A computational framework for cortical learning

Originally published in:
Biological Cybernetics 90(6), https://doi.org/10.1007/s00422-004-0487-1

Biol. Cybern. 3, 1–10 (2004)
DOI 10.1007/s00422-004-0487-1

Biological Cybernetics

A computational framework for cortical learning

Roland E. Suri
3641 Midvale Ave. #205, Los Angeles, CA 90034, USA

Received: 4 February 2004 / Accepted: 28 April 2004 / Published online: ■

Abstract. Recent physiological findings have revealed that long-term adaptation of the synaptic strengths between cortical pyramidal neurons depends on the temporal order of presynaptic and postsynaptic spikes, which is called spike-timing-dependent plasticity (STDP) or temporally asymmetric Hebbian (TAH) learning. Here I prove by analytical means that a physiologically plausible variant of STDP adapts synaptic strengths such that the presynaptic spikes predict the postsynaptic spikes with minimal error. This prediction error model of STDP implies a mechanism for cortical memory: cortical tissue learns temporal spike patterns if these spike patterns are repeatedly elicited in a set of pyramidal neurons. The trained network finishes these patterns if their beginnings are presented, thereby recalling the memory. Implementations of the proposed algorithms may be useful for applications in voice recognition and computer vision.

1 Introduction

Animal learning psychologists, engineers, and philosophers have long been speculating that basic aspects of cognition may be explained with the hypothesis that the brain learns and uses internal models (Sutton and Barto (1981); Wolpert et al. (1995); Suri (2002)). A correct internal model emulates the experience of an agent in the real world by providing the sensory experience to the agent that would have resulted from his actions without really executing them. This enables the agent to evaluate the consequences of potential actions and to select the action with the best predicted outcome. Internal models have widespread applications in engineering sciences for the control of physical systems (Ljung and Soderstrom (1983)). Internal models are often represented by a set of linear differential equations. Several methods have been developed that compute the correct model parameter values for a set of linear differential equations by using previous experience, which are usually called system identification methods. One of these methods is correlation analysis, which uses temporal correlations between the inputs and the outputs to compute the values of the model parameters that minimize the difference between the predicted and the actual outputs.

Recent theoretical and experimental studies have demonstrated long-term potentiation (LTP) of synaptic strengths between cortical pyramidal neurons if the presynaptic spike precedes the postysynaptic spike, and long-term depression (LTD) if the temporal order is reversed, which is called spike-timing-dependent plasticity (STDP) (Gerstner et al. (1996); Markram et al. (1997); Bi and Poo (1998); Debanne et al. (1998); Feldman (2000); Sjostrom et al. (2001); Froemke and Dan (2002)). Several simulation studies have suggested that STDP adapts synaptic strength such that spike sequences can be learned by chains of neurons (Levy (1996); Gerstner and Abbott (1997); August and Levy (1999); Rao and Sejnowski (2001)). As learning of temporal sequences is a salient feature of internal models, these studies raise the question whether STDP serves to learn internal models.

In the current study, I derive a physiologically plausible version of STDP assuming that STDP minimizes the errors between the actual postsynaptic spikes and the postsynaptic spikes predicted by the presynaptic spikes. Such a model neuron, or a network of such model neurons, performs system identification and thus learns an internal model.

2 Derivation of STDP rule

A spiking pyramidal neuron is modeled according to the simple spike response model SRM_0 (Gerstner and van Hemmen (1992, 1994); Gerstner and Kistler (2002)). Neuronal inputs are transmitted by M excitatory synapses of synaptic strengths w_i. Presynaptic spikes at synapse i cause excitatory postsynaptic potentials (EPSPs) that are superimposed to generate the signal $w_i \varepsilon_i(t)$. Starting with the reset potential after the previous spike occurring at time $\hat{t}$, the signals $w_i \varepsilon_i(t)$ of all activated synapses are superimposed until the membrane potential reaches a spike thresh-

Correspondence to: R. E. Suri
(e-mail: suri@salk.edu,
Tel.: +1-310-8158310)

4 2 2 0 0 4 8 7 | B | Despatch: **8/7/2004** Journal : **B C Y B** No. of Pages : **10**
Journal No. Article No. | Author's Disk received ☑ Used ☑ Corrupted ☐ Mismatch ☐ Keyed ☐

old. A nonlinear function f generates a spike by adding a spike-shaped function to the membrane potential, and the spike time $\hat{t}$ is updated to the time of the new spike. After the spike, the nonlinear function f resets the membrane potential to a reset potential, and the integration of upcoming EPSPs starts with this reset value. The membrane potential $y(t)$ can thus be written as (Gerstner and van Hemmen 1992, 1994; Gerstner and Kistler 2002)

$$y(t) = f\left(\hat{t}, \sum_{i=1}^{M} w_i \varepsilon_i(t)\right) \quad . \tag{1}$$

The principal analytical result of the current paper is that system identification can be achieved if the change in synaptic strength Δw_i is computed using the cross correlation between the membrane potential $y(t)$ and the derivative of the normalized and superimposed EPSPs $x_i(t) \equiv \mathrm{d}\varepsilon_i(t)/\mathrm{d}t$ according to

$$\Delta w_i \cong c_i \sum_{t=1}^{N} y(t) x_i(t) \quad , \tag{2}$$

where c_i is a scalar constant that depends on the average firing rates of the presynaptic and the postsynaptic neuron. Equation (2) is derived by using correlation analysis (Appendix) and by assuming that synchronous EPSP arrivals are infrequent.

2.1 Proof of Eq. (2)

To derive (2), the level of the membrane potential just before the arrival of the presynaptic spike is treated as a pseudorandom variable. The spiking neuron is thus treated as a probabilistic unit, although the used neuron model is deterministic. It may be surprising that a deterministic process is treated as a pseudorandom process, but this is actually quite common: many neural networks use pseudorandom numbers for search processes (such as "simulated annealing"), although the pseudorandom numbers are computed by deterministic algorithms. The pseudorandom numbers can be treated as if they were random numbers because the deterministic process that generates the pseudorandom numbers is independent of the processes in the neural network. For the same reason, the level of the membrane potential just before arrival of a spike can be treated as a pseudorandom variable because it is assumed that the arrival times of the presynaptic spikes are not correlated with the values of the postsynaptic membrane potential (the direct influence of the synaptic activation on the membrane potential is not taken into account). To treat the spiking neuron as a probabilistic unit, function f is separated in a term that represents a linear estimate of the membrane potential and in an error term $e\left(\hat{t}, t\right)$ (a correction for nonlinearity is described in the last paragraph in Sect. 4.1)

$$f\left(\hat{t}, \sum_{i=1}^{M} w_i \varepsilon_i(t)\right) = \left[\sum_{i=1}^{M} \frac{w_i}{\tilde{\vartheta}} x_i(t)\right] + e\left(\hat{t}, t\right) \quad . \tag{3}$$

The term in the square brackets serves as an estimate of the membrane potential change that may be caused by an upcoming postsynaptic spike, assuming that the current membrane potential is an unknown pseudorandom variable. The value of the scalar $\tilde{\vartheta}$ is estimated such that the normalization of the synaptic strengths w_i on the right side of (3) corresponds to that on the left side. The error term $e\left(\hat{t}, t\right)$ is not small, as it corresponds to the value of the membrane potential at the beginning of a given time step and represents the unpredicted portion of spikes. It is assumed that the error term $e\left(\hat{t}, t\right)$ is independent of the correct synaptic strengths w_i. In Sect. 4.1, I will show that this approximation leads to small relative errors of the estimated synaptic strengths (1.7%), as the probability that a presynaptic spike will elicit a postsynaptic spike is not exactly proportional to the synaptic strength. I will propose a method that would avoid these inaccuracies. The functions representing presynaptic spikes by $x_i(t)$ are chosen with zero mean and such that they can be computed from the normalized and superimposed EPSPs $\varepsilon_i(t)$. There is considerable freedom in choosing the representation of presynaptic spikes by $x_i(t)$. As explained after (8), I choose to represent spikes by the derivatives of the normalized and superimposed EPSP by setting $x_i(t) \equiv \mathrm{d}\varepsilon_i(t)/\mathrm{d}t$. Thus, each function $x_i(t)$ consists of a train of brief bumps of equal shape and amplitude that indicate the arrival of presynaptic spikes at synapse i.

The relationship between (3) and the description of the neuron according to (1) can also be described in the terms of system identification. The methods for the identification of dynamic systems can be used to mimic the input–output relationship of a nonlinear physical system by treating the nonlinearity as a random disturbance. The model structure, which is the functional form of the model, has to be chosen appropriately for the modeled system. Once the correct model parameters are determined, the model can be used to predict the average output for a given input. The description of the neuron according to (1) corresponds to the nonlinear physical system, and (3) corresponds to its model. The model structure is chosen by defining $x_i(t)$ and by assuming a linear relationship between the synaptic strengths w_i and their influence on the membrane potential $y(t)$. The correct model parameters w_i minimize the error $e\left(\hat{t}, t\right)$ between the neuron's membrane potential $y(t)$ and the predicted membrane potential $\sum_{i=1}^{M} w_i x_i(t)/\tilde{\vartheta}$. Since the error $e\left(\hat{t}, t\right)$ is dominated by the contributions of the postsynaptic spikes, the correct model parameters w_i reflect the probability that a presynaptic spike at synapse i will elicit a postsynaptic spike. In contrast to traditional system identification, the time course of the spike prediction does not need to match the shape of the actual spike since the model only needs to predict whether or not a postsynaptic spike occurs. Therefore it is sufficient to choose $x_i(t)$ such that $x_i(t)$ is correlated with the membrane potential $y(t)$ for nonzero values of the synaptic strength w_i (see also Sect. 2.1.1).

By using the definition of $y(t)$ according to (1), (3) can be written as

$$y(t) = \frac{\mathbf{x}^{\mathbf{T}}(t)\mathbf{w}}{\tilde{\vartheta}} + e(\hat{t}, t) \quad . \tag{4}$$

The vector signal $\mathbf{x^T}(t)$ denotes the transpose of the column vector $\mathbf{x}(t) \equiv (x_1(t), x_2(t), \ldots, x_M(t))$, the synaptic strengths are written as the column vector $\mathbf{w}(t) \equiv (w_1(t), w_2(t), \ldots, w_M(t))$, and $\mathbf{x^T}(t)\mathbf{w}$ denotes the scalar product. To derive the following two equations, I follow the derivation of correlation analysis as described in the appendix. The squared error term in (4) is minimized with respect to the synaptic strengths $\mathbf{w}$ by setting the gradient to zero, leading to

$$\frac{\mathbf{w}}{\tilde{\vartheta}} = \left[\frac{1}{N}\sum_{t=1}^{N}\mathbf{x}(t)\mathbf{x}^T(t)\right]^{-1}\frac{1}{N}\sum_{t=1}^{N}\mathbf{x}(t)y(t) \quad . \tag{5}$$

The term in the square brackets is a matrix whose elements are the correlations between the inputs $x_i(t)$ and $x_k(t)$ on synapses i and k. Only the spike representations that are aligned within a few milliseconds substantially contribute to these correlations, as the width of the EPSP derivatives in $x_i(t)$ is only a few milliseconds. Since the number of EPSPs from different synapses that are *not* aligned to each other is much larger than the number of those that *are* aligned, the cross correlations are much smaller than the autocorrelations

$$\sum_{t=1}^{N} x_i(t)x_k(t) << \sum_{t=1}^{N} x_i(t)x_i(t) \quad . \tag{6}$$

Due to (6), the $M \times M$ matrix in the square brackets in (5) is approximately diagonal (M is the number of synapses). For a synapse with number i, (5) can thus be written as

$$w_i \cong w_i^* \equiv c_i \sum_{t=1}^{N} y(t)x_i(t) \quad , \tag{7}$$

with $c_i = \frac{\tilde{\vartheta}}{\sum_{t=1}^{N} x_i(t)x_i(t)}$. For the shown simulations, c_i is set to the same constant value $c_i = c$ for all synapses, as all presynaptic firing rates are assumed to be equal. If the synapse is retrained during the time interval $[t', N]$ with $1 < t' < N$, the change $\Delta w_i = w_i(N) - w_i(t')$ of the synaptic strength is

$$\Delta w_i \cong c_i \sum_{t=t'}^{N} y(t)x_i(t) \quad , \tag{8}$$

which is equivalent to (2, 3). Equation (7) is the explicit solution for the synaptic strengths w_i that minimize the least square prediction error. Therefore, (8) guarantees that the synaptic strengths w_i minimize the prediction error if the approximations of (3, 6) are accurate.

2.1.1 Choice of the spike representations. Since it is assumed that all information provided by a spike is given by the time of its occurrence and not by its shape, there is considerable freedom in choosing functions that represent spikes. This leads to a wide range of potential functions for the Hebbian learning window. Let us assume that the shapes of the spike representations in the input and the output signals are modified by linear or nonlinear operators U and O with $\tilde{x}_i(t) = U(x_i(t))$ and $\tilde{y}(t) = U(y(t))$ such that (6) is satisfied for $\tilde{x}_i(t) = U(x_i(t))$. According to (7), it is sufficient to require that the correlation of $\tilde{x}_i(t)$ with $\tilde{y}(t)$ be proportional to the synaptic strength w_i (or that there be a monotonous relationship using the modification explained in Sect. 4.1). Therefore, it seems that a sufficient condition for the operators U and O is that for synaptic strengths $w_i \neq 0$ the correlation between the transformed representations is nonzero, i.e., $\sum_{t=1}^{N} O(y(t))U(x_i(t)) = \sum_{t=1}^{N} \tilde{y}(t)\tilde{x}_i(t) \neq 0$. For simplicity, it is assumed that $\tilde{y}(t) = y(t)$.

One may suggest representing the input spikes in $x_i(t)$ by spike-shaped functions. This method is not optimal as only synaptic strengths would get adapted whose presynaptic spikes predict the postsynaptic spike with a temporal precision in the order of the spike duration. The spiking mechanism of the assumed neuron makes it unlikely that such a high temporal precision can be achieved, as an EPSP can elicit a postsynaptic spike during a much longer time period. Therefore, the functions in $x_i(t)$ that represent presynaptic spike arrivals should be broader, such that they are more tolerant to the temporal inaccuracy of the spike prediction. Their values should be large when presynaptic EPSPs are likely to elicit postsynaptic spikes. If the membrane potential remains approximately constant during the EPSP, as in the simulations shown here, the probability that an EPSP will elicit a postsynaptic spike (also called cross correlation) is roughly proportional to the derivative of the EPSP (Gerstner (2001)). Thus, each EPSP contributes to the spike probability at a certain time t with a term roughly proportional to $w_i(\mathrm{d}\varepsilon_i(t)/\mathrm{d}t)$. A postsynaptic spike becomes certain if the amplitude of the EPSP is equal to the difference ϑ between the spike threshold and the reset potential, suggesting that the spike probability is approximately equal to $w_i(\mathrm{d}\varepsilon_i(t)/\mathrm{d}t)/\tilde{\vartheta}$. The parameter $\tilde{\vartheta}$ would be equal to ϑ if the subthreshold membrane potential had been exactly uniformly distributed between the reset potential and the firing threshold. Since this was not the case in our simulations, the parameter $\tilde{\vartheta}$ was estimated in preliminary simulations. (A more accurate fit would use $\tilde{\vartheta}(w_i)$, as described in the last paragraph of Sect. 4.1). For the sum of all incoming EPSPs, the spike probability is thus roughly equal to $\sum_{i=1}^{M} \frac{w_i}{\tilde{\vartheta}} \frac{\mathrm{d}\varepsilon_i(t)}{\mathrm{d}t}$. Since the spike amplitude is normalized to the value of one, this term corresponds to the term in the square brackets in (3). Therefore, $x_i(t)$ is defined as the derivative of the normalized and superimposed EPSPs by setting $x_i(t) \equiv \mathrm{d}\varepsilon_i(t)/\mathrm{d}t$.

3 Comparison with STDP

Typical EPSPs and spike trajectories are simulated to compute the change in the synaptic strength Δw_i. The time course of the EPSP is simulated with $[\exp(-t/50\,\mathrm{ms}) - \exp(-t/2\,\mathrm{ms})]$ to approximate EPSPs shown by Bi and Poo (1998). The EPSP is normalized such that the maximum is 0.1 (Fig. 1a, top), and its derivative $x_i(t)$ is computed (Fig. 1a, middle). To simulate a spike, the membrane potential $y(t)$ is linearly increased to the value of one

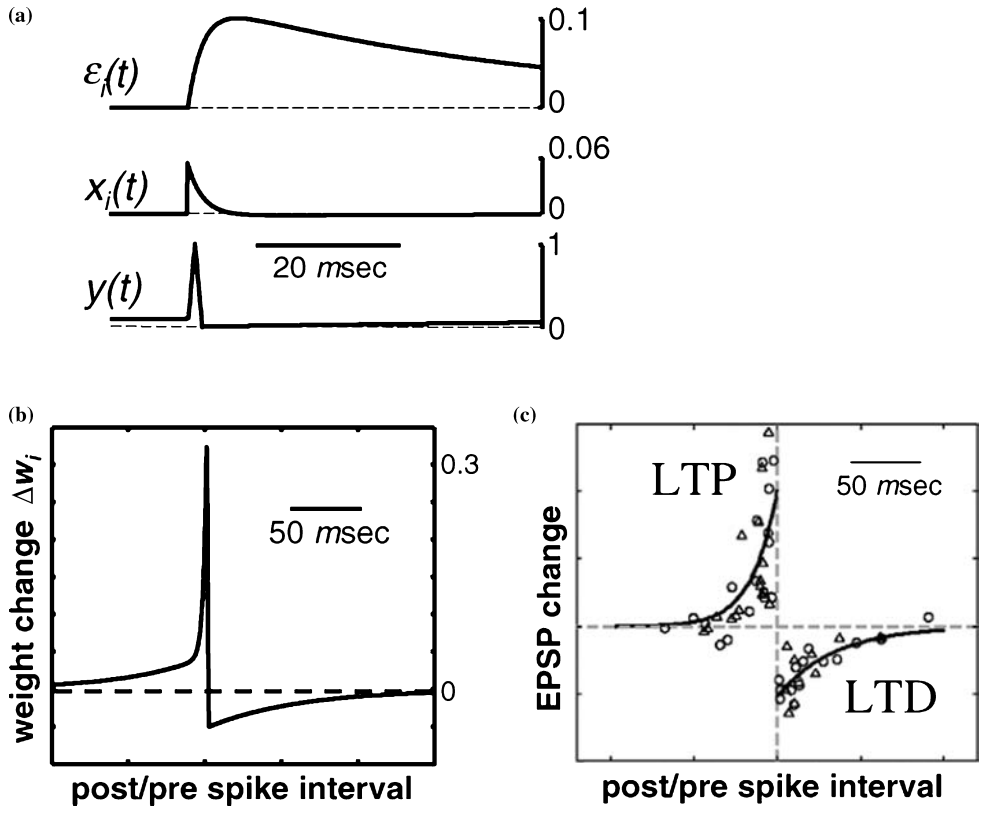

Fig. 1a–c. Comparison between change in synaptic strength proposed by the model and long-term adaptation of synaptic strength by spike-timing-dependent plasticity (STDP). **a** A simulated signal $\varepsilon_i(t)$, representing an EPSP (*top line*), and its derivative $x_i(t)$ (*middle line*) are shown for a presynaptic spike at synapse *i*. *The signal* $y(t)$ represents the postsynaptic membrane potential with a spike (*bottom line*). The EPSP and the spike are shown in the temporal relationship that leads to the maximal weight change. **b** The change in synaptic strengths Δw_i is shown for different intervals between the presynaptic EPSP onset and the postsynaptic spike according to (8). The change Δw_i becomes maximal if the presynaptic spike arrives just in time to cause the postsynaptic spike and negative if the presynaptic spike arrives later. **c** Experimental studies have demonstrated that long-term changes of the synaptic strengths between pyramidal neurons depend on the time difference between the presynaptic EPSP and the postsynaptic spike (figure adapted with permission of Nature from Froemke RC, Dan Y (2002) Nature 416 (6879):433–438; the time scale was inverted; one vertical unit corresponds to a 50% change in the normalized EPSP slope)

during 2 ms and then linearly decreased during 2 ms to a reset value of zero. After this reset, $y(t)$ recovers with a time constant of 60 ms to the baseline value, which is set to 0.1 (Fig. 1a, bottom). This shape approximately reproduces the measured membrane potential if spikes are elicited by brief input currents (see Fig. 5d in Feldman 2000 or Fig. 3a in Markram et al. 1997).

The change in synaptic strength Δw_i as a function of the interval between the presynaptic EPSP onset and the postsynaptic spike is computed according to (8) (Fig. 1b, time steps of 0.1 ms). The change in synaptic strength is strikingly similar to that of the experimentally established adaptation of synaptic strength due to STDP (Fig. 1c). STDP is characterized by LTP if a presynaptic spike precedes a postsynaptic spike and by LTD if the temporal order of the spikes is reversed (Markram et al. 1997; Bi and Poo 1998; Debanne et al. 1998; Feldman 2000; Froemke and Dan 2002). Strong LTP occurs if the presynaptic spike occurs just in time to elicit the postsynaptic spike (Bi and Poo 1998). Note that according to (8) the synaptic strength can become negative. Since pyramidal neurons are excitatory, connections between pyramidal neurons are not able to represent associations that would require such negative synaptic strengths, such as a presynaptic spike following a postsynaptic spike. This suggests that such negative cross correlations are likely to be mediated by inhibitory neurons. The presence of inhibitory neurons may reduce the chances that a correct representation of the long-term experience of a pyramidal neuron would require negative synaptic strengths.

4 Simulations

Equation (8) can be used to identify the parameters of an internal model. Let us assume that only the input signals $\varepsilon_i(t)$ (or their EPSP onsets) and the resulting output signal $y(t)$ of a physical process are known, and we would like to use the model neuron according to (1) to emulate this physical process. Equation (8) can be used to determine the parameters w_i from the observed input signals $\varepsilon_i(t)$ and the observed output signal $y(t)$ of the physical process. For the input signals $\varepsilon_i(t)$ and the desired output signal $y(t)$, (8) computes the values of the synaptic strengths w_i such that the input signals $\varepsilon_i(t)$ elicit the desired output signal $y(t)$.

4.1 Accuracy of analytical result (offline learning)

The accuracy of (8) was investigated in computer simulations of one SRM_0 neuron with 500 synapses [(1)]. The presynaptic inputs were Poisson distributed with average firing rates of 10 Hz. EPSPs were simulated as shown in Fig. 1a. Time courses of typical signals are shown in Fig. 2a. The values of the synaptic strengths w_i were uniformly distributed between 0 and 0.07. When the membrane potential reached a spike threshold of 0.1 at time t, it was set to the value of one at time t and reset to the value of zero at time $t+1$ (2 ms per time step). Furthermore, all the EPSPs that started before the spike were reset to zero, which mimics reduced excitability after spikes (Troyer and Miller 1997; Hausser et al. 2001). For each run, the membrane potential $y(t)$ was computed during 10,000 s of simulated time (5,000,000 time steps per run, 10.6 Hz output firing rate). The value of $\tilde{\vartheta} = 0.0649$ was estimated in ten preliminary runs by using linear regression between the true synaptic strengths w_i and the values computed with (7). In each of ten further runs the synaptic strengths w_i^* were computed according to (7). Eight synapses of the true strengths 0, 0.01, 0.02, 0.03, 0.04, 0.05, 0.06, and 0.07 were estimated (Fig. 2b) (mean of ten runs $\pm$ standard error of the mean). The mean relative error was defined as the mean absolute difference between the true synaptic strengths w_i and the estimated synaptic strengths w_i^* relative to the maximal synaptic strength of 0.07. This relative error was $1.7 \pm 1.1\%$ (mean $\pm$ standard deviation). The accuracy did not improve if (5) was explicitly solved for $\mathbf{w}$ (2.2% relative error), indicating that the approximation proposed in (6) is sufficiently accurate.

Note that the small nonzero synaptic strengths were consistently underestimated whereas the large synaptic strengths were consistently overestimated (Fig. 2). For the present simulations the subthreshold values of the membrane potential were not uniformly distributed between the reset potential and the firing threshold. Therefore, the chances for EPSPs to elicit action potentials were not exactly proportional to the synaptic strengths. Such linearity is assumed in (3) by requiring that the error term $e(\hat{t}, t)$ be independent of the synaptic strengths w_i. This linear approximation caused these small systematic errors. One could avoid these inaccuracies by replacing in (7) the constant parameter $\tilde{\vartheta}$ by a monotonously increasing function $\tilde{\vartheta}(w_i)$, which depends on the synaptic strength. The function $\tilde{\vartheta}(w_i)$ can be fitted from Fig. 2. To derive a learning rule without using this linear approximation, one would minimize the squared error $\sum_{t=1}^{N} (e(t))^2$ with respect to $\tilde{w}_i = w_i / \tilde{\vartheta}(w_i)$, estimate $\Delta\tilde{w}_i \cong c_i \sum_{t=1}^{N} y(t) x_i(t)$, and compute the synaptic strengths with $w_i = \tilde{w}_i \tilde{\vartheta}(w_i)$. In the current paper, this correction with $\tilde{\vartheta}(w_i)$ was not used, since the errors were already small (1.7%).

4.2 Illustration using online learning

To illustrate the information processing capability of one neuron trained with a physiologically plausible version of STDP, a model neuron is simulated that learns to predict its own postsynaptic spike. Five nonadaptive synapses are initialized with the strengths of 0.8, and 495 adaptive weights are initialized with small random strengths (uniformly distributed between 0 and 0.03). Poisson distributed presynaptic firing rates are set to 5 Hz. When the membrane potential reached a spike threshold of 0.1, it was set to the value of one at time t and reset to the reset potential of value zero at time $t+1$ (2 ms per time step). Based on experimental evidence, several update rules for spike-timing-dependent learning have been proposed (Rubin et al. 2001; Gutig et al. 2003). To avoid negative synaptic strengths, I used the rule $w_i = w_i + \eta(w_i^* - w_i) w_i$, with a desired value w_i^* set according to (7), and a learning rate $\eta = 0.3$. (To achieve rapid learning, the learning rate η was set to a value that was much larger than what is physiologically plausible.) This learning rule is consistent with empirical evidence, suggesting that the synaptic adaptation is proportional to the synaptic strength and that there is an upper limit for the synaptic strengths (Bi and Poo 1998).

During learning the neuron is forced to produce postsynaptic spikes by presynaptic spikes arriving at a set of nonadaptive and strong synapses. These strong synapses define the external process the neuron is supposed to learn and thus determine the desired neuronal output signal. These strong synapses are assumed to be driven by salient sensory stimuli. A second set of synapses is initially weak and learns to predict the postsynaptic spikes if their activation provides sufficient predictive information. These weak synapses may be activated by internal information arriving from higher cortical areas or by neighboring neurons in the same area. These weak synapses learn with the proposed version of STDP (Fig. 3a). In each of four simulated trials, all synapses are activated with the same temporal spike pattern, except that the strong synapses are not activated in trial four. In each of the first three trials, simultaneous activation of the sensory synapses elicits a postsynaptic spike, and the synaptic strengths of the internal synapses are adapted with the proposed learning rule. In trial four, the internal synapses elicit the postsynaptic spike, although the sensory synapses are not activated. The internal synapses learned to predict the occurrence of the postsynaptic spike (Fig. 3b), as the learning rule had increased the synaptic strengths of the few synapses that were consistently activated just before the output spike. In other words, the synaptic strengths of the internal synapses are adapted such that they are suitable to emulate the sensory experience driving the neuron in the first three trials. The postsynaptic spike elicited via the strong synapses thus serves as a teaching signal to adapt the synaptic strengths of the internal synapses. (The Matlab source code for the results of this paper is available at www.cnl.salk.edu/~suri/hebb.)

Note that this simulation serves only as an illustration due to two limitations. First, the synaptic strengths would only be guaranteed to converge to the desired values w_i* if the adaptive synapses did not themselves influence the output of the neuron. Second, the learning rule ensures that the synaptic strengths remain positive, although strictly positive correlations are not guaranteed.

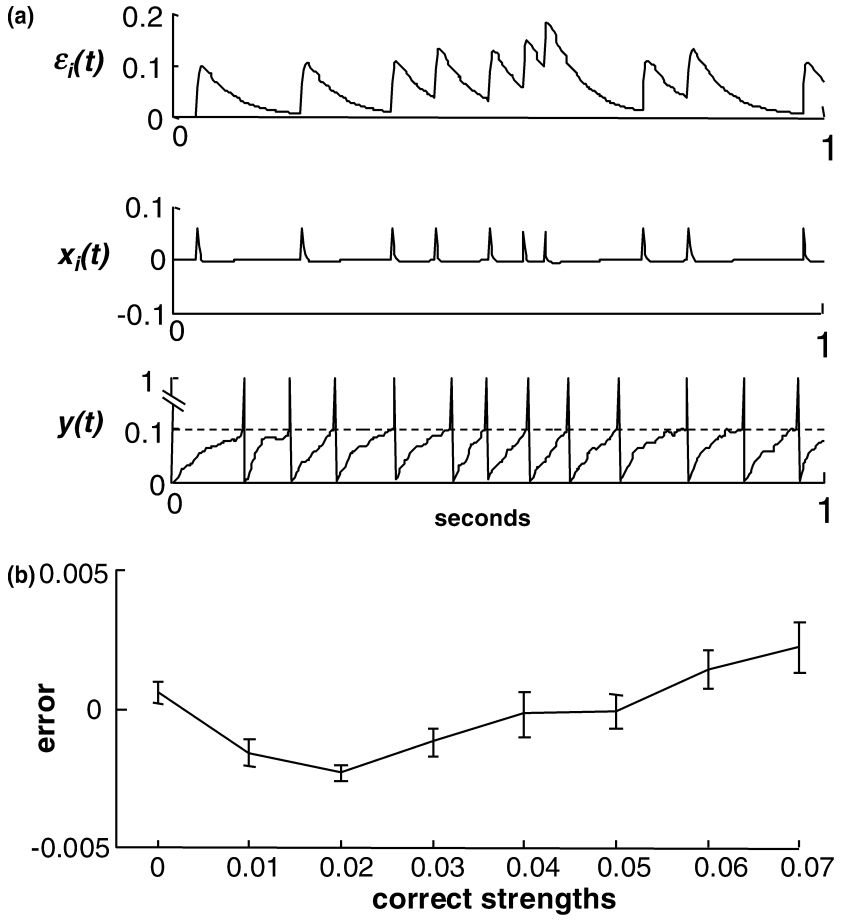

Fig. 2a, b. Accuracy of analytical result. **a** Typical time courses of the normalized and superimposed EPSP $\varepsilon_i(t)$ (*line 1*) and its derivative $x_i(t)$ (*line 2*) are shown for a synapse with number *i*. A typical time course of the membrane potential $y(t)$ is shown whereby the spike threshold is indicated by a *dashed line* (*line 3*, *spikes truncated*). **b** Correct synaptic strengths vs. error in estimated strengths. For a neuron with 500 synapses, the synaptic strengths were estimated with (7). Mean synaptic strengths were computed by averaging over ten runs. The errors of the mean synaptic strengths were defined as the differences between the correct synaptic strengths and the estimated synaptic strengths. These errors are shown for eight synapses of the correct strengths 0, 0.01, 0.02, 0.03, 0.04, 0.05, 0.06, and 0.07. *The error bars* indicate the standard error of the mean

5 Generalization to a network of neurons

These equations for one neuron generalize to a recurrent neural network. Therefore, the neurons of the "higher areas" (Fig. 3a) become a component of the modeled system. According to (3, 4), the membrane potential $y_k(t)$ of a neuron k can be written as the sum of the input signals $x_{ki}(t)$ weighted with the synaptic strengths w_{ki}

$$y_k(t) = \frac{1}{\tilde{\vartheta}_k} \sum_i w_{ki} x_{ki}(t) + e_k(t) \quad , \tag{9}$$

where $e_k(t)$ is an error term and $\tilde{\vartheta}_k$ is a scalar parameter. During learning the neural output signals $y_k(t)$ are assumed to be determined by a set of nonadaptive strong synapses. Since this assumption ensures that the neurons are independent during learning, adaptation of the synapses according to the proposed learning rule (2, 7, 8) minimizes the total error.

To achieve recurrent interactions, an operator D_{ki} is defined, which transforms output spikes to EPSPs and delays the signal to reflect delays due to the axonal spike propagation. Each input signal $x_{ki}(t+1)$ of a neuron k is computed from an output signal $y_i(t)$ of a neuron i with

$$x_{ki}(t+1) = D_{ki} y_i(t) \quad . \tag{10}$$

The time step is assumed to be sufficiently brief to ensure that the discrete-time representation corresponds to the continuous-time representation. It follows from (9, 10) that

$$y_k(t+1) = \frac{1}{\tilde{\vartheta}_k} \sum_i w_{ki} D_{ki} y_i(t) + e_k(t+1) \quad . \tag{11}$$

This equation describes the evolution of the network once the strong nonadaptive inputs are omitted. Assuming that the used approximations are accurate [(3) and (6)], (11) is a special case of the ARX model. Since the ARX model has been extensively analyzed (Ljung and Soderstrom 1983), I only briefly outline its most relevant properties. Equation (10) can be successively applied to estimate the future values of $y_k(t)$. Due to the delays in the linear operator D_{ki}, the neuronal output signals $y_k(t)$ depend on spikes that happened in the past. Let us assume that the values of $y_k(t)$ are determined by strong nonadaptive synapses during the beginning of a learned spike pattern over a time period that corresponds to the longest delay in the linear operator D_{ki}. Once these strong nonadaptive inputs are omitted, the network finishes the learned spike pattern, thereby recalling the memorized sequence.

Several conditions have to be fulfilled to guarantee correct learning of high-dimensional spike sequences by a network of neurons. (a) If the unsupervised evolution of the network is computed with (11), the proposed learning rule

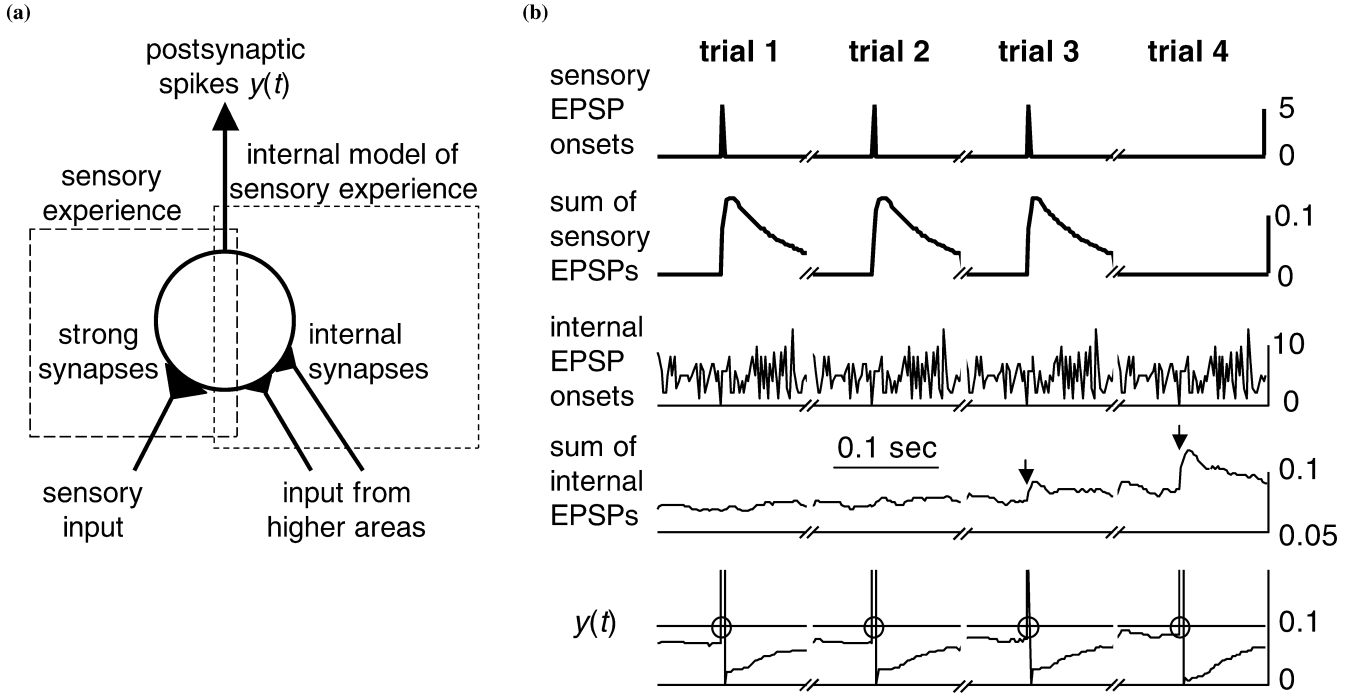

Fig. 3a, b. Model neuron learns to predict its own output spike with the proposed learning rule. **a** The neuron learns a model of its sensory experience. During learning, presynaptic spikes coding for salient sensory experience elicit postsynaptic spikes via strong synapses. Higher areas project via initially weak synapses to the same neuron. The appropriate internal synapses are adapted such that they learn to associate presynaptic spikes with postsynaptic spikes. After learning, the presynaptic spikes arriving from higher areas replace postsynaptic spikes if sensory inputs are omitted. The neuron thus learns an internal model of the relationship between presynaptic spikes coding for sensory input and postsynaptic spikes. **b** In trials 1, 2, and 3, the model neuron is driven by five EPSPs simultaneously arriving at five strong synapses (*lines 1 and 2*). In all four trials also 495 initially weak and adaptive synapses are excited by an internally generated spike pattern that is identical for each trial. For each time step of 2 ms, the number of EPSP onsets for these internal synapses is shown in *line 3*. The superposition of the internal EPSPs is shown in *line 4*. In the course of learning, the superposition of all the EPSPs generated by the internal synapses progressively increases before the postsynaptic spike (*line 4*, *arrows*). Although the sensory synapses are not activated in trial four, the internal synapses elicit a spike (*line 5* and spikes are truncated at the value of 0.2; the firing threshold is indicated by the *horizontal line*). As shown by *circles* that mark the same time within each trial, the spikes occur for each trial at almost the same time

is insufficient for the adaptation of the synaptic strengths, as correct adaptation of synaptic strengths would depend on other synaptic strengths [the gradient computed for (5) would become more complicated]. Therefore, learning may have to be avoided in this unsupervised mode. (b) The error term $e_k(t)$ is typically large. Large errors $e_k(t)$ cannot be avoided if a modeled system is probabilistic (Sect. 4.1). If the modeled system is not probabilistic but deterministic, these errors $e_k(t)$ become small if the synaptic strengths, or a synchronously activated set of synaptic strengths, are sufficiently large to reliably elicit postsynaptic spikes. However, in the case of small synaptic strengths and unsynchronized input spikes, the neuronal output signals $y_k(t)$ are only accurate predictions of a learned deterministic pattern if they are averaged over many trials or over a neuron population. (c) It is assumed that the spike probability is approximately proportional to the synaptic strength [(3), Sect. 4.1]. Although the linearity is currently needed to guarantee that a network minimizes the total error after training, slightly nonlinear neurons may be able to solve a much richer range of nonlinear system identification problems.

6 Discussion

These findings demonstrate for the first time that the information processing properties of neurons trained with a physiologically plausible version of spike-timing-dependent learning closely corresponds to those of correlation analysis, which is an established system identification method. A spike-timing-dependent learning rule is derived by assuming that the desired synaptic strengths minimize the errors between the actually occurring output spikes and the postsynaptic spikes predicted by the presynaptic spikes. According to this prediction error model of STDP, the synaptic strengths of neurons are adapted such that they predict their output spikes. The comparison of this prediction error model with system identification methods indicates that a temporal spike pattern is learned if a spike pattern is repeatedly elicited in a set of pyramidal neurons and recalled by presenting the beginning of the spike pattern. In these conditions, the same neurons serve as the input neurons during learning and as the output neurons during recall. The proposed prediction error model offers a mathematical foundation for previous

simulation studies that demonstrated that spike-timing-dependent learning can be used for the learning of spike sequences (Levy 1996; Gerstner and Abbott 1997; August and Levy 1999; Rao and Sejnowski 2001).

In the current study the spike-timing-dependent learning rule is represented as the cross correlation between representations of the neuron's input and output spikes. Without hampering the capability of the neurons for system identification, there is considerable flexibility on the functional form of spike representations for this cross-correlation rule, which influences the learning window. This robust performance is a direct consequence of spike coding: since the information is only coded by the time a spike occurs, a spike can be represented by any brief function that reliably indicates the time of spike occurrence.

In an independent study by Saudargiene et al. (2004), spike-timing-dependent learning was also represented by a cross-correlation rule. Their model suggests that spike-timing-dependent learning is proportional to the correlation between the NMDA conductance of the synapse and the temporal derivative of the membrane potential. According to the current findings, their model is likely to have the same system identification properties as the model proposed here. The proposed mathematical analysis can be applied to their model by defining the synaptic input $x_i(t)$ as the NMDA conductance and the neuronal output signal $y(t)$ as the derivative of the membrane potential.

Several experimental findings support the proposed cross-correlation hypothesis for adapting the synaptic strengths [(2)]. It was suggested that the adaptation of the synaptic strengths due to spike-timing-dependent learning may be mediated by the action potential that propagates back to the synapses (Watanabe et al. 2002). Such a backpropagating action potential would be particularly suitable to adapt the synaptic strengths according to a cross-correlation rule. Furthermore, the proposed cross-correlation hypothesis implies that long-term adaptation should be diminished proportionally to the amount of short-term depression, as synaptic short-term depression reduces the amplitude of the EPSP. This is consistent with experimental findings suggesting that the amplitudes of LTP and LTD are diminished proportionally to the amount of short-term depression (Froemke and Dan 2002). The integral of the LTP window appears to be smaller than that of the LTD window (Feldman 2000; Sjostrom et al. 2001), which seems to explain the finding that low-frequency pre-/postpairings with random delays induce LTD (Feldman 2000). This is consistent with the proposed cross-correlation rule because the correct synaptic strengths become zero if the presynaptic inputs are uncorrelated with the postsynaptic spikes.

The learning window of spike-timing-dependent learning is reproduced by representing presynaptic and postsynaptic spikes in the cross-correlation rule by the derivative of the normalized EPSP and the membrane potential, respectively. If this is taken as an exact description of spike-timing-dependent learning, elevated postsynaptic membrane potentials should boost its correlation with the EPSP derivative and thus enhance LTP (Fig. 1a). Indeed, LTP induction was found to be more effective if pairings are repeated at high frequencies such that the postsynaptic membrane potential does not repolarize back to rest between spikes (Markram et al. 1997; Sjostrom et al. 2001). Furthermore, LTP is also enhanced if the postsynaptic spike is preceded by a moderate depolarization, which can be provided by concurrent synaptic inputs or by current injection (Sjostrom et al. 2001). However, there is evidence suggesting that the chosen spike representations in the learning rule do not provide an exact description of spike-timing-dependent learning. For some pyramidal neurons the membrane potential remains elevated after spikes (Bi and Poo 1998), which would drastically change the learning window. Furthermore, the shape of the spike propagated back into the dendrite substantially differs from that in the cell body (Watanabe et al. 2002), suggesting that the representation of the postsynaptic spike, and thus the learning window, may depend on the location of the synapse (Saudargiene et al. 2004).

Previous studies have already attempted to demonstrate relationships of spike-timing-dependent learning with theoretical models of learning without reaching a consensus. Pfister et al. (2003) derived a spike-timing-dependent learning rule from an optimization criterion for a neuron with a noisy membrane potential that receives input by a single synapse. They assumed that the desired synaptic strengths maximize the likelihood of observing a postsynaptic spike train with a desired timing, given the firing rate. Like the current study, their results suggest that the LTD part of the learning window is a consequence of the hyperpolarization that often follows spikes. Unfortunately, their learning rule is not fully consistent with experimentally observed spike-timing-dependent learning, since their learning rule predicts LTD for unpaired spikes, and their learning window is wider than the experimentally established learning window. The latter seems to be a consequence of representing the input spikes by their EPSPs. Roberts (1999) demonstrated that a signal resembling the prediction signal of temporal difference (TD) learning can be acquired with spike-timing-dependent learning. Spike-timing-dependent learning was also related to TD learning because both use the time difference between the presynaptic spike and the postsynaptic spike (Rao and Sejnowski (2001)). However, the temporal difference learning rule proposed by Rao and Sejnowski (2001) is different from the TD learning rule if the presynaptic EPSP onset follows the occurrence of the postsynaptic spike: whereas this situation causes no change in the synaptic strength using TD learning, the temporal difference learning rule proposed by Rao and Sejnowski reduces the synaptic strengths. Indeed, these authors did not claim that any signal that is typical for TD learning can be learned with spike-timing-dependent learning. Simulation studies suggested various further functions of spike-timing-dependent learning including learning of precise temporal coding (Gerstner et al. 1996), spike synchronization (Suri and Sejnowski 2002), input correlations (Gutig et al. 2003), and spike sequences (Levy 1996; Gerstner and Abbott 1997; August and Levy 1999; Rao and Sejnowski 2001).

Assuming the accuracy of the used approximations (3, 6), the proposed neural network is mathematically equivalent to methods that estimate the parameters of systems of linear differential equations (Sect. 5, Appendix). Such algorithms have been used for computer vision, voice recognition, and image recognition and for learning internal models of temporal processes (Godfrey 1980; Ljung and Soderstrom 1983; Wolpert et al. 1995). Neural networks trained with such algorithms remove noise from familiar pictures and learn receptive fields resembling those of neurons in the visual cortex (Rao and Ballard 1997). How could such neural networks be used for computing predictions even though the temporal evolution of the signals of the internal model is only as fast as that of the observation? A possible solution is to train the model with a temporally compressed observation sequence. The finding that hippocampal spike patterns reflect temporally compressed representations of the animal's sensory experience (Skaggs et al. 1996) suggests that the cortex may use a similar strategy (August and Levy 1999).

Acknowledgements. I thank Mike Arnold, Jose Bargas, Andrew Barto, and Paola Fabrizio for their comments to earlier versions of the manuscript. This work was made possible by fellowships of the Swiss NSF and the EMDO Stiftung Zurich.

Appendix

Correlation analysis

I briefly describe the derivation of the well-established prediction error method correlation analysis used for system identification by following Ljung and Soderstrom (1983). This derivation is similar to that used for the proof of (2). Given a linear dynamic system

$$y(t) = \mathbf{x}^{\mathbf{T}}(t)\mathbf{w} + e(t) \ ,$$

where the system output $y(t)$ is a scalar signal, $\mathbf{x}^{\mathbf{T}}(t)\mathbf{w}$ is the scalar product between the system input signals $x_i(t)$ and system parameters w_i, and a scalar signal $e(t)$ is a noise term that does not need to be Gaussian distributed. The squared error term $e(t)$ is minimized with respect to the system parameters w_i with

$$\frac{\mathrm{d}}{\mathrm{d}\mathbf{w}} \frac{1}{N} \sum_{t=1}^{N} e(t)e(t) = \frac{\mathrm{d}}{\mathrm{d}\mathbf{w}} \frac{1}{N} \sum_{t=1}^{N} \left[y(t) - \mathbf{x}^{\mathbf{T}}(t)\mathbf{w} \right]^2 = 0 \ .$$

Computing the derivative of the middle term and resolving for $\mathbf{w}$ leads to

$$\mathbf{w} = \left[\frac{1}{N} \sum_{t=1}^{N} \mathbf{x}(t)\mathbf{x}^{\mathbf{T}}(t) \right]^{-1} \frac{1}{N} \sum_{t=1}^{N} \mathbf{x}(t) y(t) \ .$$

This solution for the model parameters $\mathbf{w}$ minimizes the least square prediction error $\sum_{t=1}^{N} (e(t))^2$ and can be seen as a special case of the maximum likelihood method. Correlation analysis can be used if the cross correlation between the input signals $x_i(t)$ is much smaller than their autocorrelation, i.e., $\sum_{t=1}^{N} x_i(t)x_k(t) << \sum_{t=1}^{N} x_i(t)x_i(t)$. In this case, w_i can be approximated with

$$w_i \cong \frac{\sum_{t=1}^{N} x_i(t) y(t)}{\sum_{t=1}^{N} (x_i(t))^2} \ .$$

Correlation analysis uses this equation to compute the system parameters w_i from the input signals $x_i(t)$ and output signal $y(t)$. This equation estimates the system parameters w_i as the slope of the linear regression computed for the N points $(x_i(t), y(t))$.

References

August DA, Levy WB (1999) Temporal sequence compression by an integrate-and-fire model of hippocampal area CA3. J Comput Neurosci 6(1):71–90

Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci 18(24): 10464–10472

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

Feldman DE (2000) Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. Neuron 27(1):45–56

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

Gerstner W (2001) Coding properties of spiking neurons: reverse and cross-correlations. Neural Netw 14(6–7):599–610

Gerstner W, Abbott LF (1997) Learning navigational maps through potentiation and modulation of hippocampal place cells. J Comput Neurosci 4(1):79–94

Gerstner W, van Hemmen JL (1992) Associative memory in a network of 'spiking' neurons, network. Comput Neural Syst 3:139–164

Gerstner W, van Hemmen JL (1994) Coding and information processing in neural networks. In: Domany E, van Hemmen JL, Schulten K (eds) Models of neural networks II. Springer, Berlin Heidelberg New York, pp 1–93

Gerstner W, Kistler WM (2002) Mathematical formulations of Hebbian Learning. Biol Cybern 87(5–6):404–415

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

Godfrey KR (1980) Correlation methods. Automatica 16:527–534

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

Hausser M, Major G, Stuart GJ (2001) Differential shunting of EPSPs by action potentials. Science 291(5501):138–141

Levy WB (1996) A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. Hippocampus 6(6):579–590

Ljung L, Soderstrom TL (1983) Theory and practice of recursive identification. MIT Press, Cambridge, MA

Markram H, Lubk, J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275(5297):213–215

Pfister J-P, Barber D, Gerstner W (2003) Optimal Hebbian learning: a probabilistic point of view. In: Kaynak et al (eds) Proceedings of joint international conference ICANN/ICONIP. Lecture notes in computer science, vol 2714. Springer, Berlin Heidelberg New York, pp 92–98

Rao RP, Ballard DH (1997) Dynamic model of visual recognition predicts neural response properties in the visual cortex. Neural Comput 9(4):721–763

Rao RP, Sejnowski TJ (2001) Spike-timing-dependent Hebbian plasticity as temporal difference learning. Neural Comput 13(10):2221–2237

Roberts PD (1999) Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. J Comput Neurosci 7(3):2235–2346

Rubin J, Lee DD, Sompolinsky H (2001) Equilibrium properties of temporally asymmetric Hebbian plasticity. Phys Rev Lett 86(2):364–367

Saudargiene A, Porr B, Worgotter F (2004) How the shape of pre- and postsynaptic signals can influence STDP: a biophysical model. Neural Comput 16(3):595–625

Sjostrom PJ, Turrigiano GG, Nelson SB (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 20 32(6):1149–1164

Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6(2):149–172

Suri RE (2002) TD models of reward predictive responses in dopamine neurons. Neural Netw 15(4–6):523–533

Suri RE, Sejnowski TJ (2002) Spike propagation synchronized by temporally asymmetric Hebbian learning. Biol Cybern 87(5–6):440–445

Sutton RS, Barto AG (1981) An adaptive network that constructs and uses an internal model of its world. Cogn Brain Theory 4:217–246

Troyer TW, Miller KD (1997) Physiological gain leads to high ISI variability in a simple network of a cortical regular spiking cell. Neural Comput 9(5):971–983

Watanabe S, Hoffman DA, Migliore M, Johnston D (2002) Dendritic K^+ channels contribute to spike-timing dependent long-term potentiation in hippocampal pyramidal neurons. Proc Natl Acad Sci USA 99:8366–8371

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. Science 269:1880–1882