation paradigm in which the whole network was stimulated using "moving bars" of activity and in which the lateral connectivity adapts to the conditioning experience through mSTDP. The conditioning bars are typically one-third the length of the network size and can start from any point and propagate in any direction (see Materials and Methods). Each receptive field is defined by the spread and strength of connected afferents relative to the postsynaptic cell. The impact on the receptive field is measured by potentiation or depression of afferent synapses: STDP favors potentiation of synaptic links in which the presynaptic cell is "hit" by the bar before the target cell. STDP depresses synaptic links in which the postsynaptic cell firing is evoked before the presynaptic one. We first used an isotropic distribution of stimulus directions by drawing values from a uniform polar distribution between 0 and 2π. After stimulating for 1500 s, structured domains of polar asymmetry bias emerged across the network (Fig. 11*A*, unbiased directional stimulation). Results show that these asymmetric structures in the spatial distribution of afferent presynaptic cells, observed in different target cells for every possible direction, were preserved once the moving bars were no longer presented and the network settled to a new level of spontaneous activity. These patches tended to form unidirectional preference domains at a more mesoscopic level. In a second phase, the same network was stimulated using moving bars strongly biased toward one direction value ($\phi = 0$; Fig. 11*A*,*B*, red arrow) using a wrapped Gaussian distribution. The direction domains were consequently strongly affected and reflected the new conditioning stimulus preference after 1500 s of learning (Fig. 11*A*,*B*, biased stimulation). This sharp directional bias in the synaptic weight distribution is visible when looking at the individual spatial profiles of receptive fields that are all strongly biased toward directions opposite to the stimulus bias, i.e., which are stimulated on the path ahead of the postsynaptic receptive field location. (Fig. 11*C*). The simulations showed that the receptive fields are biased in two halves: one corresponding to a set of presynaptic cells whose synaptic impact has been potentiated, the other to presynaptic cells whose synaptic contribution has been depressed. The demarcation line between the "potentiated" and "depressed" presynaptic sources goes through the postsynaptic cell position and is cross-oriented to the motion path of the stimulus. Such asymmetric effects have been ascribed to STDP in the visual system by Fu et al. (2002, 2004) and to the place cells in hippocampus by Mehta et al. (2000). Interestingly, the network firing statistics after biased and unbiased stimulations were similar and significantly different from the initial state (Fig. 11*D*). In particular, the mean Pearson's correlation coefficient measured over 2000 distinct neuron pairs decreased to almost 0 after learning, although nearby neurons shared similar functional properties. As demonstrated by this model, stable and flexible learning can take place in recurrent networks expressing sustained and irregular patterns of spontaneous activity.

Discussion

We have introduced here a metaplastic STDP model that generalizes a number of previous models and accounts for most experiments on associative synaptic plasticity: it provides a common framework to integrate the main phenomenology of STDP together with metaplasticity mechanisms. Previous models also included an activity-dependent induction threshold (Bienenstock et al., 1982), or metaplasticity mechanisms (Lisman, 1989; Fusi et al., 2005; Graupner and Brunel, 2007), but these models were not consistent with STDP or with priming experiments. Similarly, previous models of higher-order STDP interactions (Senn et al., 2001; Shouval et al., 2002; Badoual et al., 2006; Pfister and Gerstner, 2006; Benuskova and Abraham, 2007; Clopath et al., 2010) did not include metaplasticity as an emergent property of the model. The mSTDP model aims at being consistent with these different timescales. It accounts for the frequency dependence of STDP but also for slower timescale priming experiments. Most importantly, our mSTDP model can lead to stable learning in the presence of spontaneous activity without disturbing the network AI regime, therefore providing a possible solution to the catastrophic forgetting problem.

One original component of the mSTDP model is the LTP and LTD activity-dependent thresholds. Because of their dependence on presynaptic spiking activity, these thresholds can be understood as the context represented by the input statistics evaluated over a presynaptic ensemble. In the most extreme case, in which averages are performed over all synapses, identical thresholds are shared by all synapses, a behavior that recovers the heterosynaptic nature of the BCM rule. However, when less trivial presynaptic ensembles are defined, competition can occur more locally, resulting in the emergence of multimodal receptive fields (see related biological evidence in the study by Jia et al., 2010). This can be achieved by local diffusion of the individual thresholds, with an underlying topographical organization added. Our model offers a possible framework for detailed compartmental description of different integrative subfields underlying distinct STDP-based operations in which the validating postsynaptic signal can be either the sodium spike or the local calcium spike boosted by backpropagation of the action potential in the distal dendrite (Frégnac, 1999; Bar Ilan et al., 2011). STDP models have been developed along these lines showing the emergence of spatially distinct clusters of synaptic "engrams" on the dendrites (Iannella et al., 2010).

For network dynamics *in vivo*, this novel form of plasticity is able to serve simultaneously two main functions: (1) achieve long-term stabilization of synaptic changes induced by recurring correlation patterns independently of ongoing disruptions produced by background activity; and (2) provide a plausible mechanism of homeostasis and self-normalization for synaptic learning during changes of local or global context. Its compatibility with a plausible identified biophysical mechanism is thus testable in future *in vitro* and *in vivo* experiments. Several experiments could be designed to test specifically some of the key assumptions of the mSTDP model.

A straightforward functional prediction is that successions of short and repeated stimulations of the synapse, interleaved with long resting periods, should elicit more plasticity at the synapse than a single stimulation period of the same cumulative time. This form of "spaced training" protocol has simple predictions that have been experimentally tested (Zhou et al., 2003; Dunfield and Haas, 2009). During a single long-lasting stimulation, the thresholds for LTP and LTD are reported to slide to protect the

Directional Bias Map
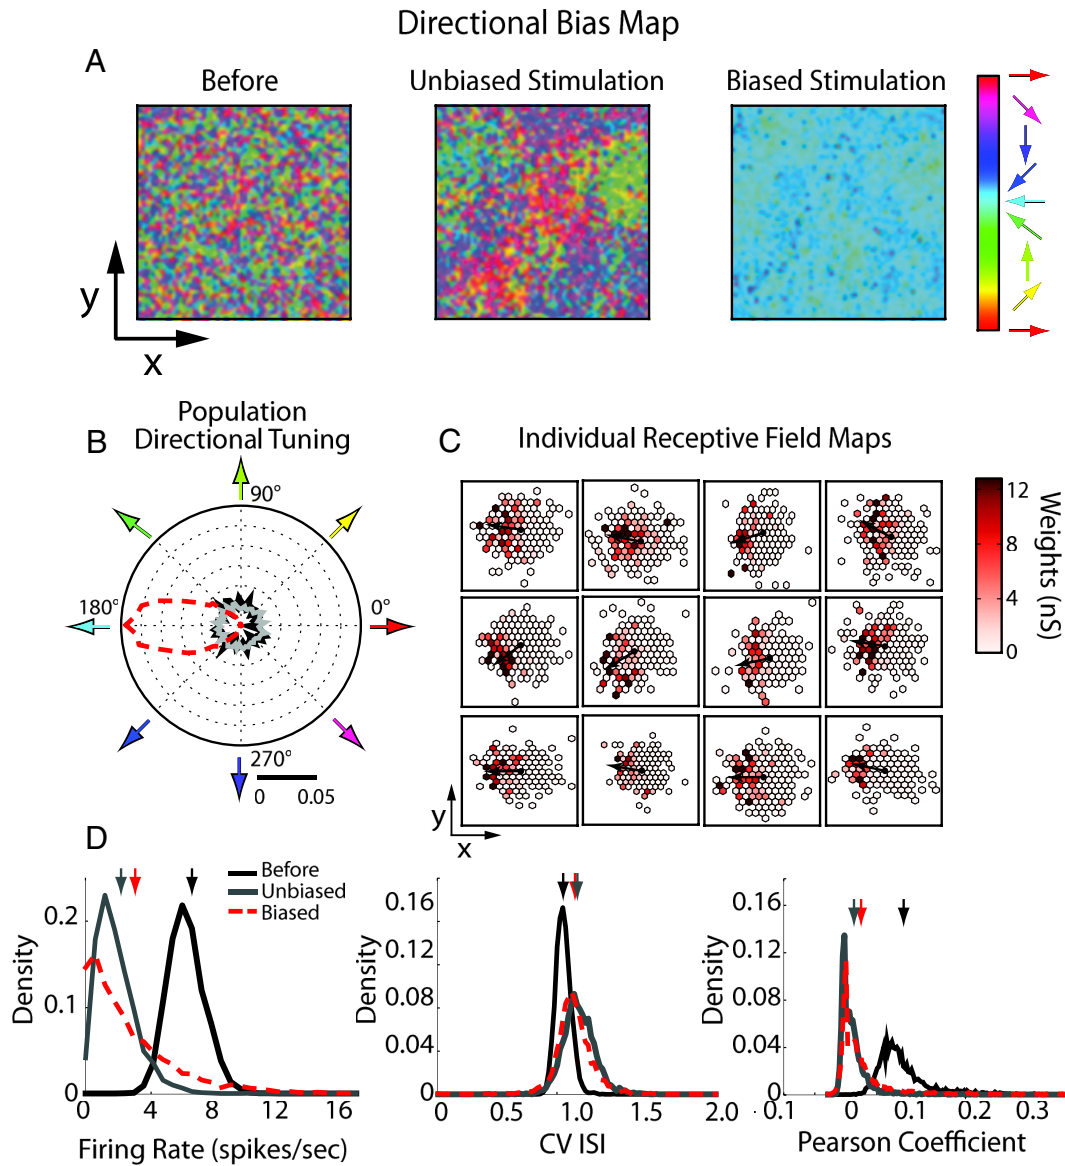
Figure 11. Stable learning in recurrent retinotopically organized networks stimulated with moving bars. **A**, Maps of direction preferences, color coded and computed for each neuron in the network from their individual receptive fields. For one neuron, the receptive field corresponds to the spatial distribution of its synaptic sources (locations of the presynaptic neurons afferent to the same target cell). Individual direction preferences were computed by identifying the direction, in each receptive field, in which recurrent presynaptic weights were the strongest (see Materials and Methods). The left shows the map of the direction preferences in the initial, randomly connected network under spontaneous activity. Middle and right shows these maps induced by plasticity respectively after unbiased (middle) or directionally biased (right, red arrow) conditioning stimulations. **B**, Distribution of the direction preferences gathered across all neurons during spontaneous (black curve), directionally unbiased (gray curve), and biased conditioning (red dashed curve) stimulation. The direction tunings are measured from individual receptive fields as spatial biases in the recruitment of presynaptic sources. **C**, Examples of individual receptive fields for several neurons randomly selected in the network after the biased stimulation. The maps are centered on the reference postsynaptic neurons (black dots), and individual direction preferences are indicated by black arrows. Incoming synaptic weights are color coded. **D**, Network firing statistics before stimulation and after the different stimulation phases: left, firing rate distribution; middle, ISI CV; right, mean Pearson's coefficient computed over 2000 distinct neuron pairs. Arrows indicate mean values for each phase.

synapse against too large changes, which may not be the case in short learning episodes. If those episodes are spaced enough in time, thresholds go back to their resting values, and therefore additional stimulations still trigger changes. This type of dynamics is reminiscent of the alternation of "up" and "down" states during slow-wave sleep, which consists of periods of activity interleaved with periods of silence (for review, see Steriade, 2001; Destexhe et al., 2007). This period of sleep may play a role in memory consolidation (Gais et al., 2000; Stickgold et al., 2001; Ji and Wilson, 2007), but no precise mechanism has been proposed yet. We suggest that the purpose of down states may be to relax the sliding thresholds, thereby enabling the activity during the up

states to efficiently trigger long-term plasticity at neocortical synapses when the network endogenously replays the transiently stored patterns. This is a challenging problem to investigate in the future, by using networks models displaying up–down states or by studying the influence of nonstationarities in the background activity on STDP by imposing a simulated bombardment of conductances using dynamic-clamp techniques.

Another experimental test, more specific to the mSTDP model, would be to characterize putative changes in the postsynaptic plasticity threshold across the neurite of the cell, as implied by the diffusion hypothesis. A feasible protocol would be to use glutamate uncaging under two-photon laser stimulation to selec-

tively target and activate, in a non-uniform manner, the dendritic branches of a labeled neuron. Similarly to what has been shown in Figure 9, local diffusion of the LTP and LTD thresholds should promote local interactions between neighboring synapses, leading to the emergence of spatially clustered synaptic patterns along the dendrite. Crosstalk between nearby synapses during plasticity processes has already been reported *in vitro* (Harvey et al., 2008), with a spatial scale up to few micrometers. We hypothesize that diffusion related to the slow variables in mSTDP would act on slower timescales (several hours) and affect a larger extent of the dendrite. The quantification of their time and spatial constants would provide some insight about the underlying mechanisms.

From a functional point of view, the presynaptic dependence of the plasticity thresholds could also be revealed by introducing stimulus-driven competition within or across dendrites of the same postsynaptic cell *in vivo*. It has become possible, by using two-photon imaging techniques, to map distinct preferred-orientation subdomains along a single dendritic tree or across several branches of the same postsynaptic neuron (Jia et al., 2010). Differential priming stimulations by visually driven activation of locally clustered or spatially segregated orientation preference domains could be used to measure the spread of the plasticity threshold changes on a single or on distinct dendritic branches.

It is important to note that the model ingredients are consistent with known molecular pathways involved in associative synaptic plasticity. A likely candidate to mediate metaplasticity processes was suggested to be an autophosphorylated state of calcium/calmodulin kinase II (CaMKII) (Bear, 1995; Mayford et al., 1995; Giese et al., 1998). In particular, autophosphorylation at site Thr305/306 has a suppressive effect on the calmodulin kinase binding, acting on a very slow timescale as described in our model (Elgersma et al., 2002; Zhang et al., 2005). The phosphorylation state at these sites has been shown recently to play a pivotal role in the induction of LTP and LTD, thus being an ideal candidate for the regulation of the BCM sliding threshold (Pi et al., 2010). The CaMKII mechanism has been profusely studied in theoretical models (Lisman, 1989, 1994; Lisman et al., 2002; Graupner and Brunel, 2007), showing that it is well suited to induce multi-stability states. However, these models did not incorporate the slow LTP downregulation necessary to produce the BCM sliding threshold. Following Wang et al. (2005), LTD could be mediated through calmodulin-dependent activation of calcineurin after calcium influx trough L-type channels. We made here the additional assumption regarding the increased availability of calmodulin for LTD when priming stimulations elicit LTP downregulation. Moreover, in this biological interpretation, local competition could result from the spatial spread of the autophosphorylated state of CaMKII among nearby synapses. These critical assumptions can be tested experimentally to directly verify the implication of these predicted molecular mechanisms.

There is ample biological evidence for synaptic consolidation processes that are distinct from the plasticity induction process itself and possibly involve mechanisms that have not been considered here (for review, see McGaugh, 2000; Dudai, 2004). For instance, protein synthesis inhibition has been shown to interfere with LTP (Gold, 2008) and is necessary for reconsolidation of old memories (Tronson and Taylor, 2007). However, the time constants to which we refer here (a few tens of seconds to minutes) are typically one or two orders of magnitude below those implied in distributed spacing of learning episodes, which justifies the limited scope of our biophysical metaplasticity scenario.

Finally, a possible extension of the model is that the dynamics of plasticity thresholds could be used in the context of reinforcement learning. There is biological evidence in both vertebrates and invertebrates that STDP is modulated by neuromodulators (Seol et al., 2007; Pawlak and Kerr, 2008; Shen et al., 2008; Zhang et al., 2009; Pawlak et al., 2010; S. Cassenaer, and G. Laurent, personal communication), which are known to be key factors in memory retention (Kasamatsu and Pettigrew, 1976; Romo and Schultz, 1990). Models using STDP modulated by biofeedback signals can reproduce reward-based paradigms (Legenstein et al., 2008), and, as Gu and Yan (2004) showed, the activation of dopamine D_4 receptors in prefrontal cortex exerts a complex regulation of CaMKII. This mSTDP model could be modified to include a regulation of the internal thresholds by reward-based signals, thus creating temporal windows during which the synapses could be modified and sensitive to STDP.

References

Abraham WC (2008) Metaplasticity: tuning synapses and networks for plasticity. Nat Rev Neurosci 9:387.

Abraham WC, Bear MF (1996) Metaplasticity: the plasticity of synaptic plasticity. Trends Neurosci 19:126–130.

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

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.

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

Bar Ilan L, Gidon A, Segev I (2011) Inter-regional synaptic competition in neurons with multiple STDP-inducing signals. J Neurophysiol 105:989–998.

Bear MF (1995) Mechanism for a sliding synaptic modification threshold. Neuron 15:1–4.

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

Benucci A, Frazor RA, Carandini M (2007) Standing waves and traveling waves distinguish two circuits in visual cortex. Neuron 55:103–117.

Benuskova L, Abraham WC (2007) STDP rule endowed with the BCM sliding threshold accounts for hippocampal heterosynaptic plasticity. J Comput Neurosci 22:129–133.

Bi GQ, Poo MM (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 2:32–48.

Bringuier V, Chavane F, Glaeser L, Frégnac Y (1999) Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. Science 283:695–699.

Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J Comput Neurosci 8:183–208.

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.

Bush D, Philippides A, Husbands P, O'Shea M (2010) Reconciling the STDP and BCM models of synaptic plasticity in a spiking recurrent neural network. Neural Comput 22:2059–2085.

Christie BR, Abraham WC (1992) Priming of associative long-term depression in the dentate gyrus by theta frequency synaptic activity. Neuron 9:79–84.

Clopath C, Büsing L, Vasilaki E, Gerstner W (2010) Connectivity reflects coding: a model of voltage-based stdp with homeostasis. Nat Neurosci 13:344–352.

Davison AP, Brüderle D, Eppler J, Kremkow J, Muller E, Pecevski D, Perrinet

L, Yger P (2008) PyNN: a common interface for neuronal network simulators. Front Neuroinform 2:11.

Deisseroth K, Bito H, Schulman H, Tsien RW (1995) Synaptic plasticity: a molecular mechanism for metaplasticity. Curr Biol 5:1334–1338.

Destexhe A, Hughes SW, Rudolph M, Crunelli V (2007) Are corticothalamic "up" states fragments of wakefulness? Trends Neurosci 30:334–342.

Dudai Y (2004) The neurobiology of consolidations, or, how stable is the engram? Annu Rev Psychol 55:51–86.

Dunfield D, Haas K (2009) Metaplasticity governs natural experience-driven plasticity of nascent embryonic brain circuits. Neuron 64:240–250.

El Boustani S, Destexhe A (2009) A master equation formalism for macroscopic modeling of asynchronous irregular activity states. Neural Comput 21:46–100.

El Boustani S, Pospischil M, Rudolph-Lilith M, Destexhe A (2007) Activated cortical states: experiments, analyses and models. J Physiol Paris 101:99–109.

Elgersma Y, Fedorov NB, Ikonen S, Choi ES, Elgersma M, Carvalho OM, Giese KP, Silva AJ (2002) Inhibitory autophosphorylation of CaMKII controls PSD association, plasticity, and learning. Neuron 36:493–505.

Faas GC, Raghavachari S, Lisman JE, Mody I (2011) Calmodulin as a direct detector of Ca^{2+} signals. Nat Neurosci 14:301–304.

Frégnac Y (2002) Hebbian synaptic plasticity. In: Handbook of brain theory and neural networks (Arbib MA, ed), pp 512–522. Cambridge, MA: MIT.

Frégnac Y (1999) A tale of two spikes. Nat Neurosci 2:299–301.

Frégnac Y, Shulz D, Thorpe S, Bienenstock E (1988) A cellular analogue of visual cortical plasticity. Nature 333:367–370.

Fu YX, Djupsund K, Gao H, Hayden B, Shen K, Dan Y (2002) Temporal specificity in the cortical plasticity of visual space representation. Science 296:1999–2003.

Fu YX, Shen Y, Gao H, Dan Y (2004) Asymmetry in visual cortical circuits underlying motion-induced perceptual mislocalization. J Neurosci 24:2165–2171.

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

Gais S, Plihal W, Wagner U, Born J (2000) Early sleep triggers memory for early visual discrimination skills. Nat Neurosci 3:1335–1339.

Gerstner W, Kempter R, van Hemmen JL, Wagner H (1996) A neuronal learning rule for sub-millisecond temporal coding. Nature 383:76–81.

Giese KP, Fedorov NB, Filipkowski RK, Silva AJ (1998) Autophosphorylation at thr286 of the alpha calcium–calmodulin kinase II in LTP and learning. Science 279:870–873.

Gilson M, Burkitt A, van Hemmen LJ (2010) Stdp in recurrent neuronal networks. Front Comput Neurosci 4.pii:23.

Gold PE (2008) Protein synthesis inhibition and memory: formation vs amnesia. Neurobiol Learn Mem 89:201–211.

Graef IA, Mermelstein PG, Stankunas K, Neilson JR, Deisseroth K, Tsien RW, Crabtree GR (1999) L-type calcium channels and gsk-3 regulate the activity of nf-atc4 in hippocampal neurons. Nature 401:703–708.

Graupner M, Brunel N (2007) Stdp in a bistable synapse model based on CaMKII and associated signaling pathways. PLoS Comput Biol 3:e221.

Gu Z, Yan Z (2004) Bidirectional regulation of Ca^{2+}/calmodulin-dependent protein kinase II activity by dopamine D4 receptors in prefrontal cortex. Mol Pharmacol 66:948–955.

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

Harvey CD, Yasuda R, Zhong H, Svoboda K (2008) The spread of ras activity triggered by activation of a single dendritic spine. Science 321:136–140.

Huang YY, Colino A, Selig DK, Malenka RC (1992) The influence of prior synaptic activity on the induction of long-term potentiation. Science 255:730–733.

Iannella NL, Launey T, Tanaka S (2010) Spike timing-dependent plasticity as the origin of the formation of clustered synaptic efficacy engrams. Front Comput Neurosci 4.pii:21.

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.

Jia H, Rochefort NL, Chen X, Konnerth A (2010) Dendritic organization of sensory input to cortical neurons in vivo. Nature 464:1307–1312.

Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat Neurosci 10:100–107.

Kang S, Kitano K, Fukai T (2008) Structure of spontaneous up and down transitions self-organizing in a cortical network model. PLoS Comput Biol 4:e1000022.

Kasamatsu T, Pettigrew JD (1976) Depletion of brain catecholamines: failure of ocular dominance shift after monocular occlusion in kittens. Science 194:206–209.

Kumar A, Schrader S, Aertsen A, Rotter S (2008) The high-conductance state of cortical networks. Neural Comput 20:1–43.

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

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

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

Lisman J (1994) The CaM kinase II hypothesis for the storage of synaptic memory. Trends Neurosci 17:406–412.

Lisman J, Schulman H, Cline H (2002) The molecular basis of CaMKII function in synaptic and behavioural memory. Nat Rev Neurosci 3:175–190.

Lubenov EV, Siapas AG (2008) Decoupling through synchrony in neuronal circuits with propagation delays. Neuron 58:118–131.

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

Marre O, Yger P, Davison AP, Frégnac Y (2009) Reliable recall of spontaneous activity patterns in cortical networks. J Neurosci 29:14596–14606.

Mayford M, Wang J, Kandel ER, O'Dell TJ (1995) CaMKII regulates the frequency-response function of hippocampal synapses for the production of both LTD and LTP. Cell 81:891–904.

McGaugh JL (2000) Memory: a century of consolidation. Science 287:248–251.

Mehta MR, Quirk MC, Wilson MA (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. Neuron 25:707–715.

Mockett B, Coussens C, Abraham WC (2002) NMDA receptor-mediated metaplasticity during the induction of long-term depression by low-frequency stimulation. Eur J Neurosci 15:1819–1826.

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

Okamoto H, Ichikawa K (1994) Model for an enzymatic reaction-diffusion system realizing storage of graded information. Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics 50:1704–1707.

Pawlak V, Kerr JN (2008) Dopamine receptor activation is required for corticostriatal spike-timing-dependent plasticity. J Neurosci 28:2435–2446.

Pawlak V, Wickens JR, Kirkwood A, Kerr JN (2010) Timing is not everything: neuromodulation opens the stdp gate. Front Synaptic Neurosci 2:146.

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

Pi HJ, Otmakhov N, Lemelin D, De Koninck P, Lisman J (2010) Autonomous camkii can promote either long-term potentiation or long-term depression, depending on the state of t305/t306 phosphorylation. J Neurosci 30:8704–8709.

Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) The asynchronous state in cortical circuits. Science 327:587–590.

Romo R, Schultz W (1990) Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. J Neurophysiol 63:592–606.

Sejnowski TJ (1977) Statistical constraints on synaptic plasticity. J Theor Biol 69:385–389.

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

Seol GH, Ziburkus J, Huang S, Song L, Kim IT, Takamiya K, Huganir RL, Lee HK, Kirkwood A (2007) Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. Neuron 55:919–929.

Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci 18:3870–3896.

Shen W, Flajolet M, Greengard P, Surmeier DJ (2008) Dichotomous do-paminergic control of striatal synaptic plasticity. Science 321:848–851.

Shouval HZ, Castellani GC, Blais BS, Yeung LC, Cooper LN (2002) Converging evidence for a simplified biophysical model of synaptic plasticity. Biol Cybern 87:383–391.

Sjöström PJ, Turrigiano GG, Nelson SB (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 32:1149–1164.

Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. J Neurosci 13:334–350.

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

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

Steriade M (2001) Impact of network activities on neuronal properties in corticothalamic systems. J Neurophysiol 86:1–39.

Stickgold R, Hobson JA, Fosse R, Fosse M (2001) Sleep, learning, and dreams: off-line memory reprocessing. Science 294:1052–1057.

Tronson NC, Taylor JR (2007) Molecular mechanisms of memory recon-solidation. Nat Rev Neurosci 8:262–275.

van Rossum MC, Bi GQ, Turrigiano GG (2000) Stable hebbian learning from spike timing-dependent plasticity. J Neurosci 20:8812–8821.

Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. Science 287:1273–1276.

Vogels TP, Abbott LF (2005) Signal propagation and logic gating in net-works of integrate-and-fire neurons. J Neurosci 25:10786–10795.

Wang HX, Gerkin RC, Nauen DW, Bi GQ (2005) Coactivation and timing-dependent integration of synaptic potentiation and depression. Nat Neu-rosci 8:187–193.

Wang Y, Wu J, Rowan MJ, Anwyl R (1998) Role of protein kinase c in the induction of homosynaptic long-term depression by brief low frequency stimulation in the dentate gyrus of the rat hippocampus in vitro. J Physiol 513:467–475.

Yu X, Shouval HZ, Knierim JJ (2008) A biophysical model of synaptic plas-ticity and metaplasticity can account for the dynamics of the backward shift of hippocampal place fields. J Neurophysiol 100:983–992.

Zhang JC, Lau PM, Bi GQ (2009) Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. Proc Natl Acad Sci U S A 106:13028–13033.

Zhang L, Kirschstein T, Sommersberg B, Merkens M, Manahan-Vaughan D, Elgersma Y, Beck H (2005) Hippocampal synaptic metaplasticity re-quires inhibitory autophosphorylation of Ca^{2+}/calmodulin-dependent kinase II. J Neurosci 25:7697–7707.

Zhou Q, Tao HW, Poo MM (2003) Reversal and stabilization of synap-tic modifications in a developing visual system. Science 300:1953–1957.