

Optimal Spike-Timing-Dependent Plasticity for Precise Action Potential Firing in Supervised Learning

Jean-Pascal Pfister

jean-pascal.pfister@epfl.ch

Taro Toyoizumi

taro@sat.t.u-tokyo.ac.jp

David Barber

david.barber@idiap.ch

Wulfram Gerstner

wulfram.gerstner@epfl.ch

Laboratory of Computational Neuroscience, School of Computer and Communication Sciences and Brain-Mind Institute, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland

In timing-based neural codes, neurons have to emit action potentials at precise moments in time. We use a supervised learning paradigm to derive a synaptic update rule that optimizes by gradient ascent the likelihood of postsynaptic firing at one or several desired firing times. We find that the optimal strategy of up- and downregulating synaptic efficacies depends on the relative timing between presynaptic spike arrival and desired postsynaptic firing. If the presynaptic spike arrives before the desired postsynaptic spike timing, our optimal learning rule predicts that the synapse should become potentiated. The dependence of the potentiation on spike timing directly reflects the time course of an excitatory postsynaptic potential. However, our approach gives no unique reason for synaptic depression under reversed spike timing. In fact, the presence and amplitude of depression of synaptic efficacies for reversed spike timing depend on how constraints are implemented in the optimization problem. Two different constraints, control of postsynaptic rates and control of temporal locality, are studied. The relation of our results to spike-timing-dependent plasticity and reinforcement learning is discussed.

1 Introduction

Experimental evidence suggests that precise timing of spikes is important in several brain systems. In the barn owl auditory system, for example, coincidence-detecting neurons receive volleys of temporally precise spikes from both ears (Carr & Konishi, 1990). In the electrosensory system of mormyrid electric fish, medium ganglion cells receive input at precisely

timed delays after electric pulse emission (Bell, Han, Sugawara, & Grant, 1997). Under the influence of a common oscillatory drive as present in the rat hippocampus or olfactory system, the strength of a constant stimulus is coded in the relative timing of neuronal action potentials (Hopfield, 1995; Brody & Hopfield, 2003; Mehta, Lee, & Wilson, 2002). In humans, precise timing of first spikes in tactile afferents encodes touch signals at the fingertips (Johansson & Birznieks, 2004). Similar codes have also been suggested for rapid visual processing (Thorpe, Delorme, & Van Rullen, 2001), and for the rat's whisker response (Panzeri, Peterson, Schultz, Lebedev, & Diamond, 2001).

The precise timing of neuronal action potentials also plays an important role in spike-timing-dependent plasticity (STDP). If a presynaptic spike arrives at the synapse before the postsynaptic action potential, the synapse is potentiated; if the timing is reversed, the synapse is depressed (Markram, Lübke, Frotscher, & Sakmann, 1997; Zhang, Tao, Holt, Harris, & Poo, 1998; Bi & Poo, 1998, 1999, 2001). This biphasic STDP function is reminiscent of a temporal contrast or temporal derivative filter and suggests that STDP is sensitive to the temporal features of a neural code. Indeed, theoretical studies have shown that given a biphasic STDP function, synaptic plasticity can lead to a stabilization of synaptic weight dynamics (Kempster, Gerstner, & van Hemmen, 1999, 2001; Song, Miller, & Abbott, 2000; van Rossum, Bi, & Turrigiano, 2000; Rubin, Lee, & Sompolinsky, 2001) while the neuron remains sensitive to temporal structure in the input (Gerstner, Kempster, van Hemmen, & Wagner, 1996; Roberts, 1999; Kempster et al., 1999; Kistler & van Hemmen, 2000; Rao & Sejnowski, 2001; Gerstner & Kistler, 2002a).

While the relative firing time of pre- and postsynaptic neurons, and hence temporal aspects of a neural code, play a role in STDP, it is less clear whether STDP is useful to learn a temporal code. In order to elucidate the computational function of STDP, we ask in this letter the following question: What is the ideal form of an STDP function in order to generate action potentials of the postsynaptic neuron with high temporal precision?

This question naturally leads to a supervised learning paradigm: the task to be learned by the neuron is to fire at a predefined desired firing time t^{des} . Supervised paradigms are common in machine learning in the context of classification and prediction problems (Minsky & Papert, 1969; Haykin, 1994; Bishop, 1995), but have more recently also been studied for spiking neurons in feedforward and recurrent networks (Legenstein, Naeger, & Maass, 2005; Rao & Sejnowski, 2001; Barber, 2003; Gerstner, Ritz, & van Hemmen, 1993; Izhikevich, 2003). Compared to unsupervised or reward-based learning paradigms, supervised paradigms on the level of single spikes are obviously less relevant from a biological point, since it is questionable what type of signal could tell the neuron about the "desired" firing time. Nevertheless, we think it is worth addressing the problem of supervised learning—first, as a problem in its own right, and second, as a starting point of spike-based reinforcement learning (Xie & Seung, 2004;

Seung, 2003). Reinforcement learning in a temporal coding paradigm implies that certain sequences of firing times are rewarded, whereas others are not. The “desired firing times” are hence defined indirectly via the presence or absence of a reward signal. The exact relation of our supervised paradigm to reward-based reinforcement learning will be presented in section 4. Section 2 introduces the stochastic neuron model and coding paradigm, which are used to derive the results presented in section 3.

2 Model

2.1 Coding Paradigm. In order to explain our computational paradigm, we focus on the example of temporal coding of human touch stimuli (Johansson & Birznieks, 2004), but the same ideas would apply analogously to the other neuronal systems with temporal codes already mentioned (Carr & Konishi, 1990; Bell et al., 1997; Hopfield, 1995; Brody & Hopfield, 2003; Mehta et al., 2002; Panzeri et al., 2001). For a given touch stimulus, spikes in an ensemble of N tactile afferents occur in a precise temporal order. If the same touch stimulus with identical surface properties and force vector is repeated several times, the relative timing of action potentials is reliably reproduced, whereas the spike timing in the same ensemble of afferents is different for other stimuli (Johansson & Birznieks, 2004). In our model, we assume that all input lines, labeled by the index j with $1 \leq j \leq N$, converge onto one or several postsynaptic neurons. We think of the postsynaptic neuron as a detector for a given spatiotemporal spike pattern in the input. The full spike pattern detection paradigm will be used in section 3.3. As a preparation and first steps toward the full coding paradigm, we also consider the response of a postsynaptic neuron to a single presynaptic spike (section 3.1) or to one given spatiotemporal firing pattern (section 3.2).

2.2 Neuron Model. Let us consider a neuron i that is receiving input from N presynaptic neurons. Let us denote the ensemble of all spikes of neuron j by $x_j = \{t_j^1, \dots, t_j^{N_j}\}$, where t_j^k denotes the time when neuron j fired its k th spike. The spatiotemporal spike pattern of all presynaptic neurons $1 \leq j \leq N$ will be denoted by boldface $\mathbf{x} = \{x_1, \dots, x_N\}$.

A presynaptic spike elicited at time t_j^f evokes an excitatory postsynaptic potential (EPSP) of amplitude w_{ij} and time course $\epsilon(t - t_j^f)$. For simplicity, we approximate the EPSP time course by a double exponential,

$$\epsilon(s) = \epsilon_0 \left[\exp\left(-\frac{s}{\tau_m}\right) - \exp\left(-\frac{s}{\tau_s}\right) \right] \Theta(s), \quad (2.1)$$

with a membrane time constant of $\tau_m = 10$ ms and a synaptic time constant of $\tau_s = 0.7$ ms, which yields an EPSP rise time of 2 ms. Here $\Theta(s)$ denotes the Heaviside step function with $\Theta(s) = 1$ for $s > 0$ and $\Theta(s) = 0$ otherwise.

We set $\epsilon_0 = 1.3$ mV such that a spike at a synapse with $w_{ij} = 1$ evokes an EPSP with amplitude of approximately 1 mV. Since the EPSP amplitude is a measure of the strength of a synapse, we refer to w_{ij} also as the efficacy (or “weight”) of the synapse between neuron j and i .

Let us further suppose that the postsynaptic neuron i receives an additional input $I(t)$ that could arise from either a second group of neurons or from intracellular current injection. We think of the second input as a teaching input that increases the probability that the neuron fires at or close to the desired firing time t^{des} . For simplicity, we model the teaching input as a square current pulse $I(t) = I_0 \Theta(t - t^{\text{des}} + 0.5\Delta T) \Theta(t^{\text{des}} + 0.5\Delta T - t)$ of amplitude I_0 and duration ΔT . The effect of the teaching current on the membrane potential is

$$u_{\text{teach}}(t) = \int_0^\infty k(s)I(t-s)ds \quad (2.2)$$

with $k(s) = k_0 \exp(-s/\tau_m)$, where k_0 is a constant that is inversely proportional to the capacitance of the neuronal membrane.

In the context of the human touch paradigm discussed in section 2.1, the teaching input could represent some preprocessed visual information (“object touched by fingers starts to slip now”), feedback from muscle activity (“strong counterforce applied now”), cross-talk from other detector neurons in the same population (“your colleagues are active now”), or unspecific modulatory input due to arousal or reward (“be aware—something interesting happening now”).

In the context of training of recurrent networks (e.g., Rao & Sejnowski, 2001), the teaching input consists of a short pulse of an amplitude that guarantees action potential firing.

The membrane potential of the postsynaptic neuron i (spike response model; Gerstner & Kistler, 2002b) is influenced by the EPSPs evoked by all afferent spikes of stimulus \mathbf{x} , the “teaching” signal, and the refractory effects generated by spikes t_i^f of the postsynaptic neuron

$$u_i(t|\mathbf{x}, y_t^i) = u_{\text{rest}} + \sum_{j=1}^N w_{ij} \sum_{t_i^f \in x_j} \epsilon(t - t_i^f) + \sum_{t_i^f \in y_i^i} \eta(t - t_i^f) + u_{\text{teach}}(t), \quad (2.3)$$

where $u_{\text{rest}} = -70$ mV is the resting potential, $y_t^i = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ is the set of postsynaptic spikes that occurred before t , and t_i^F always denotes the last postsynaptic spike before t . On the right-hand side of equation 2.3, $\eta(s)$ denotes the spike afterpotential generated by an action potential. We take

$$\eta(s) = \eta_0 \exp\left(-\frac{s}{\tau_m}\right) \Theta(s), \quad (2.4)$$




Figure 1: (A) Escape rate $g(u) = \rho_0 \exp\left(\frac{u-\vartheta}{\Delta u}\right)$. (B) Firing rate of the postsynaptic neuron as a function of the amplitude I_0 of a constant stimulation current (arbitrary units). (C) Interspike interval (ISI) distribution for different input currents.

where $\eta_0 < 0$ is a reset parameter that describes how much the voltage is reset after each spike (for the relation to integrate-and-fire neurons, see Gerstner & Kistler, 2002b). The spikes themselves are not modeled explicitly but reduced to formal firing times. Unless specified otherwise, we take $\eta_0 = -5$ mV.

In a deterministic version of the model, output spikes would be generated whenever the membrane potential u_i reaches a threshold ϑ . In order to account for intrinsic noise and also for a small amount of synaptic noise generated by stochastic spike arrival from additional excitatory and inhibitory presynaptic neurons that are not modeled explicitly, we replace the strict threshold by a stochastic one. More precisely we adopt the following procedure (Gerstner & Kistler, 2002b). Action potentials of the postsynaptic neuron i are generated by a point process with time-dependent stochastic intensity $\rho_i(t) = g(u_i(t))$ that depends nonlinearly on the membrane potential u_i . Since the membrane potential in turn depends on both the input and the firing history of the postsynaptic neuron, we write:

$$\rho_i(t|\mathbf{x}, \mathbf{y}_i^j) = g(u_i(t|\mathbf{x}, \mathbf{y}_i^j)). \quad (2.5)$$

We take an exponential to describe the stochastic escape across threshold: $g(u) = \rho_0 \exp\left(\frac{u-\vartheta}{\Delta u}\right)$ where $\vartheta = -50$ mV is the formal threshold, $\Delta u = 3$ mV is the width of the threshold region and therefore tunes the stochasticity of the neuron, and $\rho_0 = 1/\text{ms}$ is the stochastic intensity at threshold (see Figure 1). Other choices of the escape function g are possible with no qualitative change of the results. For $\Delta u \rightarrow 0$, the model is identical to the deterministic leaky integrate-and-fire model with synaptic current injection (Gerstner & Kistler, 2002b).

We note that the stochastic process, defined in equation 2.5, is similar to but different from a Poisson process since the stochastic intensity depends on the set y_t of the previous spikes of the postsynaptic neuron. Thus, the neuron model has some memory of previous spikes.

2.3 Stochastic Generative Model. The advantage of the probabilistic framework introduced above via the noisy threshold is that it is possible to describe the probability density¹ $P_i(y|\mathbf{x})$ of an entire spike train² $Y(t) = \sum_{t_i^f \in y} \delta(t - t_i^f)$ (see appendix A for details):

$$\begin{aligned} P_i(y|\mathbf{x}) &= \left(\prod_{t_i^f \in y} \rho_i(t_i^f | \mathbf{x}, y_{t_i^f}) \right) \exp \left(- \int_0^T \rho_i(s | \mathbf{x}, y_s) ds \right) \\ &= \exp \left(\int_0^T \log(\rho_i(s | \mathbf{x}, y_s)) Y(s) - \rho_i(s | \mathbf{x}, y_s) ds \right). \end{aligned} \quad (2.6)$$

Thus, we have a generative model that allows us to describe explicitly the likelihood $P_i(y|\mathbf{x})$ of emitting a set of spikes y for a given input \mathbf{x} . Moreover, since the likelihood in equation 2.6 is a smooth function of its parameters, it is straightforward to differentiate it with respect to any variable. Let us differentiate $P_i(y|\mathbf{x})$ with respect to the synaptic efficacy w_{ij} , since this is a quantity that we will use later,

$$\frac{\partial \log P_i(y|\mathbf{x})}{\partial w_{ij}} = \int_0^T \frac{\rho_i'(s | \mathbf{x}, y_s)}{\rho_i(s | \mathbf{x}, y_s)} [Y(s) - \rho_i(s | \mathbf{x}, y_s)] \sum_{t_j^f \in x_j} \epsilon(s - t_j^f) ds, \quad (2.7)$$

where $\rho_i'(s | \mathbf{x}, y_s) = \frac{dg}{du} |_{u=u_i(s|\mathbf{x}, y_s)}$.

In this letter, we propose three different optimal models: A, B, and C (see Table 1). The models differ in the stimulation paradigm and the specific task of the neuron. In section 3, the task and hence the optimality criteria are supposed to be given explicitly. However, the task in model C could also be defined indirectly by the presence or absence of a reward signal, as discussed in section 4.1. The common idea behind all three approaches is the notion of optimal performance. Optimality is defined by an objective function L that is directly related to the likelihood formula of equation 2.6 and that can be maximized by changes of the synaptic weights. Throughout

¹ For simplicity, we denoted the set of postsynaptic spikes from 0 to T by y instead of y_T .

² Capital Y is the spike train generated by the ensemble (lowercase) y .

Table 1: Summary of the Optimality Criterion L for the Unconstrained Scenarios (A_u, B_u, C_u) and the Constrained Scenarios (A_c, B_c, C_c).

Unconstrained Scenarios	Constrained Scenarios
A_u —Postsynaptic spike imposed: $L^{A_u} = \log(\rho(t^{\text{des}}))$	A_c —No activity: $L^{A_c} = L^{A_u} - \int_0^T \rho(t) dt$
B_u —Postsynaptic spike imposed + spontaneous activity: $L^{B_u} = \log(\bar{\rho}(t^{\text{des}}))$	B_c —Stabilized activity: $L^{B_c} = L^{B_u} - \frac{1}{T\sigma^2} \int_0^T (\bar{\rho}(t) - v_0)^2 dt$
C_u —Postsynaptic spike patterns imposed: $L^{C_u} = \log\left(\prod_i P_i(y^i x^i) \prod_{k \neq i} P_i(0 x^k)^{\frac{1}{M-1}}\right)$	C_c —Temporal locality constraint: $L^{C_c} = L^{C_u}, P_{\Delta\Delta'} = a \delta_{\Delta\Delta'} (\Delta - \bar{T}_0)^2$

Notes: The constraint for scenario C is not included in the likelihood function L^{C_c} itself, but rather in the deconvolution with a matrix P that penalizes quadratically the terms that are nonlocal in time. See appendix C for more details.

the article, this optimization is done by a standard technique of gradient ascent,

$$\Delta w_{ij} = \alpha \frac{\partial L}{\partial w_{ij}}, \quad (2.8)$$

with a learning rate α . Since the three models correspond to three different tasks, they have a slightly different objective function. Therefore, gradient ascent yields slightly different strategies for synaptic update. In the following, we start with the simplest model with the aim of illustrating the basic principles that generalize to the more complex models.

3 Results

In this section, we present synaptic updates rules derived by optimizing the likelihood of postsynaptic spike firing at some desired firing time t^{des} . The essence of the argument is introduced in a particularly simple scenario, where the neuron is stimulated by one presynaptic spike and the neuron is inactive except at the desired firing time t^{des} . This is the raw scenario that is developed in several directions.

First, we may ask how the postsynaptic spike at the desired time t^{des} is generated. The spike could simply be given by a supervisor. As always in maximum likelihood approaches, we then optimize the likelihood that this spike could have been generated by the neuron model (i.e., the generative model) given the known input. Or the spike could have been generated by a strong current pulse of short duration applied by the supervisor (teaching input). In this case, the a priori likelihood that the generative model fires at or close to the desired firing time is much higher. The two conceptual paradigms give slightly different results, as discussed in scenario A.

Second, we may, in addition to the spike at the desired time t^{des} , allow for other postsynaptic spikes generated spontaneously. The consequences of spontaneous activity for the STDP function are discussed in scenario B. Third, instead of imposing a single postsynaptic spike at a desired firing time t^{des} , we can think of a temporal coding scheme where the postsynaptic neuron responds to one (out of M) presynaptic spike pattern with a desired output spike train containing several spikes while staying inactive for the other $M - 1$ presynaptic spike patterns. This corresponds to a pattern classification task, which is the topic of scenario C.

Moreover, optimization can be performed in an unconstrained fashion or under some constraint. As we will see in this section, the specific form of the constraint influences the results on STDP, in particular, the strength of synaptic depression for post-before-pre timing. To emphasize this aspect, we discuss two constraints. The first constraint is motivated by the observation that neurons have a preferred working point defined by a typical mean firing rate that is stabilized by homeostatic synaptic processes (Turrigiano & Nelson, 2004). Penalizing deviations from a target firing rate is the constraint that we will use in scenario B. For a very low target firing rate, the constraint reduces to the condition of “no activity,” which is the constraint implemented in scenario A.

The second type of constraint is motivated by the notion of STDP itself: changes of synaptic plasticity should depend on the relative timing of pre- and postsynaptic spike firing and not on other factors. If STDP is to be implemented by some physical or chemical mechanisms with finite time constants, we must require the STDP function to be local in time, that is, the amplitude of the STDP function approaches zero for large time differences. This is the temporal locality constraint used in scenario C. While the unconstrained optimization problems are labeled with the subscript u (A_u, B_u, C_u), the constrained problems are marked by the subscript c (A_c, B_c, C_c) (see Table 1).

3.1 Scenario A: One Postsynaptic Spike Imposed. Let us start with a particularly simple model, which consists of one presynaptic neuron and one postsynaptic neuron (see Figure 2A). Let us suppose that the task of the postsynaptic neuron i is to fire a single spike at time t^{des} in response to the input, which consists of a single presynaptic spike at time t^{pre} , that is, the input is $x = \{t^{\text{pre}}\}$ and the desired output of the postsynaptic neuron is $y = \{t^{\text{des}}\}$. Since there is only a single pre- and a single postsynaptic neuron involved, we drop in this section the indices j and i of the two neurons.

3.1.1 Unconstrained Scenario A_u : One Spike at t^{des} . In this section, we assume that the postsynaptic neuron has not been active in the recent past, that is, refractory effects are negligible. In this case, we have $\rho(t|x, y_t) = \rho(t|x)$ because of the absence of previous spikes. Moreover, since there is




Figure 2: (A) Scenario A: a single presynaptic neuron connected to a postsynaptic neuron with a synapse of weight w . (B) Optimal weight change given by equation 3.2 for scenario A_u . This weight change is exactly the mirror image of an EPSP.

only a single presynaptic spike (i.e., $\mathbf{x} = \{t^{\text{pre}}\}$), we write $\rho(t|t^{\text{pre}})$ instead of $\rho(t|\mathbf{x})$.

Since the task of the postsynaptic neuron is to fire at time t^{des} , we can define the optimality criterion L^{A_u} as the log likelihood of the firing intensity at time t^{des} ,

$$L^{A_u} = \log \left(\rho(t^{\text{des}}|t^{\text{pre}}) \right). \quad (3.1)$$

The gradient ascent on this function leads to the following STDP function,

$$\Delta w^{A_u} = \alpha \frac{\partial L^{A_u}}{\partial w} = \alpha \frac{\rho'(t^{\text{des}}|t^{\text{pre}})}{\rho(t^{\text{des}}|t^{\text{pre}})} \epsilon (t^{\text{des}} - t^{\text{pre}}), \quad (3.2)$$

where $\rho'(t|t^{\text{pre}}) \equiv \frac{d\rho}{du} \Big|_{u=\rho(t|t^{\text{pre}})}$. Since this optimal weight change Δw^{A_u} can be calculated for any presynaptic firing time t^{pre} , we get an STDP function that depends on the time difference $\Delta t = t^{\text{pre}} - t^{\text{des}}$ (see Figure 2B). As we can see directly from equation 3.2, the shape of the potentiation is exactly a mirror image of an EPSP. This result is independent of the specific choice of the function $g(u)$.

The drawback of this simple model becomes apparent if the STDP function given by equation 3.2 is iterated over several repetitions of the experiment. Ideally, it should converge to an optimal solution given by $\Delta w^{A_u} = 0$ in equation 3.2. However, the optimal solution given by $\Delta w^{A_u} = 0$ is problematic: for $\Delta t < 0$, the optimal weight tends toward ∞ , whereas for $\Delta t \geq 0$,

there is no unique optimal weight ($\Delta w^{A_u} = 0, \forall w$). The reason for this problem is that the model describes only potentiation and includes no mechanisms for depression.

3.1.2 Constrained Scenario A_c : No Other Spikes Than at t^{des} . In order to get some insight into where the depression could come from, let us consider a small modification of the previous model. In addition to the fact that the neuron has to fire at time t^{des} , let us suppose that it should not fire anywhere else. This condition can be implemented by an application of equation 2.6 to the case of a single input spike $x = \{t^{\text{pre}}\}$ and a single output spike $y = \{t^{\text{des}}\}$. In terms of notation, we set $P(y|x) = P(t^{\text{des}}|t^{\text{pre}})$ and similarly $\rho(s|x, y) = \rho(s|t^{\text{pre}}, t^{\text{des}})$ and use equation 2.6 to find

$$P(t^{\text{des}}|t^{\text{pre}}) = \rho(t^{\text{des}}|t^{\text{pre}}) \exp \left[- \int_0^T \rho(s|t^{\text{pre}}, t^{\text{des}}) ds \right]. \quad (3.3)$$

Note that for $s \leq t^{\text{des}}$, the firing intensity does not depend on t^{des} ; hence, $\rho(s|t^{\text{pre}}, t^{\text{des}}) = \rho(s|t^{\text{pre}})$ for $s \leq t^{\text{des}}$. We define the objective function L^{A_c} as the log likelihood of generating a single output spike at time t^{des} , given a single input spike at t^{pre} . Hence, with equation 3.3,

$$\begin{aligned} L^{A_c} &= \log(P(t^{\text{des}}|t^{\text{pre}})) \\ &= \log(\rho(t^{\text{des}}|t^{\text{pre}})) - \int_0^T \rho(s|t^{\text{pre}}, t^{\text{des}}) ds, \end{aligned} \quad (3.4)$$

and the gradient ascent $\Delta w^{A_c} = \alpha \partial L^{A_c} / \partial w$ rule yields

$$\Delta w^{A_c} = \alpha \frac{\rho'(t^{\text{des}}|t^{\text{pre}})}{\rho(t^{\text{des}}|t^{\text{pre}})} \epsilon(t^{\text{des}} - t^{\text{pre}}) - \alpha \int_0^T \rho'(s|t^{\text{pre}}, t^{\text{des}}) \epsilon(s - t^{\text{pre}}) ds. \quad (3.5)$$

Since we have a single postsynaptic spike at t^{des} , equation 3.5 can directly be plotted as a STDP function. In Figure 3 we distinguish two different cases. In Figure 3A we optimize the likelihood L^{A_c} in the absence of any teaching input. To understand this scenario, we may imagine that a postsynaptic spike has occurred spontaneously at the desired firing time t^{des} . Applying the appropriate weight update calculated from equation 3.5 will make such a timing more likely the next time the presynaptic stimulus is repeated. The reset amplitude η_0 has only a small influence.

In Figure 3B, we consider a case where firing of the postsynaptic spike at the appropriate time was made highly likely by a teaching input of duration $\Delta T = 1$ ms centered around the desired firing t^{des} . The form of the STDP function depends on the amount η_0 of the reset. If there is no reset $\eta_0 = 0$, the

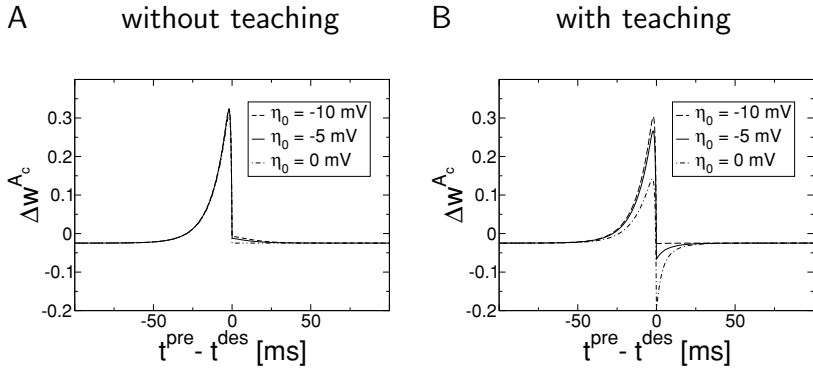


Figure 3: Optimal weight adaptation for scenario A_c given by equation 2.7 in the absence of a teaching signal (A) and in the presence of a teaching signal (B). The weight change in the post-before-pre is governed by the spike afterpotential $u_{AP}(t) = \eta(t) + u_{\text{teach}}(t)$. The duration of the teaching input is $\Delta T = 1$ ms. The amplitude of the current I_0 is chosen so that $\max_t u_{\text{teach}}(t) = 5$ mV. u_{rest} is chosen such that the spontaneous firing rate $g(u_{\text{rest}})$ matches the desired firing rate $1/T$: $u_{\text{rest}} = \Delta u \log(\frac{1}{T\rho_0}) + \theta \simeq -60$ mV. The weight strength is $w = 1$.

STDP function shows strong synaptic depression of synapses that become active after the postsynaptic spike. This is due to the fact that the teaching input causes an increase of the membrane potential that decays back to rest with the membrane time constant τ_m . Hence, the window of synaptic depression is also exponential with the same time constant. Qualitatively the same is true if we include a weak reset. The form of the depression window remains the same, but its amplitude is reduced. The inverse of the effect occurs only for strong reset to or below resting potential. A weak reset is standard in applications of integrate-and-fire models to in vivo data and is one of the possibilities for explaining the high coefficient of variation of neuronal spike trains in vivo (Bugmann, Christodoulou, & Taylor, 1997; Troyer & Miller, 1997).

A further property of the STDP functions in Figure 3 is a negative offset for $|t^{\text{pre}} - t^{\text{des}}| \rightarrow \infty$. The amplitude of the offset can be calculated for $w \simeq 0$ and $\Delta t > 0$, that is, $\Delta w_0 \simeq -\rho'(u_{\text{rest}}) \int_0^\infty \epsilon(s) ds$. This offset is due to the fact that we do not want spikes at other times than t^{des} . As a result, the optimal weight w^* (the solution of $\Delta w^{A_c} = 0$) should be as negative as possible ($w^* \rightarrow -\infty$ or $w^* \rightarrow w^{\text{min}}$ in the presence of a lower bound) for $\Delta t > 0$ or $\Delta t \ll 0$.

3.2 Scenario B: Spontaneous Activity. The constraint in scenario A_c of having strictly no other postsynaptic spikes than the one at time t^{des} may seem artificial. Moreover, it is this constraint that leads to the negative




Figure 4: Scenario B. (A) $N = 200$ presynaptic neurons are firing one after the other at time $t_j = j\delta t$ with $\delta t = 1$ ms. (B) The optimal STDP function of scenario B_u .

offset of the STDP function discussed at the end of the previous paragraph. In order to relax the constraint of no spiking, we allow in scenario B for a reasonable spontaneous activity. As above, we start with an unconstrained scenario B_u before we turn to the constrained scenario B_c .

3.2.1 Unconstrained Scenario B_u : Maximize the Firing Rate at t^{des} . Let us start with the simplest model, which includes spontaneous activity. Scenario B_u is the analog of the model A_u , but with two differences. First, we include spontaneous activity in the model. Since $\rho(t|\mathbf{x}, y_t)$ depends on the spiking history for any given trial, we have to define a quantity that is independent of the specific realizations y of the postsynaptic spike train. Second, instead of considering only one presynaptic neuron, we consider $N = 200$ presynaptic neurons, each emitting a single spike at time $t_j = j\delta t$, where $\delta t = 1$ ms (see Figure 4A). The input pattern will therefore be described by the set of delayed spikes $\mathbf{x} = \{x_j = \{t_j\}, j = 1, \dots, N\}$. As long as we consider only a single spatiotemporal spike pattern in the input, it is always possible to relabel neurons appropriately so that neuron $j + 1$ fires after neuron j .

Let us define the instantaneous firing rate $\bar{\rho}(t)$ that can be calculated by averaging $\rho(t|y_t)$ over all realizations of postsynaptic spike trains:

$$\bar{\rho}(t|\mathbf{x}) = \langle \rho(t|\mathbf{x}, y_t) \rangle_{y_t|\mathbf{x}}. \tag{3.6}$$

Here the notation $\langle \cdot \rangle_{y_t|\mathbf{x}}$ means taking the average over all possible configuration of postsynaptic spikes up to t for a given input \mathbf{x} . In analogy to a Poisson process, a specific spike train with firing times $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ is generated with probability $P(y_t|\mathbf{x})$ given by equation 2.6. Hence, the

average $\langle \cdot \rangle_{y_t|\mathbf{x}}$ of equation 3.6 can be written as follows (see appendix B for numerical evaluation of $\bar{\rho}(t)$):

$$\bar{\rho}(t|\mathbf{x}) = \sum_{F=0}^{\infty} \frac{1}{F!} \int_0^t \dots \int_0^t \rho(t|\mathbf{x}, y_t) P(y_t|\mathbf{x}) dt_1^F, \dots, dt_F^1. \quad (3.7)$$

Analogous to model A_u , we can define the quality criterion as the log likelihood L^{B_u} of firing at the desired time t^{des} :

$$L^{B_u} = \log(\bar{\rho}(t^{\text{des}}|\mathbf{x})). \quad (3.8)$$

Thus, the optimal weight adaptation of synapse j is given by

$$\Delta w_j^{B_u} = \alpha \frac{\partial \bar{\rho}(t^{\text{des}}|\mathbf{x}) / \partial w_j}{\bar{\rho}(t^{\text{des}}|\mathbf{x})}, \quad (3.9)$$

where $\frac{\partial \bar{\rho}(t|\mathbf{x})}{\partial w_j}$ is given by

$$\frac{\partial \bar{\rho}(t|\mathbf{x})}{\partial w_j} = \bar{\rho}'(t|\mathbf{x}) \epsilon(t - t_j) + \left\langle \rho(t|\mathbf{x}, y_t) \frac{\partial}{\partial w_j} \log P(y_t|\mathbf{x}) \right\rangle_{y_t|\mathbf{x}}, \quad (3.10)$$

$\frac{\partial}{\partial w_j} \log P(y_t|\mathbf{x})$ is given by equation 2.7 and $\bar{\rho}'(t|\mathbf{x}) = \langle \frac{d\bar{\rho}}{du} |_{u=\bar{\rho}(t|\mathbf{x}, y_t)} \rangle_{y_t|\mathbf{x}}$.

Figure 4B shows that for our standard set of parameters, the differences to scenario A_u are negligible.

Figure 5A depicts the STDP function for various values of the parameter Δu at a higher postsynaptic firing rate. We can see a small undershoot in the pre-before-post region. The presence of this small undershoot can be understood as follows: enhancing a synapse of a presynaptic neuron that fires too early would induce a postsynaptic spike that arrives before the desired firing time and because of refractoriness would therefore prevent the generation of a spike at the desired time. The depth of this undershoot decreases with the stochasticity of the neuron and increases with the amplitude of the refractory period (if $\eta_0 = 0$, there is no undershoot). In fact, correlations between pre- and postsynaptic firing reflect the shape of an EPSP in the high-noise regime, whereas they show a trough for low noise (Poliakov, Powers, & Binder, 1997; Gerstner, 2001). Our theory shows that the pre-before-post region of the optimal plasticity function is a mirror image of these correlations.

3.2.2 Constrained Scenario B_c : Firing Rate Close to v_0 . In analogy to model A_c we introduce a constraint. Instead of imposing strictly no spikes at times $t \neq t^{\text{des}}$, we can relax the condition and minimize deviations of the




Figure 5: (A) The optimal STDP functions of scenario B_u for different levels of stochasticity described by the parameter Δu . The standard value ($\Delta u = 3$ mV) is given by the solid line; decreased noise ($\Delta u = 1$ mV and $\Delta u = 0.5$ mV) is indicated by dot-dashed and dashed lines, respectively. In the low-noise regime, enhancing a synapse that fires slightly too early can prevent the firing at the desired firing time t^{des} due to refractoriness. To increase the firing rate at t^{des} , it is advantageous to decrease the firing probability time before t^{des} . Methods: For each value of Δu , the initial weight w_0 is set such that the spontaneous firing rate is $\bar{\rho} = 30$ Hz. In all three cases, Δw has been multiplied by Δu in order to normalize the amplitude of the STDP function. Reset: $\eta_0 = -5$ mV. (B) Scenario B_c . Optimal STDP function for scenario B_c given by equation 3.13 for a teaching signal of duration $\Delta T = 1$ ms. The maximal increase of the membrane potential after 1 ms of stimulation with the teaching input is $\max_t u_{\text{teach}}(t) = 5$ mV. Synaptic efficacies w_{ij} are initialized such that $u_0 = -60$ mV, which gives a spontaneous rate of $\bar{\rho} = v_0 = 5$ Hz. Standard noise level: $\Delta u = 3$ mV.

instantaneous firing rate $\bar{\rho}(t|\mathbf{x}, t^{\text{des}})$ from a reference firing rate v_0 . This can be done by introducing into equation 3.8 a penalty term P_B given by

$$P_B = \exp\left(-\frac{1}{T} \int_0^T \frac{(\bar{\rho}(t|\mathbf{x}, t^{\text{des}}) - v_0)^2}{2\sigma^2} dt\right). \quad (3.11)$$

For small σ , deviations from the reference rate yield a large penalty. For $\sigma \rightarrow \infty$, the penalty term has no influence. The optimality criterion is a combination of a high firing rate $\bar{\rho}$ at the desired time under the constraint of small deviations from the reference rate v_0 . If we impose the penalty as a multiplicative factor and take as before the logarithm, we get

$$L^{B_c} = \log\left(\bar{\rho}(t^{\text{des}}|\mathbf{x})P_B\right). \quad (3.12)$$




Figure 6: Scenario C. N presynaptic neurons are fully connected to M postsynaptic neurons. Each postsynaptic neuron is trained to respond to a specific input pattern and not respond to $M - 1$ other patterns as described by the objective function of equation 3.14.

Hence the optimal weight adaptation is given by

$$\Delta w_j^{\text{Bc}} = \alpha \frac{\partial \bar{\rho}(t^{\text{des}}|\mathbf{x}) / \partial w_j}{\bar{\rho}(t^{\text{des}}|\mathbf{x})} - \frac{\alpha}{T\sigma^2} \int_0^T (\bar{\rho}(t|\mathbf{x}, t^{\text{des}}) - v_0) \frac{\partial}{\partial w_j} \bar{\rho}(t|\mathbf{x}, t^{\text{des}}) dt. \quad (3.13)$$

Since in scenario B, each presynaptic neuron j fires exactly once at time $t_j = j\delta t$ and the postsynaptic neuron is trained to fire at time t^{des} , we can interpret the weight adaptation Δw_j^{Bc} of equation 3.13 as an STDP function Δw_j^{Bc} that depends on the time difference $\Delta t = t^{\text{pre}} - t^{\text{des}}$. Figure 5 shows this STDP function for different values of the free parameter σ of equation 3.11. The higher the standard deviation σ , the less effective is the penalty term. In the limit of $\sigma \rightarrow \infty$, the penalty term can be ignored, and the situation is identical to that of scenario B_u.

3.3 Scenario C: Pattern Detection

3.3.1 Unconstrained Scenario C_u: Spike Pattern Imposed. This last scenario is a generalization of scenario A_c. Instead of restricting the study to a single pre- and postsynaptic neuron, we consider N presynaptic neurons and M postsynaptic neurons (see Figure 6). The idea is to construct M independent detector neurons. Each detector neuron $i = 1, \dots, M$, should respond best to a specific prototype stimulus, say \mathbf{x}^i , by producing a desired spike train y^i , but should not respond to other stimuli, $y^i = 0, \forall \mathbf{x}^k, k \neq i$ (see Figure 7). The aim is to find a set of synaptic weights that maximizes the probability

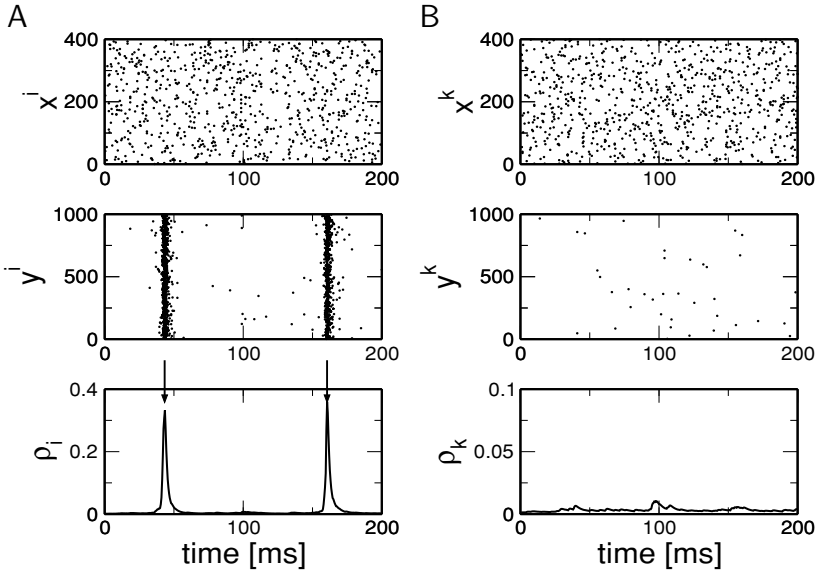


Figure 7: Pattern detection after learning. (Top) The left raster plot represents the input pattern the i th neuron has to be sensitive to. Each line corresponds to one of the $N = 400$ presynaptic neurons. Each dot represents an action potential. The right figure represents one of the patterns the i th neuron should not respond to. (Middle) The left raster plot corresponds to 1000 repetitions of the output of neuron i when the corresponding pattern \mathbf{x}^i is presented. The right plot is the response of neuron i to one of the patterns it should not respond to. (Bottom) The left graph represents the probability density of firing when pattern \mathbf{x}^i is presented. This plot can be seen as the PSTH of the middle graph. Arrows indicate the supervised timing neuron i learned. The right graph describes the probability density of firing when pattern \mathbf{x}^k is presented. Note the different scales of vertical axis.

that neuron i produces y^i when \mathbf{x}^i is presented and produces no output when $\mathbf{x}^k, k \neq i$ is presented. Let the likelihood function L^{C_u} be

$$L^{C_u} = \log \left(\prod_{i=1}^M P_i(y^i | \mathbf{x}^i) \prod_{k=1, k \neq i}^M P_i(0 | \mathbf{x}^k)^{\frac{1}{M-1}} \right) \tag{3.14}$$

where $P_i(y^i | \mathbf{x}^i)$ (see equation 2.6) is the probability that neuron i produces the spike train y^i when the stimulus \mathbf{x}^i is presented. The parameter γ

characterizes the relative importance of the patterns that should not be learned compared to those that should be learned. We get

$$L^{C_u} = \sum_{i=1}^M \log(P_i(y^i | \mathbf{x}^i)) + \gamma \langle \log(P_i(0 | \mathbf{x}^k)) \rangle_{\mathbf{x}^k \neq \mathbf{x}^i}, \quad (3.15)$$

where the notation $\langle \cdot \rangle_{\mathbf{x}^k \neq \mathbf{x}^i} \equiv \frac{1}{M-1} \sum_{k \neq i}^M$ means taking the average over all patterns other than \mathbf{x}^i . The optimal weight adaptation yields

$$\Delta w_{ij}^C = \alpha \frac{\partial}{\partial w_{ij}} \log(P_i(y^i | \mathbf{x}^i)) + \alpha \gamma \left\langle \frac{\partial}{\partial w_{ij}} \log(P_i(0 | \mathbf{x}^k)) \right\rangle_{\mathbf{x}^k \neq \mathbf{x}^i}. \quad (3.16)$$

The learning rule of equation 3.16 gives the optimal weight change for each synapse and can be evaluated after presentation of all pre- and postsynaptic spike patterns; it is a “batch” update rule. Since each pre- and postsynaptic neuron emits many spikes in the interval $[0, T]$, we cannot directly interpret the result of equation 3.16 as a function of the time difference $\Delta t = t^{\text{pre}} - t^{\text{des}}$ as we did in scenario A or B.

Ideally, we would like to write the total weight change of the optimal rule given by equation 3.16 as a sum of contributions

$$\Delta w_{ij}^C = \sum_{\substack{t^{\text{pre}} \in \mathbf{x}_j^i \\ t^{\text{des}} \in \mathbf{y}^i}} \Delta W^{C_u}(t^{\text{pre}} - t^{\text{des}}), \quad (3.17)$$

where $\Delta W^{C_u}(t^{\text{pre}} - t^{\text{des}})$ is an STDP function and the summation runs over all pairs of pre- and postsynaptic spikes. The number of pairs of pre- and postsynaptic spikes with a given time shift is given by the correlation function, which is best defined in discrete time. We assume time steps of duration $\delta t = 0.5$ ms. Since the correlation will depend on the presynaptic neuron j and the postsynaptic neuron i under consideration, we introduce a new index, $k = N(i - 1) + j$. We define the correlation in discrete time by its matrix elements $C_{k\Delta}$ that describe the correlation between the presynaptic spike train $X_j^i(t)$ and the postsynaptic spike train $Y^i(t - T_0 + \Delta\delta t)$. For example, $C_{3\Delta} = 7$ implies that seven spike pairs of presynaptic neuron $j = 3$ with postsynaptic neuron $i = 1$ have a relative time shift of $T_0 - \Delta\delta t$. With this definition, we can rewrite equation 3.17 in vector notation (see section C.1 for more details) as

$$\Delta \mathbf{w}^C \stackrel{!}{=} \mathbf{C} \Delta \mathbf{W}^{C_u}, \quad (3.18)$$




Figure 8: (A) Optimal weight change for scenario C_u . In this case, no locality constraint is imposed, and the result is similar to the STDP function of scenario A_c (with $\eta_0 = 0$ and $u_{\text{teach}}(t) = 0$) represented on Figure 3. (B) Optimal weight change for scenario C_c as a function of the locality constraint characterized by a . The stronger the importance of the locality constraint, the narrower is the spike-spike interaction. For A and B , $M = 20$, $\eta_0 = -5$ mV. The initial weights w_{ij} are chosen so that the spontaneous firing rate matches the imposed firing rate.

where $\Delta \mathbf{w}^C = (\Delta w_{11}^C, \dots, \Delta w_{1N}^C, \Delta w_{21}^C, \dots, \Delta w_{MN}^C)^T$ is the vector containing all the optimal weight change given by equation 3.16 and $\Delta \mathbf{W}^{C_u}$ is the vector containing the discretized STDP function with components $\Delta W_{\Delta}^{C_u} = \Delta W^{C_u}(-T_0 + \Delta \delta t)$ for $1 \leq \Delta \leq 2\tilde{T}_0$ with $\tilde{T}_0 = T_0/\delta$. In particular, the center of the STDP function ($t^{\text{pre}} = t^{\text{des}}$) corresponds to the index $\Delta = \tilde{T}_0$. The symbol $\stackrel{!}{=}$ expresses the fact that we want to find $\Delta \mathbf{W}^{C_u}$ such that $\Delta \mathbf{w}^C$ is as close as possible to $C \Delta \mathbf{W}^{C_u}$. By taking the pseudo-inverse $C^+ = (C^T C)^{-1} C^T$ of C , we can invert equation 3.18 and get

$$\Delta \mathbf{W}^{C_u} = C^+ \Delta \mathbf{w}^C. \tag{3.19}$$

The resulting STDP function is plotted in Figure 8A. As it was the case for the scenario A_u , the STDP function exhibits a negative offset. In addition to the fact the postsynaptic neuron i should not fire at other times than the ones given by y^i , it should also not fire whenever pattern \mathbf{x}^k , $k \neq i$ is presented. The presence of the negative offset is due to those two factors.

3.3.2 Constrained Scenario C_c : Temporal Locality. In the previous paragraph, we obtained a STDP function with a negative offset. This negative offset does not seem realistic because it implies that the STDP function is not localized in time. In order to impose temporal locality (finite memory




Figure 9: (A) Optimal STDP function as a function of the number of input patterns M . ($a = 0.04$, $N = 400$). (B) Optimal weight change as a function of the weight w . If the weights are small (dashed line), potentiation dominates, whereas if they are big (dotted line), depression dominates.

span of the learning rule), we modify equation 3.19 in the following way (see section C.2 for more details):

$$\Delta \mathbf{W}^C = (\mathbf{C}^T \mathbf{C} + \mathbf{P})^{-1} \mathbf{C}^T \Delta \mathbf{w}^C, \quad (3.20)$$

where \mathbf{P} is a diagonal matrix that penalizes nonlocal terms. In this article, we take a quadratic suppression of terms that are nonlocal in time. With respect to a postsynaptic spike at t^{des} , the penalty term is proportional to $(t - t^{\text{des}})^2$. In matrix notation and using our convention that the postsynaptic spike corresponds to $\Delta = \tilde{T}_0$, we have:

$$P_{\Delta \Delta'} = a \delta_{\Delta \Delta'} (\Delta - \tilde{T}_0)^2. \quad (3.21)$$

The resulting STDP functions for different values of a are plotted in Figure 8B. The higher the parameter a , the more nonlocal terms are penalized, the narrower is the STDP function.

Figure 9A shows the STDP functions for various number of patterns M . No significant change can be observed for different numbers of input patterns M . This is due to the appropriately chosen normalization factor $1/(M - 1)$ in the exponent of equation 3.14.

The target spike trains y^j have a certain number of spikes during the time window T ; they set a target value for the mean rate. Let $\nu^{\text{post}} = \frac{1}{TM} \times \sum_{i=1}^M \int_0^T y^i(t) dt$ be the imposed firing rate. Let w_0 denote the amplitude of the synaptic strength such that the firing rate $\bar{\rho}_{w_0}$ given by those weights is identical to the imposed firing rate: $\bar{\rho}_{w_0} = \nu^{\text{post}}$. If the actual weights are




Figure 10: Correlation plot between the optimal synaptic weight change $\Delta \mathbf{w}^{\text{opt}} = \Delta \mathbf{w}^{\text{Cu}}$ and the reconstructed weight change $\Delta \mathbf{w}^{\text{rec}} = C \Delta \mathbf{W}^{\text{C}_c}$ using the temporal locality constraint. (A) No locality constraint; $a = 0$. Deviations from the diagonal are due to the fact that the optimal weight change given by equation 3.16 cannot be perfectly accounted for the sum of pair effects. The mean deviations are given by equation C.7. (B) A weak locality constraint ($a = 0.04$) almost does not change the quality of the weight change reconstruction. (C) Strong locality constraint ($a = 0.4$). The horizontal lines arise since most synapses are subject to a few strong updates induced by pairs of pre- and postsynaptic spike times with small time shifts.

smaller than w_0 , almost all the weights should increase, whereas if they are bigger than w_0 , depression should dominate (see Figure 9B). Thus, the exact form of the optimal STDP function depends on the initial weight value w_0 . Alternatively, homeostatic process could ensure that the mean weight value is always in the appropriate regime.

In equations 3.17 and 3.18, we imposed that the total weight change should be generated as a sum over pairs of pre- and postsynaptic spikes. This is an assumption that has been made in order to establish a link to standard STDP theory and experiments where spike pairs have been in the center of interest (Gerstner et al., 1996; Kempster et al., 1999; Kistler & van Hemmen, 2000; Markram et al., 1997; Bi & Poo, 1998; Zhang et al., 1998). It is, however, clear by now that the timing of spike pairs is only one of several factors contributing to synaptic plasticity. We therefore asked how much we miss if we attribute the optimal weight changes calculated in equation 3.16 to spike pair effects only. To answer this question, we compared the optimal weight change Δw_{ij}^{C} from equation 3.16 with that derived from the pair-based STDP rule $\Delta w_{ij}^{\text{rec}} = \sum_{t^{\text{pre}} \in x_i^t} \sum_{t^{\text{des}} \in y_j} \Delta W^{\text{C}_c}(t^{\text{pre}} - t^{\text{des}})$ with or without locality constraint, that is, for different values of the locality parameter ($a = 0, 0.04, 0.4$) (see Figure 10). More precisely, we simulate $M = 20$ detector neurons, each having $N = 400$ presynaptic inputs, so each subplot of Figure 10 contains 8000 points. Each point in a graph corresponds to the optimal change of one weight for one detector neuron (x -axis) compared to the weight change of the same weight due to pair-based STDP

(y -axis). We found that in the absence of a locality constraint, the pair-wise contributions are well correlated with the optimal weight changes. With strong locality constraints, the quality of the correlation drops significantly. However, for a weak locality constraint that corresponds to an STDP function with reasonable potentiation and depression regimes, the correlation of the pair-based STDP rule with the optimal update is still good. This suggests that synaptic updates with an STDP function based on pairs of pre- and postsynaptic spikes is close to optimal in the pattern detection paradigm.

4 Discussion

4.1 Supervised versus Unsupervised and Reinforcement Learning.

Our approach is based on the maximization of the probability of firing at desired times t^{des} with or without constraints. From the point of view of machine learning, this is a supervised learning paradigm implemented as a maximum likelihood approach using the spike response model with escape noise as a generative model. Our work can be seen as a continuous-time extension of the maximum likelihood approach proposed in Barber (2003).

The starting point of all supervised paradigms is the comparison of a desired output with the actual output a neuron has, or would have, generated. The difference between the desired and actual output is then used as the driving signal for synaptic updates in typical model approaches (Minsky & Papert, 1969; Haykin, 1994; Bishop, 1995). How does this compare to experimental approaches? Experiments focusing on STDP have been mostly performed *in vitro* (Markram et al., 1997; Magee & Johnston, 1997; Bi & Poo, 1998). Since in typical experimental paradigms, firing of the postsynaptic neuron is enforced by strong pulses of current injection, the neuron is not in a natural unsupervised setting; but the situation is also not fully supervised, since there is never a conflict between the desired and actual output of a neuron. In one of the rare *in vivo* experiments to STDP (Frégnac, Shulz, Thorpe, & Bienenstock, 1988, 1992), the spikes of the postsynaptic neuron are also imposed by current injection. Thus, a classification of STDP experiments in terms of supervised, unsupervised, or reward based is not as clear-cut as it may seem at a first glance.

From the point of view of neuroscience, paradigms of unsupervised or reinforcement are probably much more relevant than the supervised scenario discussed here. However, most of our results from the supervised scenario analyzed in this article can be reinterpreted in the context of reinforcement learning following the approach proposed by Xie and Seung (2004). To illustrate the link between reinforcement learning and supervised learning, we define a global reinforcement signal $R(x, y)$ that depends on the spike timing of the presynaptic neurons x and the postsynaptic neuron y .

The quantity optimized in reinforcement learning is the expected reward $\langle R \rangle_{\mathbf{x},y}$ averaged over all pre- and postsynaptic spike trains:

$$\langle R \rangle_{\mathbf{x},y} = \sum_{\mathbf{x},y} R(\mathbf{x}, y) P(y|\mathbf{x}) P(\mathbf{x}). \quad (4.1)$$

If the goal of learning is to maximize the expected reward, we can define a learning rule that achieves this goal by changing synaptic efficacies in the direction of the gradient of the expected reward $\langle R \rangle_{\mathbf{x},y}$:

$$\langle \Delta w \rangle_{\mathbf{x},y} = \alpha \left\langle R(\mathbf{x}, y) \frac{\partial \log P(y|\mathbf{x})}{\partial w} \right\rangle_{\mathbf{x},y}, \quad (4.2)$$

where α is a learning parameter and $\frac{\partial \log P(y|\mathbf{x})}{\partial w}$ is the quantity we discussed in this article. Thus, the quantities optimized in our supervised paradigm re-appear naturally in a reinforcement learning paradigm.

For an intuitive interpretation of the link between reinforcement learning, and supervised learning, consider a postsynaptic spike that (spontaneously) occurred at time t_0 . If no reward is given, no synaptic change takes place. However, if the postsynaptic spike at t_0 is linked to a rewarding situation, the synapse will try to recreate in the next trial a spike at the same time, that is, t_0 has the role of the desired firing time t^{des} introduced in this article. Thus, the STDP function with respect to a postsynaptic spike at t^{des} derived in this article can be seen as the spike timing dependence that maximizes the expected reward in a spike-based reinforcement learning paradigm.

4.2 Interpretation of STDP Function. Let us now summarize and discuss our results in a broader context. In all three scenarios, we found an STDP function with potentiation for pre-before-post timing. Thus, this result is structurally stable and independent of model details. However, depression for post-before-pre timing does depend on model details.

In scenario A, we saw that the behavior of the post-before-pre region is determined by the spike afterpotential (see Table 2 for a result summary of the three models). In the presence of a teaching input and firing rate constraints, a weak reset of the membrane potential after the spike means that the neuron effectively has a depolarizing spike after potential (DAP). In experiments, DAPs have been observed by Feldman (2000), Markram et al. (1997), and Bi and Poo (1998) for strong presynaptic input. Other studies have shown that the level of depression does not depend on the postsynaptic membrane potential (Sjöström, Turrigiano, & Nelson, 2001). In any case, a weak reset (i.e., to a value below threshold rather than to the resting potential) is consistent with the findings of other researchers that used integrate-and-fire models to account for the high coefficient of

Table 2: Main Results for Each Scenario.

Unconstrained Scenarios	Constrained Scenarios
A _u —pre-before-post: LTP ~ EPSP	A _c —post-before-pre: LTD (or LTP) ~ spike afterpotential
B _u —pre-before-post: LTP/LTD ~ reverse correlation	B _c —post-before-pre: LTD ~ increased firing rate
C _u —pre-before-post: LTP ~ EPSP, LTD ~ background patterns	C _c —post-before-pre: LTD ~ background patterns ~ temporal locality

variation of spike trains in vivo (Bugmann et al., 1997; Troyer & Miller, 1997).

In the presence of spontaneous activity (scenario B), a constraint on the spontaneous firing rate causes the optimal weight change to elicit a depression of presynaptic spikes that arrive immediately after the postsynaptic one. In fact, the reason for the presence of the depression in scenario B_c is directly related to the presence of a DAP caused by the strong teaching stimulus. In both scenarios A and B, depression occurs in order to compensate the increased firing probability due to the DAP.

In scenario C, it has been shown that the best way to adapt the weights (in a task where the postsynaptic neuron has to detect a specific input pattern among others) can be described as an STDP function. This task is similar to the one in Izhikevich (2003) in the sense that a neuron is designed to be sensitive to a specific input pattern, but different since our work does not assume any axonal delays. The depression part in this scenario arises from a locality constraint. We impose that weight changes are explained by a sum of pair-based STDP functions.

There are various ways of defining objective functions, and we have used three different objective functions in this article. The formulation of an objective function gives a mathematical expression of the functional role we assign to a neuron. The functional role depends on the type of coding (temporal coding or rate coding) and hence on the information the postsynaptic neurons will read out. The functional role also depends on the task or context in which a neuron is embedded. It might seem that different tasks and coding schemes could thus give rise to a huge number of objective functions. However, the reinterpretation of our approach in the context of reinforcement learning provides a unifying viewpoint: even if the functional role of some neurons in a specific region of the brain can be different from other neurons of a different region, it is still possible to see the different objective functions as different instantiations of the same underlying concept: the maximization of the reward, where the reward is task specific.

More specifically, all objective functions used in this letter maximized the firing probability at a desired firing time t^{des} , reflecting the fact that

in the framework of timing-based codes, the task of a neuron is to fire at precise moments in time. With a different assumption on the neuron's role on signal processing, different objective functions need to be used. An extreme case is a situation where the neuron's task is to avoid firing at time t^{des} . A good illustration is given by the experiments done in the electrosensory lobe (ELL) of the electric fish (Bell et al., 1997). These cells receive two sets of input: the first one contains the pulses coming from the electric organ, and the second input conveys information about the sensory stimulus. Since a large fraction of the sensory stimulus can be predicted by the information coming from the electric organ, it is computationally interesting to subtract the predictable contribution and focus on only the unpredictable part of the sensory stimulus. In this context, a reasonable task would be to ask the neuron not to fire at time t^{des} where t^{des} is the time where the predictable simulation arrives, and this task could be defined indirectly by an appropriate reward signal. An objective function of this type would, in the end, reverse the sign of the weight change of the causal part (LTD for the pre-before-post region), and this is precisely what is seen experimentally (Bell et al., 1997).

In our framework, the definition of the objective function is closely related to the neuronal coding. In scenario C, we postulate that neurons emit a precise spike train whenever the "correct" input is presented and are silent otherwise. This coding scheme is clearly not the most efficient one. Another possibility is to impose postsynaptic neurons to produce a specific but different spike train for each input pattern, and not only for the "correct" input. Such a modification of the scenario does not dramatically change the results. The only effect is to reduce the amount of depression and increase the amount of potentiation.

4.3 Optimality Approaches versus Mechanistic Models. Theoretical approaches to neurophysiological phenomena in general, and to synaptic plasticity in particular, can be roughly grouped into three categories: biophysical models that aim at explaining the STDP function from principles of ion channel dynamics and intracellular processes (Senn, Tsodyks, & Markram, 2001; Shouval, Bear, & Cooper, 2002; Abarbanel, Huerta, & Rabinovich, 2002; Karmarkar & Buonomano, 2002); mathematical models that start from a given STDP function and analyze computational principles such as intrinsic normalization of summed efficacies or sensitivity to correlations in the input (Kempster et al., 1999; Roberts, 1999; Roberts & Bell, 2000; van Rossum et al., 2000; Kistler & van Hemmen, 2000; Song et al., 2000; Song & Abbott, 2001; Kempster et al., 2001; Gütig, Aharonov, Rotter, & Sompolinsky, 2003); and models that derive "optimal" STDP properties for a given computational task (Chechik, 2003; Dayan & Häusser, 2004; Hopfield & Brody, 2004; Bohte & Mozer, 2005; Bell & Parra, 2005; Toyozumi, Pfister, Aihara, & Gerstner, 2005a, 2005b). Optimizing the likelihood of postsynaptic firing in a predefined interval, as we did in this letter, is only one possibility

among others of introducing concepts of optimality (Barlow, 1961; Atick & Redlich, 1990; Bell & Sejnowski, 1995) into the field of STDP. Chechik (2003) uses concepts from information theory but restricts his study to the classification of stationary patterns. The paradigm considered in Bohte and Mozer (2005) is similar to our scenario B_c , in that they use a fairly strong teaching input to make the postsynaptic neuron fire. Bell and Parra (2005) and Toyozumi et al. (2005a) are also using concepts from information theory, but they are applying them to the pre- and postsynaptic spike trains. The work of Toyozumi et al. (2005a) is a clearcut unsupervised learning paradigm and hence distinct from our approach. Dayan and Häusser (2004) use concepts of optimal filter theory but are not interested in precise firing of the postsynaptic neuron. The work of Hopfield and Brody (2004) is similar to our approach in that it focuses on recognition of temporal input patterns, but we are also interested in triggering postsynaptic firing with precise timing. Hopfield and Brody emphasize the repair of disrupted synapses in a network that has previously acquired its function of temporal pattern detector.

Optimality approaches such as ours will never be able to make strict predictions about the properties of neurons or synapses. Optimality criteria may, however, help to elucidate computational principles and provide insights into potential tasks of electrophysiological phenomena such as STDP.

Appendix A: Probability Density of a Spike Train

The probability density of generating a spike train $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ with the stochastic process defined by equation 2.5 can be expressed as follows,

$$P(y_t) = P(t_i^1, \dots, t_i^F)R(t|y_t), \quad (\text{A.1})$$

where $P(t_i^1, \dots, t_i^F)$ is the probability density of having F spikes at times t_i^1, \dots, t_i^F and $R(t|y_t) = \exp(-\int_{t_i^F}^t \rho(t'|y_t)dt')$ corresponds to the probability of having no spikes from t_i^F to t . Since the joint probability $P(t_i^1, \dots, t_i^F)$ can be expressed as a product of conditional probabilities,

$$P(t_i^1, \dots, t_i^F) = P(t_i^1) \prod_{f=2}^F P(t_i^f | t_i^{f-1}, \dots, t_i^1), \quad (\text{A.2})$$

Equation A.1 becomes

$$P(y_t) = \rho(t_i^1 | y_{t_i^1}) \exp\left(-\int_0^{t_i^1} \rho(t' | y_{t'}) dt'\right)$$

$$\begin{aligned}
& \cdot \left\{ \prod_{f=2}^F \rho(t_i^f | y_{t_i^f}) \exp \left(- \int_{t_i^{f-1}}^{t_i^f} \rho(t' | y_{t'}) dt' \right) \right\} \exp \left(- \int_{t_i^F}^t \rho(t' | y_{t'}) dt' \right) \\
& = \left(\prod_{t_i^f \in y_t} \rho(t_i^f | y_{t_i^f}) \right) \exp \left(- \int_0^t \rho(t' | y_{t'}) dt' \right). \tag{A.3}
\end{aligned}$$

Appendix B: Numerical Evaluation of $\bar{\rho}(t)$

Since it is impossible to numerically evaluate the instantaneous firing rate $\bar{\rho}(t)$ with the analytical expression given by equation 3.6, we have to do it in a different way. In fact, there are two ways to evaluate $\bar{\rho}(t)$. Before going into the details, let us first recall that from the law of large numbers, the instantaneous firing rate is equal to the empirical density of spikes at time t ,

$$\langle \rho(t | y_t) \rangle_{y_t} = \langle Y(t) \rangle_{Y(t)}, \tag{B.1}$$

where $Y(t) = \sum_{t_i^f \in y_t} \delta(t - t_i^f)$ is one realization of the postsynaptic spike train. Thus, the first and simpler method based on the right-hand side of equation B.1 is to build a PSTH by counting spikes in small time bins $[t, t + \delta t]$ over, say, $K = 10,000$ repetitions of an experiment. The second, and more advanced, method consists in evaluating the left-hand side of equation B.1 by Monte Carlo sampling. Instead of averaging over all possible spike trains y_t , we generate $K = 10,000$ spike trains by repetition of the same stimulus. A specific spike train $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ will automatically appear with appropriate probability given by equation 2.6. The Monte Carlo estimation $\tilde{\rho}(t)$ of $\bar{\rho}(t)$ can be written as

$$\tilde{\rho}(t) = \frac{1}{P} \sum_{m=1}^P \rho(t | y_t^m), \tag{B.2}$$

where y_t^m is the m th spike train generated by the stochastic process given by equation 2.5. Since we use the analytical expression of $\rho(t | y_t^m)$, we will call equation B.2 a semianalytical estimation. Let us note that the semianalytical estimation $\tilde{\rho}(t)$ converges more rapidly to the true value $\bar{\rho}(t)$ than the empirical estimation based on the PSTH.

In the limit of a Poisson process, $\eta_0 = 0$, the semianalytical estimation $\tilde{\rho}(t)$ given by equation B.2 is equal to the analytical expression of equation 3.6, since the instantaneous firing rate ρ of a Poisson process is independent of the firing history $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ of the postsynaptic neuron.

Appendix C: Deconvolution

C.1 Deconvolution for Spike Pairs. With a learning rule such as equation 3.16, we know the optimal weight change Δw_{ij} for each synapse, but we still do not know the corresponding STDP function.

Let us first define the correlation function $c_k(\tau)$, $k = N(i - 1) + j$ between the presynaptic spike train $X_j^i(t) = \sum_{t^{\text{pre}} \in X_j^i} \delta(t - t^{\text{pre}})$ and the postsynaptic spike train $Y^i(t) = \sum_{t^{\text{des}} \in Y^i} \delta(t - t^{\text{des}})$,

$$c_k(\tau) = \int_0^T X_j^i(s) Y^i(s + \tau) ds, \quad k = 1, \dots, NM, \quad (\text{C.1})$$

where we allow a range $-T_0 \leq \tau \leq T_0$, with $T_0 \ll T$. Since the sum of the pair-based weight change ΔW should be equal to the total adaptation of weights Δw_k , we can write

$$\int_{-T_0}^{T_0} c_k(s) \Delta W(s) ds \stackrel{!}{=} \Delta w_k \quad k = 1, \dots, NM. \quad (\text{C.2})$$

If we want to express equation C.1 in a matrix form, we need to discretize time in small bins δt and define the matrix element,

$$C_{k\Delta} = \int_{\Delta\delta t - T_0}^{(\Delta+1)\delta t - T_0} c_k(s) ds. \quad (\text{C.3})$$

Now equation C.2 becomes

$$\Delta \mathbf{w} \stackrel{!}{=} \mathbf{C} \Delta \mathbf{W}, \quad (\text{C.4})$$

where $\Delta \mathbf{w} = (\Delta w_{11}, \dots, \Delta w_{1N}, \Delta w_{21}, \dots, \Delta w_{MN})^T$ is the vector containing all the optimal weight change and $\Delta \mathbf{W}$ is the vector containing the discretized STDP function: $\Delta W_\Delta = \Delta W(-T_0 + \Delta\delta t)$, for $\Delta = 1, \dots, 2\tilde{T}_0$ with $\tilde{T}_0 = T_0/\delta t$.

In order to solve the last matrix equation, we have to compute the inverse of the nonsquare $NM \times 2\tilde{T}_0$ matrix \mathbf{C} , which is known as the Moore-Penrose inverse (or the pseudo-inverse),

$$\mathbf{C}^+ = (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T, \quad (\text{C.5})$$

which exists only if $(\mathbf{C}^T \mathbf{C})^{-1}$ exists. In fact, the solution given by

$$\Delta \mathbf{W} = \mathbf{C}^+ \Delta \mathbf{w} \quad (\text{C.6})$$

minimizes the square distance

$$D = \frac{1}{2}(C\Delta\mathbf{W} - \Delta\mathbf{w})^2. \quad (\text{C.7})$$

C.2 Temporal Locality Constraint. If we want to impose a constraint of locality, we can add a term in the minimization process of equation C.7 and define the following,

$$E = D + \frac{1}{2}\Delta\mathbf{W}^T P \Delta\mathbf{W}, \quad (\text{C.8})$$

where P is a diagonal matrix that penalizes nonlocal terms. In this article, we take a quadratic suppression of terms that are nonlocal in time:

$$P_{\Delta\Delta'} = a\delta_{\Delta\Delta'} (\Delta - \tilde{T}_0)^2. \quad (\text{C.9})$$

\tilde{T}_0 corresponds to the index of the vector $\Delta\mathbf{W}$ in equations C.4 and C.8 for which $t^{\text{pre}} - t^{\text{des}} = 0$. Calculating the gradient of E given by equation C.8 with respect to $\Delta\mathbf{W}$ yields

$$\nabla_{\Delta\mathbf{W}} E = C^T(C\Delta\mathbf{W} - \Delta\mathbf{w}) + P\Delta\mathbf{W}. \quad (\text{C.10})$$

By looking at the minimal value of E , that is, $\nabla_{\Delta\mathbf{W}} E = 0$, we have

$$\Delta\mathbf{W} = (C^T C + P)^{-1} C^T \Delta\mathbf{w}. \quad (\text{C.11})$$

By setting $a = 0$, we recover the previous case.

Acknowledgments

This work was supported by the Swiss National Science Foundation (200020-103530/1 and 200020-108093/1). T.T was supported by the Research Fellowships of the Japan Society for the Promotion of Science for Young Scientists and a Grant-in-Aid for JSPS Fellows.

References

- Abarbanel, H., Huerta R., & Rabinovich, M. (2002). Dynamical model of long-term synaptic plasticity. *Proc. Natl. Academy of Sci. USA*, 59, 10137–10143.
- Atick, J., & Redlich, A. (1990). Towards a theory of early visual processing. *Neural Computation*, 4, 559–572.
- Barber, D. (2003). Learning in spiking neural assemblies. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems*, 15 (pp. 149–156). Cambridge, MA: MIT Press.

- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Bell, A., & Sejnowski, T. (1995). An information maximization approach to blind separation and blind deconvolution. *Neural Computation*, *7*, 1129–1159.
- Bell, A. J., & Parra, L. C. (2005). Maximising sensitivity in a spiking network. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems*, *17* (pp. 121–128). Cambridge, MA: MIT Press.
- Bell, C., Han, V., Sugawara, Y., & Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, *387*, 278–281.
- Bi, G., & Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, *18*, 10464–10472.
- Bi, G., & Poo, M. (1999). Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature*, *401*, 792–796.
- Bi, G., & Poo, M. (2001). Synaptic modification of correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci.*, *24*, 139–166.
- Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford: Clarendon Press.
- Bohte, S. M., & Mozer, M. C. (2005). Reducing spike train variability: A computational theory of spike-timing dependent plasticity. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems*, *17* (pp. 201–208). Cambridge, MA: MIT Press.
- Brody, C., & Hopfield, J. (2003). Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron*, *37*, 843–852.
- Bugmann, G., Christodoulou, C., & Taylor, J. G. (1997). Role of temporal integration and fluctuation detection in the highly irregular firing of leaky integrator neuron model with partial reset. *Neural Computation*, *9*, 985–1000.
- Carr, C. E., & Konishi, M. (1990). A circuit for detection of interaural time differences in the brain stem of the barn owl. *J. Neurosci.*, *10*, 3227–3246.
- Chechik, G. (2003). Spike-timing-dependent plasticity and relevant mutual information maximization. *Neural Computation*, *15*, 1481–1510.
- Dayan, P., & Häusser, M. (2004). Plasticity kernels and temporal statistics. In S. Thrun, L. K. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems*, *16*. Cambridge, MA: MIT Press.
- Feldman, D. (2000). Timing-based LTP and LTD and vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*, *27*, 45–56.
- Frégnac, Y., Shulz, D. E., Thorpe, S., & Bienenstock, E. (1988). A cellular analogue of visual cortical plasticity. *Nature*, *333*(6171), 367–370.
- Frégnac, Y., Shulz, D. E., Thorpe, S., & Bienenstock, E. (1992). Cellular analogs of visual cortical epigenesis. I: Plasticity of orientation selectivity. *Journal of Neuroscience*, *12*(4), 1280–1300.
- Gerstner, W. (2001). Coding properties of spiking neurons: Reverse- and cross-correlations. *Neural Networks*, *14*, 599–610.
- Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, *383*, 76–78.
- Gerstner, W., & Kistler, W. K. (2002a). Mathematical formulations of Hebbian learning. *Biological Cybernetics*, *87*, 404–415.

- Gerstner, W., & Kistler, W. K. (2002b). *Spiking neuron models*. Cambridge: Cambridge University Press.
- Gerstner, W., Ritz, R., & van Hemmen, J. L. (1993). Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.*, *69*, 503–515.
- Gütig, R., Aharonov, R., Rotter, S., & Sompolinsky, H. (2003). Learning input correlations through non-linear temporally asymmetric Hebbian plasticity. *J. Neuroscience*, *23*, 3697–3714.
- Haykin, S. (1994). *Neural networks*. Upper Saddle River, NJ: Prentice Hall.
- Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature*, *376*, 33–36.
- Hopfield, J. J., & Brody, C. D. (2004). Learning rules and network repair in spike-timing-based computation networks. *Proc. Natl. Acad. Sci. USA*, *101*, 337–342.
- Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, *14*, 1569–1572.
- Johansson, R., & Birznieks, I. (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience*, *7*, 170–177.
- Karmarkar, U., & Buonomano, D. (2002). A model of spike-timing dependent plasticity: One or two coincidence detectors. *J. Neurophysiology*, *88*, 507–513.
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E*, *59*, 4498–4514.
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Computation*, *13*, 2709–2741.
- Kistler, W. M., & van Hemmen, J. L. (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput.*, *12*, 385–405.
- Legenstein, R., Naeger, C., & Maass, W. (2005). What target functions can be learnt with spike-timing-dependent plasticity? *Neural Computation*, *17*, 2337–2382.
- Magee, J. C., & Johnston, D. (1997). A synaptically controlled associative signal for Hebbian plasticity in hippocampal neurons. *Science*, *275*, 209–213.
- Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSP. *Science*, *275*, 213–215.
- Mehta, M. R., Lee, A. K., & Wilson, M. A. (2002). Role of experience of oscillations in transforming a rate code into a temporal code. *Nature*, *417*, 741–746.
- Minsky, M. L., & Papert, S. A. (1969). *Perceptrons*. Cambridge MA: MIT Press.
- Panzeri, S., Peterson, R., Schultz, S., Lebedev, & Diamond, M. (2001). The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron*, *29*, 769–777.
- Poliakov, A. V., Powers, R. K., & Binder, M. C. (1997). Functional identification of the input-output transforms of motoneurons in the rat and cat. *J. Physiology*, *504*, 401–424.
- Rao, R. P. N., & Sejnowski, T. J. (2001). Spike-timing-dependent Hebbian plasticity as temporal difference learning. *Neural Computation*, *13*, 2221–2237.
- Roberts, P. (1999). Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. *J. Computational Neuroscience*, *7*, 235–246.
- Roberts, P., & Bell, C. (2000). Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. *Computational Neuroscience*, *9*, 67–83.

- Rubin, J., Lee, D. D., & Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Physical Review Letters*, *86*, 364–367.
- Senn, W., Tsodyks, M., & Markram, H. (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation*, *13*, 35–67.
- Seung, S. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron*, *40*, 1063–1073.
- Shouval, H. Z., Bear, M. F., & Cooper, L. N. (2002). A unified model of NMDA receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA*, *99*, 10831–10836.
- Sjöström, P., Turrigiano, G., & Nelson, S. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, *32*, 1149–1164.
- Song, S., & Abbott, L. (2001). Column and map development and cortical re-mapping through spike-timing dependent plasticity. *Neuron*, *32*, 339–350.
- Song, S., Miller, K., & Abbott, L. (2000). Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience*, *3*, 919–926.
- Thorpe, S., Delorme, A., & Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Networks*, *14*, 715–725.
- Toyozumi, T., Pfister, J.-P., Aihara, K., & Gerstner, W. (2005a). Spike-timing dependent plasticity and mutual information maximization for a spiking neuron model. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in neural information processing systems*, *17* (pp. 1409–1416). Cambridge, MA: MIT Press.
- Toyozumi, T., Pfister, J.-P., Aihara, K., & Gerstner, W. (2005b). Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission. *Proc. National Academy Sciences (USA)*, *102*, 5239–5244.
- Troyer, T. W., & Miller, K. (1997). Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation*, *9*, 971–983.
- Turrigiano, G., & Nelson, S. (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. *J. Neuroscience*, *20*, 8812–8821.
- Xie, X., & Seung, S. (2004). Learning in neural networks by reinforcement of irregular spiking. *Phys. Rev. E*, *69*, 041909.
- Zhang, L., Tao, H., Holt, C., Harris, W. A., & Poo, M.-M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, *395*, 37–44.