

Reconciling the STDP and BCM Models of Synaptic Plasticity in a Spiking Recurrent Neural Network

Daniel Bush

d.bush@sussex.ac.uk

Andrew Philippides

andrewop@sussex.ac.uk

Phil Husbands

philh@sussex.ac.uk

Michael O'Shea

m.o-shea@sussex.ac.uk

*Centre for Computational Neuroscience and Robotics, University of Sussex,
Brighton, Sussex BN1 9QG, U.K.*

Rate-coded Hebbian learning, as characterized by the BCM formulation, is an established computational model of synaptic plasticity. Recently it has been demonstrated that changes in the strength of synapses *in vivo* can also depend explicitly on the relative timing of pre- and post-synaptic firing. Computational modeling of this spike-timing-dependent plasticity (STDP) has demonstrated that it can provide inherent stability or competition based on local synaptic variables. However, it has also been demonstrated that these properties rely on synaptic weights being either depressed or unchanged by an increase in mean stochastic firing rates, which directly contradicts empirical data. Several analytical studies have addressed this apparent dichotomy and identified conditions under which distinct and disparate STDP rules can be reconciled with rate-coded Hebbian learning. The aim of this research is to verify, unify, and expand on these previous findings by manipulating each element of a standard computational STDP model in turn. This allows us to identify the conditions under which this plasticity rule can replicate experimental data obtained using both rate and temporal stimulation protocols in a spiking recurrent neural network. Our results describe how the relative scale of mean synaptic weights and their dependence on stochastic pre- or postsynaptic firing rates can be manipulated by adjusting the exact profile of the asymmetric learning window and temporal restrictions on spike pair interactions respectively. These findings imply that previously disparate models of rate-coded autoassociative learning and temporally coded heteroassociative learning, mediated by symmetric and asymmetric connections respectively, can be implemented in a single network using a single plasticity rule. However, we also demonstrate that forms of STDP that can be reconciled with rate-coded Hebbian

learning do not generate inherent synaptic competition, and thus some additional mechanism is required to guarantee long-term input-output selectivity.

1 Introduction

Hebbian changes in the strength of neural connections based on correlations between pre- and postsynaptic activity are widely believed to mediate learning, memory, and the activity-dependent development of neural circuits (Hebb, 1949; Martin, Grimwood, & Morris, 2000). Empirical observations of synaptic plasticity are well characterized by the BCM model, which postulates the existence of a modification threshold that distinguishes between the long-term depression (LTD) of a synapse at lower firing rates and the long-term potentiation (LTP) of a synapse at higher firing rates (Lomo & Bliss, 1973; Bienenstock, Cooper, & Munro, 1982; Dudek & Bear, 1992; O'Connor, Wittenberg, & Wang, 2005a). More recently, it has been demonstrated that changes in the strength of synapses *in vivo* can also depend explicitly on the relative timing of pre- and postsynaptic action potentials (Markram, Lubke, Frotscher, & Sakmann, 1997; Bi & Poo, 1998; Bi, 2002; Kampa, Letzkus, & Stuart, 2007). Computational models of spike-timing-dependent plasticity (STDP) have shown that it can provide inherent stability or competition using only local synaptic variables (Abbott & Nelson, 2001; Gutig, Aharonov, Rotter, & Sompolinsky, 2003; Meffin, Besson, Burkitt, & Grayden, 2006; Morrison, Diesmann, & Gerstner, 2008). However, these properties rely on synaptic weights being either depressed or unchanged following an increase in presynaptic stimulation, which directly contradicts empirical data (Song, Miller, & Abbott, 2000; van Rossum, Bi, & Turrigiano, 2000; Senn, 2002). This apparent dichotomy has been the subject of several analytical studies, which have described distinct and disparate conditions under which STDP can be reconciled with the BCM formulation (Izhikevich & Desai, 2003; Burkitt, Meffin, & Grayden, 2004; Pfister & Gerstner, 2006). We aim to verify these results in a spiking recurrent neural network and, by manipulating each element of a standard computational STDP model in turn, provide a comprehensive description of the conditions under which this phenomenological plasticity rule can replicate experimental data obtained using both rate and temporal stimulation protocols.

The difficulty of empirically examining changes at the synaptic level means that STDP is not well characterized, and many computational interpretations of experimental data exist (Morrison et al., 2008). Phenomenological STDP rules have several distinct features that vary widely among modeling studies, including the exact profile of the asymmetric learning window, the nature of constraints placed on spike pair interactions, whether the scale of plasticity is dependent on current synaptic weight (multiplicative and additive STDP respectively), and whether the effects of multiple spike pairings sum linearly or nonlinearly (pair- and triplet-based rules

respectively). Previous analytical investigations regarding a reconciliation of STDP with the BCM formulation have examined only a small, distinct, and disparate number of this vast repertoire of possible implementations. For additive STDP, Izhikevich and Desai (2003) demonstrated that a nearest-neighbor spike pairing scheme, coupled with constraints on the ratios of parameters that define the asymmetric learning window, allows reconciliation with the BCM formulation. With multiplicative STDP, Burkitt et al. (2004) have demonstrated that an input-restricted spike pairing scheme generates the selective potentiation of inputs with elevated presynaptic firing rates. More recently, Pfister and Gerstner (2006) have demonstrated that triplet-based additive STDP can also replicate rate-coded plasticity data.

In contrast to these analytical studies, we investigate the emergent dynamics of a wide range of STDP implementations in a spiking recurrent neural network that corresponds to established autoassociative models of hippocampal mnemonic function (Burgess, 2007; Rolls, 2008). This represents an appropriate framework for this research for two principal reasons. First, the majority of neurobiological data regarding synaptic plasticity is obtained from the hippocampus (Andersen, Morris, Amaral, Bliss, & O'Keefe, 2007; Morris, 2007), and second, this region has been modeled using both rate-coded Hebbian learning and STDP to store and recall static and sequential activity patterns respectively (Marr, 1971; Hopfield, 1982; Amit & Mongillo, 2003; Lengyel, Kwag, Paulsen, & Dayan, 2005; Mongillo, Curti, Romani, & Amit, 2005; Samura & Hattori, 2005; Rolls & Kesner, 2006; Wagatsuma & Yamaguchi, 2007). The results presented here can therefore inform the development of a novel associative memory model that mediates the learning of both rate- and temporally coded activity patterns using a single plasticity rule. Within this context, manipulations of the plasticity rule or network model can also be directly related to systems-level mnemonic function.

We identify a range of conditions under which STDP can be reconciled with rate-coded Hebbian learning. For pair-based STDP, a dependence of mean synaptic weight on either pre- or postsynaptic firing rate can generally be produced if certain restrictions are placed on spike pair interactions and the profile of the asymmetric learning window. Furthermore, the dependence of mean synaptic weight on pre- or postsynaptic activity can be manipulated by adjusting the nature of temporal restrictions on spike pairing, and the position of the theoretical modification threshold postulated by the BCM formulation can be manipulated by adjusting the parameters that define the profile of the asymmetric learning window. Conversely, triplet-based STDP dictates that mean synaptic weight is dependent on postsynaptic firing rate regardless of the finer details of the plasticity rule, and the theoretical modification threshold cannot be manipulated in this case. These results support, unify, and extend the analytical findings of Izhikevich and Desai (2003), Burkitt et al. (2004), and Pfister and Gerstner (2006). We subsequently establish that this repertoire of STDP rules

can produce the rapid and selective potentiation of connections among neurons with concurrently elevated firing rates in a simple model of associative learning. These findings provide broad insight into the emergent dynamics generated by STDP in a spiking recurrent neural network and indicate that previous autoassociative network models of mnemonic function can be reconciled with spiking neural dynamics and spike-driven plasticity (Amit & Mongillo, 2003; Mongillo et al., 2005; Samura & Hattori, 2005). However, our results also indicate that synaptic competition is generally absent in these simulations, and a slow increase in the weight of all pre- or postsynaptic connections can therefore compromise the long-term input or output selectivity of single neurons in a manner analogous to previous rate-coded Hebbian learning rules (Abbott & Nelson, 2001; Desai, 2003). We discuss the implications of these findings for the modeling of both synaptic plasticity and systems-level models of hippocampal function.

2 Methods

There are two common computational models of rate-coded Hebbian learning. The most basic implementation, widely used in previous autoassociative network models of hippocampal mnemonic function, dictates that changes in synaptic strength are directly proportional to the product of pre- and postsynaptic firing rates (r_{ji}) scaled by a learning rate k (see equation 2.1; Hopfield, 1982; Burgess, 2007; Rolls, 2008):

$$\Delta w = kr_i r_j \quad (2.1)$$

However, this plasticity rule generates no competition between synaptic inputs, as any increase in synaptic weight produces an increase in postsynaptic firing rate, thus creating a positive feedback loop (Desai, 2003). The BCM model, which postulates the existence of a theoretical modification threshold (θ_m) that distinguishes between depression (at lower firing rates) and potentiation (at higher firing rates), was proposed to address this issue (see equation 2.2; Bienenstock, Cooper, & Munro, 1982):

$$\Delta w = kr_i r_j (r_i - \theta_m) \quad (2.2)$$

The value of θ_m is itself a supralinear function of pre- or postsynaptic activity, generating competition between synaptic inputs by making potentiation more difficult to achieve if the long-term average firing rate increases.

In order to identify conditions under which a standard computational STDP rule can be reconciled with rate-coded Hebbian learning, we use a spiking recurrent neural network inspired by previous autoassociative models of declarative memory function (Hopfield, 1982; Burgess, 2007; Rolls, 2008). More complex network models have previously demonstrated that rate-coded associative learning can be achieved using spike-driven

synaptic plasticity, but with a focus on systems-level function rather than the relationship between plasticity rule and emergent synaptic dynamics (Amit & Mongillo, 2003; Mongillo et al., 2005; Samura & Hattori, 2005). Conversely, this neural network model (and the rate-coded associative learning paradigm employed) are kept deliberately simple, for both elegant parsimony and to better identify the minimum necessary conditions required for STDP to produce an increase in mean synaptic weight with mean firing rate. Hence, many features of real neural networks that may complicate the observed synaptic dynamics (such as inhibitory input and short-term depression or facilitation) are omitted. However, the (minor) effects caused by manipulating each of the parameters that define the network model (i.e., the range of axonal delays, maximum synaptic weight, number of neurons and connectivity) are examined, and any significant findings discussed where appropriate.

2.1 The Network Model. Simulated neurons operate according to the Izhikevich (2004) spiking model, which dynamically calculates the membrane potential (v) and a membrane recovery variable (u) based on the values of four dimensionless constants (a , b , c , and d) and a dimensionless current input (I) according to equations 2.3 to 2.5:

$$v' = 0.04v^2 + 5v + 140 - u + I \quad (2.3)$$

$$u' = a(bv - u) \quad (2.4)$$

$$\text{if } v \geq 30 \text{ then } \begin{cases} v \rightarrow c \\ u \rightarrow u + d \end{cases} \quad (2.5)$$

This model can exhibit firing patterns of all known types of cortical neurons by variation of the magnitude of applied current and the parameters a – d . The values used to replicate spiking in a standard excitatory neuron are [$a = 0.02$, $b = 0.2$, $c = -65$, $d = 6$], which directs simulated neurons to fire single-action potentials at lower levels of stimulation and complex bursts (i.e., several spikes with interspike interval of less than 6 ms) that are realistic of pyramidal cells in the hippocampus at higher levels of stimulation (Ranck, 1973; Izhikevich, 2004; Frerking, Schulte, Wiebe, & Staubli, 2005).

Each simulated neuron has a randomly chosen axonal delay in the range [1 ms : D ms]. In the majority of simulations, the maximum axonal delay is set to $D = 5$ ms—this being realistic of the CA3 region (Miles, 1990)—although values of $D = [1, 10, 15, 20 \text{ ms}]$ are also examined to ascertain any effect on synaptic dynamics.

Unless otherwise specified, the neural network consists of $N_e = 100$ simulated excitatory neurons. However, for comparison, simulations with a value of $N_e = [20, 500]$ were also performed. The network is generally fully recurrently connected except for self-connections, although in

some simulations, the connectivity of the network is reduced such that the number of randomly assigned presynaptic connections to each excitatory neuron is proportional to 15% of the total number of excitatory neurons—this value being approximately realistic of the CA3 region (Rolls, 2008).

2.2 The Plasticity Model. Mathematically, with $s = t_{post} - t_{pre}$ being the time difference between pre- and postsynaptic spiking, the change in the weight of a synapse (Δw) according to a standard implementation of additive STDP can be calculated using equations 2.6 to 2.9 (Song et al., 2000; Bi, 2002; Izhikevich & Desai, 2003; Burkitt et al., 2004; Pfister & Gerstner, 2006):

$$\Delta w_+ = F(s) = P_+ \quad \text{for } s > 0 \quad (2.6)$$

$$\Delta w_- = F(s) = P_- \quad \text{for } s \leq 0 \quad (2.7)$$

$$P_+ = A_+ \left(1 - \frac{1}{\tau_+}\right)^s \quad (2.8)$$

$$P_- = A_- \left(1 - \frac{1}{\tau_-}\right)^{-s} \quad (2.9)$$

The parameters A_+ and A_- effectively correspond to the maximum possible change in the synaptic weight per spike pair, while τ_+ and τ_- denote the time constants of exponential decay for potentiation and depression increments, respectively (see Figure 1). In accordance with empirical data, coincident pre- and postsynaptic firing elicits maximal depression from all STDP implementations examined here (Debanne, Gähwiler, & Thompson, 1998). We examine a wide range of STDP parameter sets that include all combinations of the ratios $A_+ : A_-$ and $\tau_+ : \tau_-$. In the majority of simulations, the absolute values of A_{\pm} are scaled such that approximately 50 spike pairings (at interspike intervals corresponding to ~ 35 Hz pre- and postsynaptic firing) are sufficient to traverse the entire range of weight values, in accordance with empirical data (Bi & Poo, 1998).

A number of common spike pairing schemes are examined in this research, including the all-to-all scheme, which places no temporal restrictions on spike pair interactions; the input-restricted scheme, which dictates that each postsynaptic spike interacts with only one preceding but all subsequent presynaptic spikes; the output-restricted scheme, which dictates that each presynaptic spike interacts with only one preceding but all subsequent postsynaptic spikes; and the strict nearest-neighbor scheme, which dictates that each pre- or postsynaptic spike interacts with only the single closest opposing spike (Bi, 2002; Burkitt et al., 2004). The manner in which each of these spike pairing schemes mediates the update of P_{\pm} on afferent or efferent firing is described in Table 1.

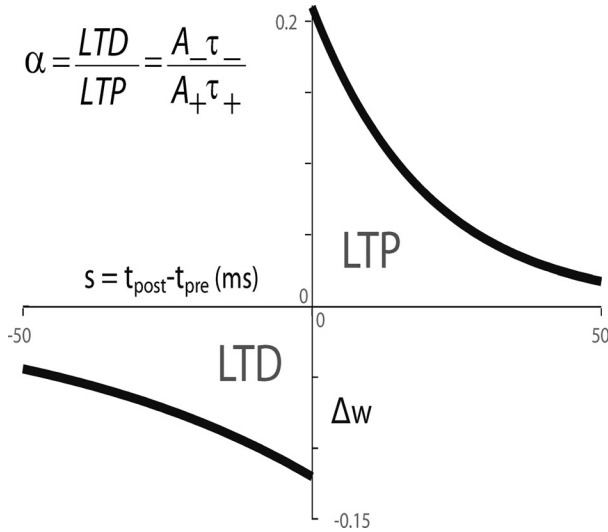


Figure 1: The asymmetric profile of the STDP learning window. The difference between the timing of pre- and postsynaptic firing ($s = t_{post} - t_{pre}$, ms) determines the degree and direction of the change in synaptic weight (Δw). The profiles of the exponential potentiation ($s > 0$) and depression ($s \leq 0$) windows, respectively, are defined by the maximum possible weight change per spike pair (A_{\pm}) and the time constants of decay (τ_{\pm}). The ratio of integrals over these windows is defined as $\alpha = A_- \tau_- / A_+ \tau_+$, such that a value of $\alpha > 1$ indicates an overall dominance of depression, while $\alpha < 1$ indicates an overall dominance of potentiation.

Table 1: Computational Details of Spike-Pairing Schemes.

Spike Pairing Scheme	Upon Afferent Firing, $P_+ \rightarrow$	Upon Efferent Firing, $P_+ \rightarrow$	Upon Efferent Firing, $P_- \rightarrow$	Upon Afferent Firing, $P_- \rightarrow$
Temporally unrestricted (all-to-all)	$P_+ + A_+$	-	$P_- + A_-$	-
Lax nearest neighbor	A_+	-	A_-	-
Strict nearest neighbor	A_+	0	A_-	0
Input restricted	A_+	-	$P_- + A_-$	0
Output restricted	$P_+ + A_+$	0	A_-	-

Several common forms of multiplicative STDP are also examined in this research (van Rossum et al., 2000; Gutig et al., 2003; Burkitt et al., 2004; Meffin et al., 2006). Equations 2.10 to 2.12 illustrate the manner in which the relevant value of P_+ or P_- is adjusted at each spike pairing event, and in each case the magnitude of weight dependency is scaled by a positive

constant f :

$$\Delta w_+ = P_+ e^{-f w_{ij}} \quad (2.10)$$

$$\Delta w_- = P_- f w_{ij} \quad (2.11)$$

$$\Delta w_- = P_- f w_{ij}^3. \quad (2.12)$$

For the triplet-based STDP rule, the absolute value of weight increase at each potentiation event is supplemented by a value equivalent to a slowly decay trace of the most recent depression increment:

$$\Delta w_+ = P_+ - \Delta w_- \left(1 - \frac{1}{\tau_{++}}\right)^s \quad (2.13)$$

In all simulations, the time constant for the decay of this additional potentiation term is $\tau_{++} = 20$ ms, in accordance with empirical data (Wang, Gerkin, Nauen, & Bi, 2005; Pfister & Gerstner, 2006).

When additive STDP is implemented, a hard limit is placed on the achievable strength of synapses, and initial weights are maintained continuously in the range $[0 : w_{\max}]$. While there is little clear biological basis to inform the choice of w_{\max} , it is known that recurrent synapses in the CA3 region are generally incapable of solely provoking postsynaptic activity (Kobayashi & Poo, 2004). In order to generate a postsynaptic spike in the simulated neurons employed here, a single 1 ms current application of $I = 16.5$ is required. The value assigned to the maximum weight limit in the majority of simulations therefore is $w_{\max} = 2.5$, although simulations are also run with values of $w_{\max} = [1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5]$ to assess any impact on emergent dynamics. When multiplicative STDP is implemented in the network, no upper limit is placed on the achievable strength of a synapse, and initial synaptic weights are assigned randomly in the range $[0 : w_{\text{init}}]$, with $w_{\text{init}} = 2.5$ in all simulations unless stated otherwise.

2.3 External Input. The first set of simulations described here aims to characterize the synaptic dynamics produced when all neurons in the network fire stochastically at an approximately equal rate, corresponding hypothetically to uncorrelated rate-coded input. Dimensionless current, sampled from a uniform random distribution in the range $[0 : I_{\max}]$, is applied to each simulated excitatory neuron at each 1 ms time step for a total of 1000 s. An examination of the neural dynamics confirms that this form of input generates a gaussian distribution of inter spike intervals (ISIs) in a single neuron (data not shown). Simulations are initialized with all synaptic weights assigned randomly in the range $[0 : w_{\max}]$ and repeated three times for each STDP implementation or level of external input to ensure that initial conditions have no effect on the equilibrium data described. Synaptic and

neural dynamics generally achieve an equilibrium state after approximately 100 s, with negligible subsequent fluctuations in mean weight and firing rate (except where discussed). To eliminate any effects of jitter, the stable excitatory mean weight and firing rate for each simulation are calculated as an average over the final 10 seconds of simulated time. Figures 2, 4, and 5 represent the average of these values over the three simulations performed for each network incarnation.

The second set of simulations described here aims to characterize the synaptic dynamics produced when a subset of neurons in the network fires at an elevated rate. This represents an abstract, parsimonious test of rate-coded associative learning that concurs with a large number of previous computational models of hippocampal mnemonic function (Hopfield, 1982; Burgess, 2007; Rolls, 2008). In these simulations, dimensionless current, sampled from a uniform random distribution in the range $[0 : I_{back}]$, generates a constant background firing rate of ~ 0.5 Hz (this value being realistic of the hippocampus) in 90% of the excitatory neurons in the network (Frerking et al., 2005; Rolls, 2008). Dimensionless current to the remaining 10% of excitatory neurons is sampled from a uniform random distribution in the range $[0 : I_{fore}]$, with the value of I_{fore} set to a different value for each simulation in order to generate mean foreground firing rates spanning a range of 0 Hz to ~ 70 Hz. Simulations are initialized with all synaptic weights assigned an equal value of $0.01 w_{max}$, proceed for 1000s, and are repeated three times for each STDP implementation and level of external input.

3 Results

Many computational implementations of STDP exist, but it has been demonstrated that forms which provide inherent competition or stability also generate emergent synaptic dynamics that directly contradict computational models and empirical observations of rate-based plasticity (Abbott & Nelson, 2001; Senn, 2002; Morrison et al., 2008). Three previous analytical studies have addressed this apparent dichotomy and described conditions under which STDP might be reconciled with rate-coded Hebbian learning rule (see equations 2.1 and 2.2). However, each of these studies employs distinct and disparate implementations of a standard computational STDP rule, and so it is unclear how the results might be related. Here we aim to verify each previous analytical result and, by examining a wider range of STDP implementations, unify these findings in order to describe a general set of conditions under which this plasticity rule can replicate rate-coded Hebbian learning in a spiking recurrent neural network. We then examine the synaptic dynamics produced by each STDP implementation in a simple model of rate-coded associative learning, where a subset of neurons fires at elevated stochastic firing rates.

3.1 Pair-Based Additive STDP. The additive STDP model dictates that potentiation and depression processes are independent of the current

synaptic weight. This has been demonstrated to provoke strong competition between inputs to a neuron at the cost of inherent stability and is characterized by a bimodal distribution of synaptic weights (Song et al., 2000; Gutig et al., 2003). Previous research has also shown that when the learning window exhibits an overall dominance of depression ($\alpha > 1$ in Figure 1) and spike pairing is temporally unrestricted (all-to-all), then additive STDP regulates postsynaptic firing rates by reducing mean synaptic weights in response to an increase in presynaptic activity (Song et al., 2000; Kempter, Gerstner, & van Hemmen, 2001). However, this homeostatic property directly contradicts computational models of rate-coded Hebbian learning and empirical observations of synaptic plasticity mediated by tetanic stimulation protocols (Senn, 2002; Morrison et al., 2008).

Conversely, it has been analytically demonstrated that STDP can be reconciled with the BCM formulation (see equation 2.2) under certain conditions, if potentiation dominates at shorter relative spike timings ($A_+ > A_-$) while depression dominates overall ($\alpha > 1$) and a temporally restricted (nearest-neighbor) spike pairing scheme is employed (Izhikevich & Desai, 2003). To establish whether this form of pair-based additive STDP could replicate rate-coded learning in a spiking recurrent neural network, we examined the stable mean weight of excitatory synapses at various mean rates of stochastic firing. Our results show that emergent synaptic dynamics that closely resemble the theoretical BCM modification curve are indeed generated (see Figure 2a). Furthermore, the position of the theoretical modification threshold (θ_m) can be adjusted by altering the specific value of α . This confirms that the constraints described by Izhikevich and Desai (2003)— $\alpha > 1$ and $A_+ > A_-$ with lax nearest-neighbor spike pairing—are sufficient to replicate the properties of rate-coded Hebbian learning in a spiking recurrent neural network.

However, further simulations using all combinations of the ratios $A_+ : A_-$ and $\tau_+ : \tau_-$ and several different spike pairing schemes demonstrate that these constraints are sufficient, but not necessary, to replicate rate-coded Hebbian learning. First, a dominance of potentiation at higher firing rates ($A_+ > A_-$) is not necessarily required, as an increase in mean synaptic weight with mean firing rate is still manifested when $A_+ = A_-$ and $\tau_+ < \tau_-$ (see Figure 2b). However, significant potentiation is incurred only when mean firing rates vastly exceed those commonly observed in the CA3 region (Frerking et al., 2005; Rolls, 2008). Second, an overall dominance of depression ($\alpha > 1$) is not necessarily required, as mean weight and firing rate are also positively correlated when the overall ratio between depression and potentiation processes is equal ($\alpha = 1$) or potentiation dominates overall ($\alpha < 1$). In these cases, however, adjusting the specific value of α does not modulate the position of the theoretical modification threshold (see Figure 2c).

Finally, qualitatively similar results are achieved with all temporally restricted (strict nearest-neighbor, input and output restricted) spike pairing

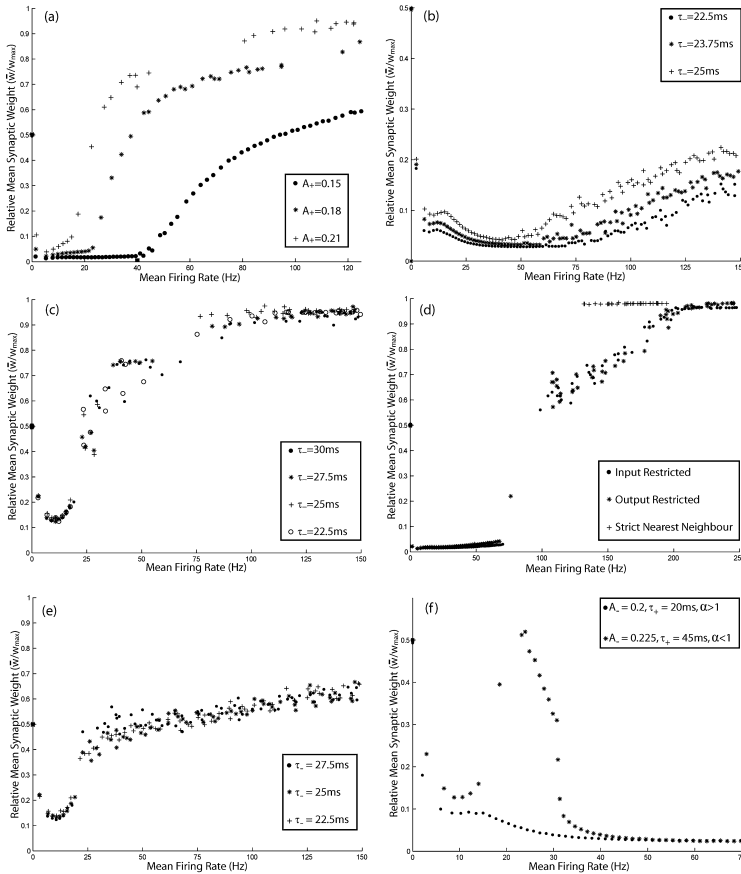


Figure 2: Typical emergent synaptic dynamics generated by pair-based additive STDP in a fully recurrently connected network of 100 excitatory spiking neurons with $w_{max} = 2.5$ and uncorrelated stochastic activity. Relative mean weight of excitatory synapses (w/w_{max}) is plotted against mean excitatory firing rate for (a) $\alpha > 1$ ($A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms) and lax nearest-neighbor spike pairing, which concurs with the conditions described by Izhikevich and Desai (2003). (b) $\alpha > 1$ but no dominance of potentiation at shorter ISIs ($A_+ = A_- = 0.15$, $\tau_+ = 20$ ms) and lax nearest-neighbor spike pairing. (c) $\alpha < 1$ ($A_+ = 0.15$, $A_- = 0.1$, $\tau_+ = 20$ ms) and lax nearest-neighbor spike pairing. (d) $\alpha > 1$ ($A_+ = 0.15$, $A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms), which concur with those described by Izhikevich and Desai (2003) and various temporally restricted spike pairing implementations. (e) $\alpha < 1$ ($A_+ = 0.15$, $A_- = 0.1$, $\tau_+ = 20$ ms) and temporally unrestricted (all-to-all) spike pairing. (f) A profile of the asymmetric learning window which dictates that the dominance of depression increases as ISIs are reduced ($A_+ < A_-$; $A_+ = 0.15$, $\tau_- = 20$ ms) and lax nearest-neighbor spike pairing.

schemes (see Figure 2d), although the transition between heterosynaptic depression at low firing rates and significant potentiation at higher firing rates is much more abrupt, such that a small increase in the level of external input provokes a significant increase in mean synaptic weights and firing rates (due to recurrent excitation). Thus, the lax nearest-neighbor STDP model used by Izhikevich and Desai (2003) is not necessary to generate an increase in mean synaptic weight with mean firing rate. Conversely, with temporally unrestricted (all-to-all) spike pair interactions, rate-coded Hebbian learning can only be replicated if $\alpha < 1$ (see Figure 2e). In this case, the scale of mean weights is generally much lower, and the position of the theoretical modification threshold (θ_m) cannot be manipulated (as observed for temporally restricted spike pairing schemes with $\alpha < 1$; see Figure 2c). All the results described here are qualitatively replicated with all absolute values of w_{\max} and the parameters A_{\pm} and τ_{\pm} examined, although the quantitative scale of relative mean weights generated in these simulations can differ significantly (data not shown).

The results of these simulations suggest that the emergent synaptic dynamics generated by pair-based additive STDP with temporally restricted spike pairing are determined by the relative scale of potentiation and depression processes evaluated around the mean interspike interval (ISI). As an increase in external excitatory input causes mean firing rates in the network to increase, the temporal restrictions on spike pairing effectively reduce the relative spike timing values that dictate plasticity, and thus the effective learning window moves closer to the y -axis (see Figure 3). Under these conditions, the properties of rate-coded Hebbian learning can generally be replicated if a dominance of depression at longer ISIs decreases as the mean ISI is reduced by external input (see Figures 2a, 2b, and 2d). Conversely, if the profile of the asymmetric learning window is set such that the dominance of depression increases as the mean ISI is reduced, with an overall value of either $\alpha > 1$ or $\alpha \leq 1$, then a homeostatic decrease in mean weights at higher firing rates analogous to that described in previous research can be produced (see Figure 2f; Song et al., 2000; Bi, 2002). The concept of an effective learning window supports the findings of Izhikevich and Desai (2003) and provides a broader conceptual framework for predicting the emergent synaptic dynamics generated by pair-based additive STDP.

3.2 Pair-Based Multiplicative STDP. Multiplicative STDP models dictate that the degree of potentiation or depression generated by a spike pairing event is dependent on the current synaptic weight, as indicated by several empirical studies of plasticity in the hippocampus (Bi & Poo, 1998; Bi, 2002; Debanne, Gahwiler, & Thompson, 1999; Wang et al., 2005). This form of synaptic plasticity is inherently stable and generates a unimodal weight distribution that approximates that observed in biological preparations (van Rossum et al., 2000; Gutig et al., 2003; Burkitt et al., 2004; Barbour,

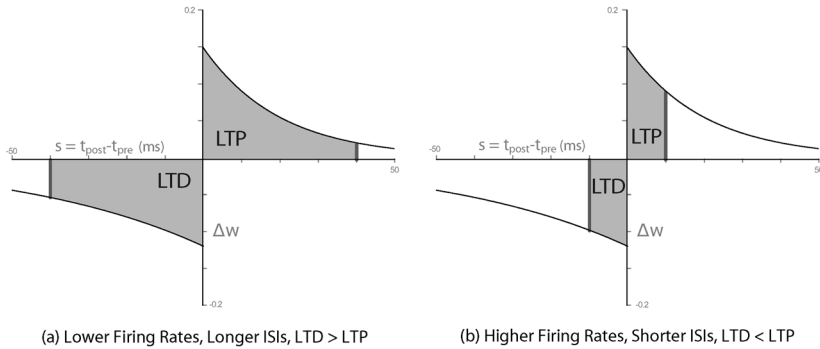


Figure 3: The effective learning window, which dictates the emergent dynamics of STDP when a temporally restricted spike pairing scheme is employed. As firing rates increase, spike pairings are restricted to shorter ISIs, and thus the effective ratio of integrals over the depression and potentiation windows is changed. Subsequently, the relative values of A_{\pm} primarily dictate plasticity at high rates (short ISIs), while the relative values of τ_{\pm} primarily dictate plasticity at low rates (long ISIs). The profile of the asymmetric learning window illustrated here corresponds to that described by Izhikevich and Desai (2003) and is depicted in Figure 2a ($A_+ = 0.15$, $A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms).

Brunel, Hakim, & Nadal, 2007; Standage, Jalil, & Trappenberg, 2007). However, previous research has shown that synaptic weights are generally unaffected by mean firing rates under these conditions, which contradicts empirical data and computational models of rate-coded Hebbian learning (van Rossum et al., 2000; Gutig et al., 2003; Morrison et al., 2008).

Several different multiplicative STDP implementations have been examined in previous computational studies, each of which adjusts the absolute scale of depression based on current synaptic weight (van Rossum et al., 2000; Gutig et al., 2003; Meffin et al., 2006). Burkitt et al. (2004) have analytically demonstrated that these multiplicative depression rules can generally replicate rate-coded Hebbian learning if an input-restricted spike pairing scheme is employed. However, to the best of our knowledge, no studies have examined a multiplicative STDP rule that solely adjusts the absolute scale of potentiation. We therefore repeated the simulations described in the previous section using two common multiplicative depression rules and one novel multiplicative potentiation rule, and again examined the relationship between mean firing rate and mean synaptic weight.

Our findings demonstrate that an increase in mean synaptic weight with mean firing rate is generally produced by each multiplicative STDP rule examined if the profile of the asymmetric learning window dictates that the dominance of depression decreases with mean ISI and any temporally

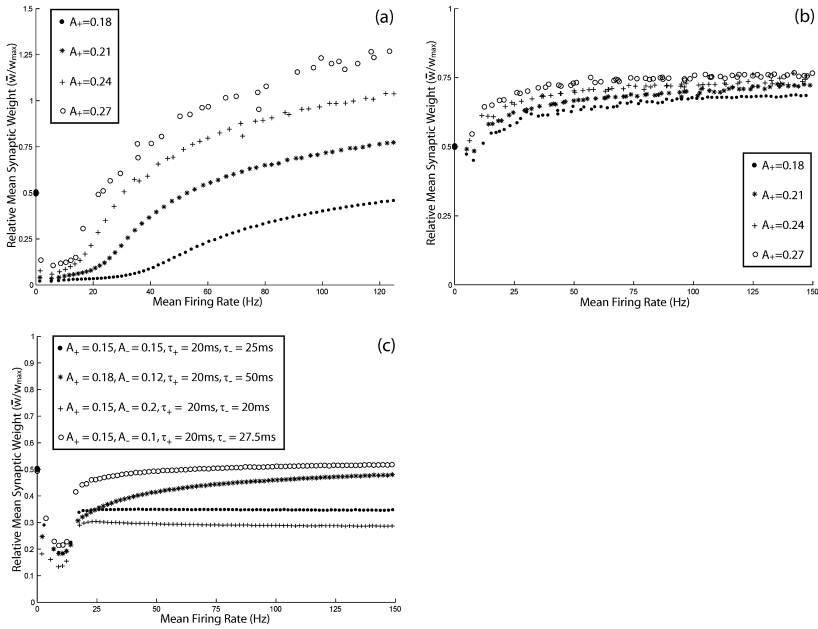


Figure 4: Typical emergent synaptic dynamics generated by pair-based multiplicative STDP rules in a fully recurrently connected network of 100 excitatory spiking neurons with uncorrelated stochastic activity. Relative mean weight (\bar{w}/w_{\max} , calculated for a value of $w_{\max} = 2.5$ to enable comparison with additive STDP data) is plotted against mean excitatory firing rate for (a) a linear dependence of potentiation upon initial synaptic weight with $f = 0.5$, $\alpha > 1$ ($A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms) and lax nearest-neighbor spike pairing; (b) a cubic dependence of depression upon initial synaptic weight with $f = 0.2$, $\alpha > 1$ ($A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms) and output-restricted spike pairing; (c) a linear dependence of depression upon initial synaptic weight with $f = 0.58$ and input-restricted spike pairing. Qualitatively similar results are observed for all temporally restricted spike pairing schemes with the multiplicative STDP rules illustrated here.

restricted spike pairing scheme is employed (see Figure 4). Conversely, there is no significant increase in mean weight at higher firing rates when no dominance of potentiation exists at shorter ISIs (see Figure 4c). For temporally unrestricted spike pairing, rate-coded Hebbian learning is replicated only when $\alpha < 1$ (data not shown), which also concurs with results obtained using additive STDP (see Figure 2e). These findings support the analytical results of Burkitt et al., (2004)—that multiplicative STDP can replicate rate-coded Hebbian learning if an input-restricted spike pairing scheme is employed—and demonstrate that this result can be extended to

include any temporally restricted spike pairing scheme. Interestingly, the profile of the asymmetric learning window implemented in that research ($A_+ = A_-$ and $\tau_+ < \tau_-$) dictates a decrease in the dominance of depression with the mean ISI, which we have demonstrated to be necessary for STDP to replicate rate-coded Hebbian learning (Burkitt et al., 2004). Thus, the concept of an effective learning window illustrated in Figure 3 can be used to predict the emergent synaptic dynamics produced by both pair-based additive and multiplicative STDP.

Clear differences in emergent synaptic dynamics exist among the different forms of multiplicative STDP examined here. Unimodal synaptic weight distributions are generated by all multiplicative rules, but only the multiplicative potentiation rule allows the position of the theoretical modification threshold and overall scale of synaptic weights to be significantly modulated by adjusting the parameters A_{\pm} and τ_{\pm} (see Figure 4a). This produces mean weight curves that are qualitatively similar to those generated by additive STDP, albeit with a soft upper weight bound. Conversely, the multiplicative depression rules generate a soft lower weight bound, which is observed in simulations with all STDP parameter sets and spike pairing schemes (see Figures 4b and 4c). Synaptic weights are not heterosynaptically depressed (i.e., saturate at the lower bounds) under any conditions, a narrower range of weight values is traversed, and the position of the theoretical modification threshold cannot be adjusted. All the results described here are qualitatively replicated with all absolute values of f and the parameters A_{\pm} and τ_{\pm} examined, although the quantitative scale of relative mean weights generated in these simulations can differ significantly (data not shown).

3.3 Triplet-Based STDP. The majority of empirical studies of STDP have used single pairs of pre- and postsynaptic spikes delivered at low frequencies (Bi & Poo, 1998). However, recent investigations have demonstrated that more complex triplets or quadruplets of pre- and postsynaptic spikes result in a nonlinear integration of modular potentiation and depression processes (Froemke & Dan, 2002; Wang et al., 2005). In the hippocampus, potentiation events effectively override any depression incurred at a synapse within a short temporal window (< 70 ms). The pair-based STDP rules examined thus far are unable to replicate the results of these experiments, as sequential potentiation and depression events sum linearly. Furthermore, it has been analytically demonstrated that a triplet-based STDP rule, which incorporates a short-term dominance of potentiation, can be reconciled with the BCM formulation of rate-coded Hebbian learning (Pfister & Gerstner, 2006).

Our simulations demonstrate that the short-term dominance of potentiation produced by triplet-based STDP generates an increase in mean synaptic weight with mean firing rate regardless of the finer details of the plasticity rule (see Figure 5). However, the position of the theoretical modification

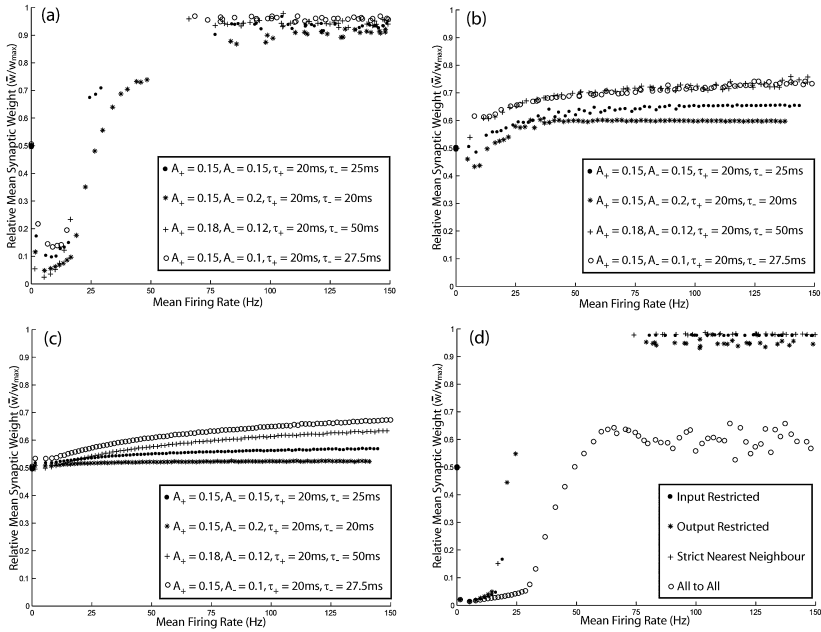


Figure 5: Typical emergent synaptic dynamics generated by triplet-based STDP rules in a fully recurrently connected network of 100 excitatory spiking neurons with uncorrelated stochastic activity. Relative mean weight (w/w_{max} , calculated with $w_{max} = 2.5$ for multiplicative STDP data) is plotted against mean excitatory firing rate for (a) additive STDP with $w_{max} = 2.5$ and lax nearest-neighbor spike pairing; (b) a cubic dependence of depression on initial synaptic weight, $f = 0.58$ and lax nearest-neighbor spike pairing; (c) a linear dependence of depression on initial synaptic weight, $f = 5$ and lax nearest-neighbor spike pairing; (d) additive STDP with $w_{max} = 2.5$, a value of $\alpha > 1$ ($A_+ = 0.15, A_- = 0.12, \tau_+ = 20\text{ms}, \tau_- = 50\text{ms}$) and various spike pairing schemes.

threshold (θ_m) cannot be manipulated by adjusting the profile of the asymmetric learning window (A_{\pm} and τ_{\pm}) under these conditions (see Figures 5a–5c). The differences in emergent dynamics generated by pair-based additive and multiplicative STDP are qualitatively similar for triplet-based rules—the range of mean weight values traversed being much narrower for multiplicative depression rules (see Figures 5a–5c). Interestingly, the position of the theoretical modification threshold (θ_m) is consistently higher, and the asymptotic mean weight consistently lower, for all-to-all spike pairing with $\alpha > 1$ (see Figure 5d). These findings provide support for the analytical results of Pfister and Gerstner (2006)—that an additive triplet-based STDP rule based on empirical data from the hippocampus can be reconciled with the BCM formulation—and extend those results to a broader repertoire of

triplet-based STDP models. All the results described here are qualitatively replicated with all absolute values of w_{\max}, f and parameters A_{\pm} and τ_{\pm} examined, although the quantitative scale of relative mean weights generated in these simulations can differ significantly (data not shown).

3.4 Rate-Coded Associative Learning Simulations. The simulations described thus far deviate significantly from the neural activity that is generally observed in computational models of declarative memory function and the CA3 region, as each synapse experiences pre- and postsynaptic firing of an approximately constant and equal rate. In order to further elucidate the emergent synaptic dynamics generated by STDP in this context, we next performed simulations in which the mean firing rate in the majority of (background) neurons was held fixed at ~ 0.5 Hz (this being realistic of the hippocampus), while the mean firing rate in the remaining (foreground) neurons was systematically varied (Frerking et al., 2005). The aim of these simulations is to more specifically assess whether STDP is compatible with rate-coded associative learning, which is generally mediated by elevating firing rates in a subset of excitatory neurons such that connections among these neurons are selectively potentiated while all other connections in the network are heterosynaptically depressed or remain unchanged (Hopfield, 1982; Burgess, 2007; Rolls, 2008). In analyzing the synaptic dynamics observed in these simulations, we divide connections in the network into different subgroups depending on whether the pre- and postsynaptic neuron fires at foreground or background rates. We then examine the stable mean weight of each subset of synapses following 1000 s of simulated activity with different foreground firing rates.

For pair-based additive and multiplicative STDP implementations, the strength of connections between background neurons is generally depressed or remains unchanged—most likely because the low pre- and postsynaptic firing rates generate very few effective spike pairings—while the mean weight of connections between foreground neurons adheres to the synaptic dynamics described above, with the same dependence on the profile of the asymmetric learning window and additive or multiplicative STDP rule employed (see Figure 6). More interestingly, the strength of connections between background and foreground neurons is generally affected in one of two ways, depending on the spike pairing scheme employed: either following a mean synaptic weight curve similar to that of the foreground synaptic weights, or (like background connections) being heterosynaptically depressed at all mean firing rates. For the lax nearest-neighbor and output-restricted spike pairing schemes, the strength of connections with elevated postsynaptic firing rate is significantly potentiated (see Figures 6a–6c). Conversely, for the input-restricted spike pairing schemes, connections with elevated presynaptic firing rate are significantly potentiated (see Figures 6d–6e). Finally, for the strict nearest-neighbor spike pairing scheme,

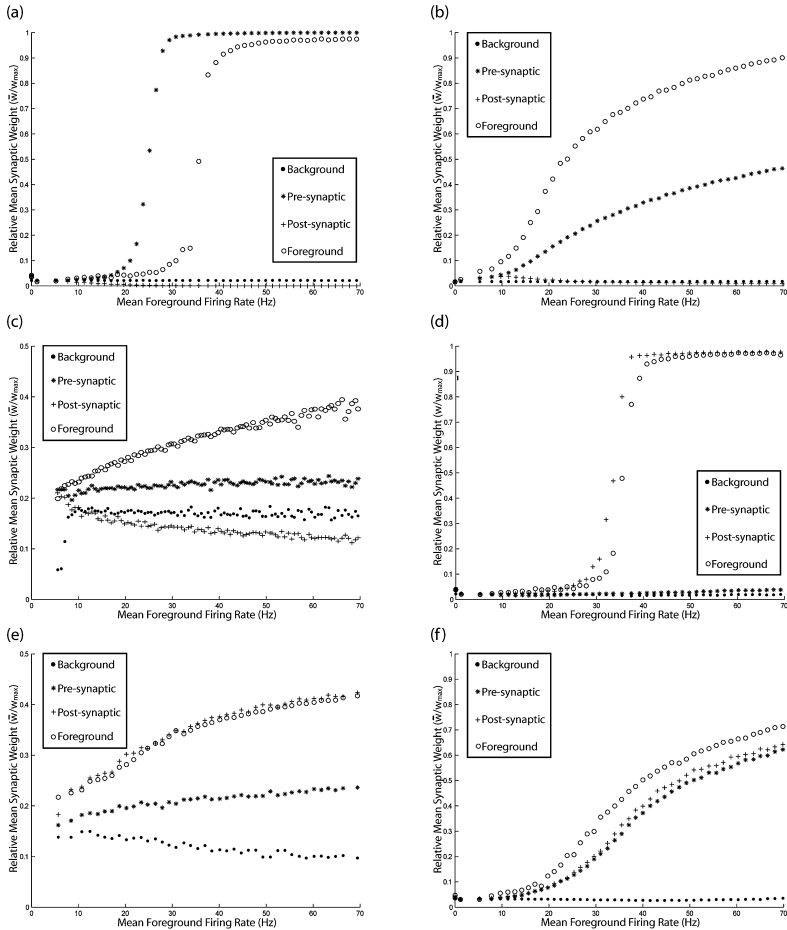


Figure 6: Typical emergent synaptic dynamics generated by pair-based STDP rules in a fully recurrently connected network of 100 excitatory spiking neurons with a constant background firing rate of ~ 0.5 Hz and various foreground firing rates. Data are illustrated for $\alpha > 1$ ($A_+ = 0.18$, $A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms) with (a) additive STDP and lax nearest-neighbor spike pairing; (b) a linear dependence of potentiation on initial synaptic weight, $f = 0.5$ and output-restricted spike pairing; (c) a cubic dependence of depression on initial synaptic weight, $f = 0.2$ and lax nearest-neighbor spike pairing; (d) additive STDP and input-restricted spike pairing; (e) a linear dependence of depression on initial synaptic weight with $f = 0.58$ and input-restricted spike pairing; and (f) a linear dependence of potentiation on initial synaptic weight and strict nearest-neighbor spike pairing. The stable relative mean synaptic weight of foreground, background, pre- and postsynaptic connections is plotted against the mean firing rate in foreground excitatory neurons.

synaptic connections with either elevated pre- or postsynaptic firing rate are potentiated at higher mean firing rates (see Figure 6f).

Critically, these results demonstrate that the synaptic weight dynamics produced by pair-based STDP differ in their dependency on pre- or postsynaptic firing rate according to the exact nature of temporal restrictions on spike pairing. Input-restricted spike pairing creates a dependence of mean synaptic weight on presynaptic firing rate, output-restricted and lax nearest-neighbor spike pairing create a dependence of mean synaptic weight on postsynaptic firing rate, and strict nearest-neighbor spike pairing creates a dependence of mean synaptic weight on either pre- or postsynaptic firing rate. This concurs with the analytical expressions derived by Burkitt et al. (2004), providing further support for the analytical reconciliation of STDP and rate-coded Hebbian learning described in that research, and demonstrates that those results also apply to pair-based additive STDP rules and more complex neural dynamics. Conversely, for temporally unrestricted spike pairing, the significant potentiation of foreground connections is achieved only for values of $\alpha < 1$, and in this case, mean synaptic weight is dependent on either pre- or postsynaptic firing rate (data not shown).

Further simulations demonstrate that mean synaptic weights are dependent on postsynaptic firing rate for triplet-based STDP with all temporally restricted spike pairing schemes, regardless of the finer details of the plasticity rule (see Figure 7). Hence, in contrast to pair-based STDP, the strength of connections with elevated postsynaptic firing rate is significantly potentiated when either input-restricted or strict nearest-neighbor spike pairing is employed (see Figures 7a and 7b). Furthermore, in contrast to pair-based STDP, significant potentiation of synapses with elevated postsynaptic firing rate also proceeds when the profile of the asymmetric learning window dictates a dominance of depression at shorter ISIs (see Figure 7c). Uniquely for triplet-based STDP, temporally unrestricted spike pairing creates a dependence of mean synaptic weight on either pre- or postsynaptic firing rate (see Figure 7d), as observed for pair-based STDP rules. The qualitative nature of all results described here (for both pair- and triplet-based STDP) is replicated with all values of maximum axonal delay (D), maximum weight limit (w_{\max}), damping factor (f), number of simulated neurons in the network (N_e), and recurrent connectivity (100% or 15%) examined (see section 2), although the quantitative scale of mean synaptic weights can differ significantly (data not shown).

An examination of the synaptic weight dynamics that proceed during these associative learning simulations demonstrates that the rate of synaptic weight change produced by all STDP implementations is generally a product of the firing rate of both pre- and postsynaptic neurons (data not shown). Hence, although the potentiation of foreground connections is limited by the fact that the temporal order of spike pairings must continually alternate in order to generate strong bidirectional connections, the mean weight of these synapses increases much more quickly than those with

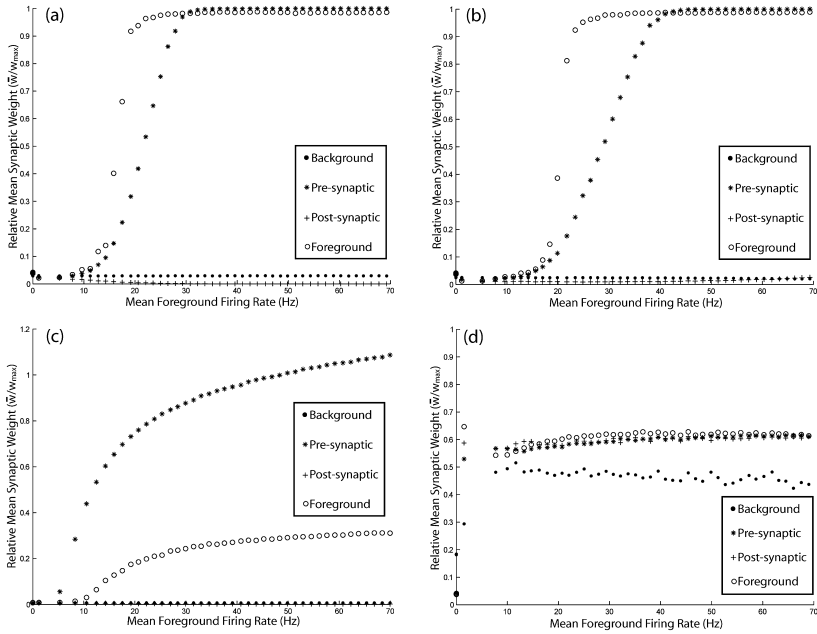


Figure 7: Typical emergent synaptic dynamics generated by triplet-based STDP rules in a fully recurrently connected network of 100 excitatory spiking neurons with a constant background firing rate of ~ 0.5 Hz and various foreground firing rates. Data are illustrated for (a) additive STDP with $\alpha > 1$ ($A_+ = 0.18$, $A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms) and input-restricted spike pairing; (b) additive STDP with $\alpha > 1$ ($A_+ = 0.18$, $A_- = 0.12$, $\tau_+ = 20$ ms, $\tau_- = 50$ ms) and strict nearest-neighbor spike pairing; (c) a linear dependence of potentiation on initial synaptic weight with $f = 1$, a dominance of depression at shorter ISIs ($A_+ = 0.1$, $A_- = 0.2$, $\tau_+ = 20$ ms, $\tau_- = 20$ ms) and output-restricted spike pairing; (d) a cubic dependence of depression on initial synaptic weight with $f = 0.2$, $\alpha > 1$ ($A_+ = 0.15$, $A_- = 0.15$, $\tau_+ = 20$ ms, $\tau_- = 25$ ms) and temporally unrestricted (all-to-all) spike pairing. The stable relative mean synaptic weight of foreground, background, and pre- and postsynaptic connections is plotted against the mean firing rate in foreground excitatory neurons.

elevated pre- or postsynaptic firing rate alone. Subsequently, connections with elevated input and output firing rate are selectively potentiated in the short term, which concurs with empirical findings, although the mean weight of synapses with elevated input or output firing rate only can approach—and in some cases exceed (see Figures 6b to 6f, 7a, and 7b)—the mean weight of these foreground connections over longer periods. This indicates that some additional mechanism that can stimulate competition between synapses over a longer timescale will be required to allow the efficient

operation of an autoassociative network model that uses these forms of STDP.

4 Discussion

We have identified a broad set of conditions sufficient for a standard computational STDP rule to replicate rate-coded Hebbian learning when pre- and postsynaptic activity is stochastic in a prototypical model of associative memory—the Hopfield network. Specifically, we have demonstrated that pair-based STDP produces an increase in mean synaptic weight with mean firing rate if a temporally restricted spike pairing scheme is combined with an asymmetric learning window that dictates a decrease in the dominance of depression as mean ISIs are reduced (see Figure 3). Adjusting the specific nature of spike pairing restrictions subsequently determines whether mean synaptic weights are primarily dependent on pre- or postsynaptic firing rate and, for additive and multiplicative potentiation STDP rules, the position of the theoretical modification threshold (θ_m) can be altered by manipulating the profile of the asymmetric learning window (see Figures 2a, 4a, 6, and 7). A dependence of mean weight on postsynaptic firing rate is generally produced by triplet-based STDP regardless of the finer details of the plasticity rule, but the position of the theoretical modification threshold (θ_m) cannot be adjusted in this case (see Figure 6). These results support, unify, and extend previous analytical findings pertaining to a reconciliation of STDP with the BCM formulation of rate-coded Hebbian learning (Izhikevich & Desai, 2003; Burkitt et al., 2004; Pfister & Gerstner, 2006).

Our findings also demonstrate that this repertoire of phenomenological plasticity rules can rapidly potentiate synaptic connections among neurons with concurrently elevated stochastic firing rates in a manner analogous to rate-coded Hebbian learning. Similarly, several more complex models have previously examined autoassociative network function using spike-driven synaptic plasticity (Amit & Mongillo, 2003; Mongillo et al., 2005; Samura & Hattori, 2005). Although the neural network used here has been kept deliberately simple, it is interesting to appraise the implications of our results for more general models of mnemonic function, where synaptic plasticity models are generally assessed on the basis of three criteria: the capacity for modulation, competition, and biological realism (Abbott & Nelson, 2001; Morrison et al., 2008). In terms of modulation, we have demonstrated that synaptic dynamics can be most significantly manipulated when pair-based additive or multiplicative potentiation rules are implemented (see Figures 2 and 4). These forms of STDP also produce the greatest distinction between the mean weight of foreground and background connections in rate-coded learning simulations (see Figures 6 and 7), and a bimodal separation of synaptic weights is generally considered preferable for efficient autoassociative network function (Fusi, Annunziato, Badoni, Salamon, & Amit, 2000; Rolls, 2008). The synaptic dynamics produced under these conditions

correspond closely to those described by the BCM formulation, allowing the scale of mean synaptic weights and their dependence on pre- or postsynaptic firing rate to be adjusted by manipulating the exact parameters of the STDP model.

While we have not specifically assessed the level of synaptic competition generated by each form of STDP examined here, it is possible to infer certain implications from the results of our simulations. For pair-based additive STDP with temporally restricted spike pairing, for example, we have demonstrated that synaptic weights saturate unimodally around the upper bounds at higher firing rates when $A_+ > A_-$ or $\alpha < = 1$ (see Figures 2a, 2c, and 2d). A similar result is observed, regardless of the profile of the asymmetric learning window, for triplet-based additive STDP with temporally restricted spike pairing (see Figures 5a and 5d); and unimodal synaptic weight distributions are generated by multiplicative STDP rules at all firing rates (see Figures 4, 5b, and 5c). This implies that strong bidirectional connections, which are considered essential for efficient rate-coded autoassociative network function, can develop under these conditions (Burgess, 2007; Rolls, 2008). However, unimodal weight distributions are also believed to indicate a lack of synaptic competition and inability to maintain learned associations once an encoding stimulus is removed (van Rossum et al., 2000; Gutig et al., 2003; Meffin et al., 2006; Billings & van Rossum, 2009). The results of our rate-coded associative learning simulations support the notion that inherent synaptic competition is absent under these conditions, as all input or output connections of a single neuron that is firing at an elevated rate can saturate unimodally over longer periods (see Figures 6 and 7).

That these forms of STDP should lack inherent synaptic competition is not surprising, considering the wealth of literature regarding the stability versus plasticity dilemma and positive feedback issues that face the rate-coded Hebbian learning rules we aim to emulate (Grossberg, 1987; Abbott & Nelson, 2001). Fortunately, several computational methods of introducing competition into the Hebbian framework exist and can easily be integrated into the STDP rules described here. For example, a process of synaptic scaling should retain the idealized weight matrix that develops during the first few seconds of associative learning simulations, while preventing a significant increase in pre- or postsynaptic weights, and is supported by a wealth of biological data (Desai, 2003; Turrigiano & Nelson, 2004; Stellwagen & Malenka, 2006). Elsewhere, the position of the theoretical BCM modification threshold might be manipulated by adjusting the profile of the asymmetric learning window, such that an increase in pre- or postsynaptic weights is not manifested at the firing rate that exists in active neurons following the (short) initial learning period (Bienenstock et al., 1982; Benuskova & Abraham, 2007; Abraham, 2008). Within this context, it would be interesting to examine whether this plasticity rule could be used to develop input selectivity when the position of the theoretical modification

threshold (θ_m) is a time-averaged function of pre- or postsynaptic activity, as described in the original BCM formulation (Bienenstock et al., 1982). Conversely, for pair-based additive STDP with temporally restricted spike pairing when an overall dominance of potentiation does not exist at low ISIs (i.e. $A_+ = A_-$), mean synaptic weights do not exceed a value of $w_{\max}/2$ in our simulations (see Figures 2b and 3b). Similar results are generally observed for temporally unrestricted spike pair interactions (see Figures 3d, 5d, and 6d). This suggests that some form of inherent synaptic competition does exist under these conditions, although this may prevent the development of strong bidirectional connections and significant potentiation at firing rates that are realistic of the hippocampus.

While the phenomenological plasticity models presented here are not intended to reproduce the full complexity and heterogeneity of empirical findings, it is useful to appraise these results with regard to biological data (Debanne et al., 1999; O'Connor, Wittenberg, & Wang, 2005b). The properties of synaptic plasticity vary widely throughout the cortex, and hence modeling studies such as this must be constrained to empirical data from a single region and the requirements for efficiently modeling the postulated mnemonic function of that region (Pfister & Gerstner, 2006). Interestingly, the profile of the asymmetric learning window that has been demonstrated to replicate rate-coded Hebbian learning most accurately in this research ($A_+ > A_-$ and $\alpha > 1$) concurs with those observed in the hippocampus (Bi & Poo, 1998; Wang et al., 2005). However, data pertaining to the nature of spike pair interactions have been obtained only from the visual cortex (Froemke, Tsay, Raad, Long, & Dan, 2006). There is also little consensus among experimental or computational models of the BCM rule with regard to its dependency on pre- or postsynaptic firing rate, although the plasticity rule presented here can effectively produce either (Standage et al., 2007; Toyozumi, Pfister, Aihara, & Gerstner, 2007). Clearly, only triplet-based STDP rules can replicate empirical data pertaining to the nonlinear integration of potentiation and depression processes with more complex spike trains (Wang et al., 2005). However, recent empirical data indicate a role for the nature of postsynaptic firing in dictating plasticity processes and a broad dichotomy between LTD at low (< 5 Hz) and LTP at higher (> 5 Hz) mean firing rates regardless of the temporal order of spike pairings—both findings that our model cannot replicate (Wittenberg & Wang, 2006). Biological data also suggest that the scale of depression processes is dependent on current synaptic weight (Bi & Poo, 1998; Debanne et al., 1999; Wang et al., 2005). Each of the multiplicative STDP rules examined here generally produces a unimodal weight distribution that adheres to that observed *in vivo*, while additive STDP generally produces a bimodal separation of synaptic weights. However, it is important to note that empirical recordings most likely neglect silent synapses and might therefore be accounted for by an additive STDP rule with a gaussian distribution of maximum weight limits. Furthermore, recent empirical data indicate that bidirectional changes

in synaptic strength in the hippocampus are discrete, all-or-none transitions between heterogeneous high and low states (O'Connor et al., 2005b).

In summary, this research has presented a variety of STDP rules that can rapidly rearrange synaptic weights in order to reflect rate-coded input correlations in a spiking recurrent neural network. If we can assume that these STDP rules are also able to encode temporal input correlations, then the findings described here provide the foundations for a novel associative network model that can integrate classic models of declarative memory function that store static, sparse, rate-coded activity patterns via autoassociation with more contemporary models of sequence learning that store dynamic theta-coded spike timing patterns via heteroassociation (Marr, 1971; Hopfield, 1982; Amit & Mongillo, 2003; Lengyel et al., 2005; Mongillo et al., 2005; Samura & Hattori, 2005; Burgess, 2007; de Almeida, Idiart, & Lisman, 2007; Wagatsuma & Yamaguchi, 2007; Rolls, 2008). The efficient operation of such a model is likely to require the inclusion of one or more long-term features that can maintain competition between synaptic inputs or outputs of a neuron and retain stability. This work should therefore elucidate the manner in which the emergent synaptic dynamics generated by different profiles of the asymmetric learning window, neural dynamics, spike pairing schemes, additive and multiplicative, pair- and triplet-based rules can contribute to associative Hebbian learning, and thus establish a dialogue between the properties of plasticity at the synaptic level and the systems-level function of a simple, biologically inspired model of declarative memory.

Acknowledgments

This work was supported by funding from the BBSRC and EPSRC (EP/C51632X/1). We also thank members of the CCNR for their support and input.

References

- Abbott, L. F., & Nelson, S. B. (2001). Synaptic plasticity: Taming the beast. *Nature Neuroscience*, 3, 1178–1183.
- Abraham, W. C. (2008). Metaplasticity: Tuning synapses and networks for plasticity. *Nature Reviews Neuroscience*, 9, 387–399.
- Amit, D. J., & Mongillo, G. (2003). Spike-driven synaptic dynamics generating working memory states. *Neural Computation*, 15, 565–596.
- Andersen, P., Morris, R., Amaral, D., Bliss, T., & O'Keefe, J. (2007). Historical perspective: Proposed functions, biological characteristics and neurobiological models of the hippocampus. In P. Andersen, R. Morris, D. Amaral, T. Bliss, & J. O'Keefe (Eds.), *The hippocampus book*. New York: Oxford University Press.
- Barbour, B., Brunel, N., Hakim, V., & Nadal, J.-P. (2007). What can we learn from synaptic weight distributions? *Trends in Neurosciences*, 30(12), 622–629.

- Benuskova, L., & Abraham, W. C. (2007). STDP rule endowed with the BCM sliding threshold accounts for hippocampal heterosynaptic plasticity. *Journal of Computational Neuroscience*, *22*, 129–133.
- Bi, G-Q. (2002). Spatiotemporal specificity of synaptic plasticity: Cellular rules and mechanisms. *Biological Cybernetics*, *87*, 319–332.
- Bi, G-Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, *18*(24), 10464–10472.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, *2*, 32–48.
- Billings, G., & Van Rossum, M. C. W. (2009). Memory retention and spike timing dependent plasticity. *Journal of Neurophysiology*, *101*, 2775–2788.
- Burgess, N. (2007). Computational models of the spatial and mnemonic functions of the hippocampus. In P. Andersen, R. Morris, D. Amaral, T. Bliss, & J O'Keefe (Eds.), *The hippocampus book*. New York: Oxford University Press.
- Burkitt, A. N., Meffin, H., & Grayden, D. B. (2004). Spike timing-dependent plasticity: The relationship to rate-based learning for models with weight dynamics determined by a stable fixed-point. *Neural Computation*, *16*(5), 885–940.
- de Almedia, L., Idiart, M., & Lisman, J. E. (2007). Memory retrieval time and memory capacity of the CA3 network: Role of gamma frequency oscillations. *Learning and Memory*, *14*, 795–806.
- Debanne, D., Gahwiler, B. H., & Thompson, S. M. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *Journal of Physiology*, *507*, 237–247.
- Debanne, D., Gahwiler, B. H., & Thompson, S. M. (1999). Heterogeneity of synaptic plasticity at unitary CA3-CA1 and CA3-CA3 connections in rat hippocampal slice cultures. *Journal of Neuroscience*, *19*(24), 10664–10671.
- Desai, N. S. (2003). Homeostatic plasticity in the CNS: Synaptic and intrinsic forms. *Journal of Physiology*, *97*, 391–402.
- Dudek, S. M., & Bear, M. F. (1992). Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *PNAS*, *89*(10), 4363–4367.
- Frerking, M., Schulte, J., Wiebe, S. P., & Staubli, U. (2005). Spike timing in CA3 pyramidal cells during behaviour: Implications for synaptic transmission. *Journal of Physiology*, *94*(2), 1528–1540.
- Froemke, R. C., & Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, *416*, 433–438.
- Froemke, R. C., Tsay, I. A., Raad, M., Long, J. D., & Dan, Y. (2006). Contribution of individual spikes in burst-induced long-term synaptic modification. *Journal of Neurophysiology*, *95*, 1620–1629.
- Fusi, S., Annunziato, M., Badoni, D., Salamon, A., & Amit, D. J. (2000). Spike-driven synaptic plasticity: Theory, simulation, VLSI implementation. *Neural Computation*, *12*, 2227–2258.
- Grossberg, S. (1987). *The adaptive brain I: Cognition, learning, reinforcement and rhythm*. Amsterdam: North-Holland.

- Gutig, R., Aharonov, R., Rotter, S., & Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *Journal of Neuroscience*, *23*(9), 3697–3714.
- Hebb, D. O. (1949). *Organization of behavior*. New York: Wiley.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *PNAS*, *79*(8), 2554–2558.
- Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, *15*(5), 1063–1070.
- Izhikevich, E. M., & Desai, N. S. (2003). Relating STDP to BCM. *Neural Computation*, *15*, 1511–1523.
- Kampa, B. M., Letzkus, J. J., & Stuart, G. J. (2007). Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in Neurosciences*, *30*(9), 456–463.
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (2001). Intrinsic stabilisation of output rates by spike-based Hebbian learning. *Neural Computation*, *13*, 2709–2741.
- Kobayashi, K., & Poo, M. M. (2004). Spike train timing-dependent associative modification of hippocampal CA3 recurrent synapses by mossy fibers. *Neuron*, *41*, 445–454.
- Lengyel, M., Kwag, J., Paulsen, O., & Dayan, P. (2005). Matching storage and recall: Hippocampal spike timing dependent plasticity and phase response curves. *Nature Neuroscience*, *8*(12), 1677–1683.
- Lomo, T. V., & Bliss, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, *232*(2), 331–356.
- Markram, H., Lubke, J., Frotscher, M., & Sakmann, P. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, *275*, 213–215.
- Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society B*, *262*(841), 23–81.
- Martin, S. J., Grimwood, P. D., & Morris, R. G. M. (2000). Synaptic plasticity and memory: An evaluation of the hypothesis. *Annual Reviews in Neuroscience*, *23*, 649–711.
- Meffin, H., Besson, J., Burkitt, A. N., & Grayden, D. B. (2006). Learning the structure of correlated synaptic subgroups using stable and competitive spike-timing-dependent plasticity. *Physical Review E*, *73*, 041911.
- Miles, R. (1990). Synaptic excitation of inhibitory cells by single CA3 pyramidal cells of the guinea pig in vitro. *Journal of Physiology*, *428*, 61–77.
- Mongillo, G., Curti, E., Romani, S., & Amit, D. J. (2005). Learning in realistic networks of spiking neurons and spike-driven plastic synapses. *European Journal of Neuroscience*, *21*, 3143–3160.
- Morris, R. (2007). Theories of hippocampal function. In P. Andersen, R. Morris, D. Amaral, T. Bliss & J. O'Keefe (Eds.), *The hippocampus book*. New York: Oxford University Press.
- Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics*, *98*, 459–478.
- O'Connor, D. H., Wittenberg, G. M., & Wang, S. S. H. (2005a). Dissection of bidirectional synaptic plasticity into saturable unidirectional processes. *Journal of Neurophysiology*, *94*, 1565–1573.

- O'Connor, D. H., Wittenberg, G. M., & Wang, S. S. H. (2005b). Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *PNAS*, *102*(27), 9679–9684.
- Pfister, J.-P. & Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience*, *26*(38), 9673–9682.
- Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats I: Behavioural correlates and firing repertoires. *Experimental Neurology*, *41*, 461–531.
- Rolls, E. T. (2008). *Memory, attention and decision making*. New York: Oxford University Press.
- Rolls, E. T., & Kesner, R. P. (2006). A computational theory of hippocampal function, and empirical tests of the theory. *Progress in Neurobiology*, *79*, 1–48.
- Samura, T., & Hattori, M. (2005). Hippocampal memory modification induced by pattern completion and spike-timing dependent synaptic plasticity. *International Journal of Neural Systems*, *15*(1–2), 13–22.
- Senn, W. (2002). Beyond spike timing: The role of nonlinear plasticity and unreliable synapses. *Biological Cybernetics*, *87*, 344–355.
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, *3*(9), 919–926.
- Standage, D., Jalil, S., & Trappenberg, T. (2007). Computational consequences of experimentally derived spike-time and weight dependent plasticity rules. *Biological Cybernetics*, *96*(6), 615–623.
- Stellwagen, D., & Malenka, R. C. (2006). Synaptic scaling mediated by glial TNF- α . *Nature*, *440*(20), 1054–1058.
- Toyoizumi, T., Pfister, J.-P., Aihara, K., & Gerstner, W. (2007). Optimality model of unsupervised spike-timing dependent plasticity: Synaptic memory and weight distribution. *Neural Computation*, *19*(3), 639–671.
- Turrigiano, G. G., & Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, *5*, 97–107.
- van Rossum, M. C. W., Bi, G.-Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience*, *20*(23), 8812–8821.
- Wagatsuma, H., & Yamaguchi, Y. (2007). Neural dynamics of the cognitive map in the hippocampus. *Cognitive Neurodynamics*, *1*, 119–141.
- Wang, H. X., Gerkin, R. C., Nauen, D. W., & Bi, G.-Q. (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nature Neuroscience*, *8*, 187–193.
- Wittenberg, G. M. & Wang, S. S-H. (2006). Malleability of spike-timing-dependent plasticity at the CA3-CA1 synapse. *Journal of Neuroscience*, *26*(24), 6610–6617.