

Spike-driven synaptic plasticity for learning correlated patterns of mean rates

S Fusi

Institute of Physiology, University of Bern, CH-3012 Bülhplatz 5, Switzerland

E-mail: fusi@cns.unibe.ch

Abstract. Long term synaptic changes induced by neural spike activity are believed to underlie learning and memory. Spike-driven long term synaptic plasticity has been investigated in simplified situations in which the patterns of mean rates to be encoded were statistically independent. An additional regulatory mechanism is required to extend the learning capability to more complex and natural stimuli. This mechanism can be provided by those effects of the action potentials that are believed to be responsible for spike-timing dependent plasticity. These effects, when combined with the dependence of synaptic plasticity on the post-synaptic depolarization, produce the non-monotonic learning rule needed for storing correlated patterns of mean rates.

Submitted to: *Rev. Neurosci.*

1. Introduction

It is widely believed that synaptic plasticity is the phenomenon underlying learning and memory. In the last years many models of spike-driven plastic synapses have been proposed (see e.g. [1, 2, 21, 23]). Most of these theoretical works were devoted to modeling experimental results on single pairs of synaptically connected neurons [2, 21, 23], to studying regulatory mechanisms based on spike-timing dependent plasticity [3, 29, 30, 27], or to relating the long term modifications to the neuronal activity to be encoded [1]. Only very recently a few investigators started to study the formation of memory and the collective dynamical behavior of networks of integrate-and-fire neurons connected by spike-driven plastic synapses [4, 5, 9]. Here we make a further step in the direction of understanding the computational relevance of spike-based synaptic dynamics by extending the model introduced in [1]. The Hebbian paradigm there implemented was good enough to embed in the synaptic matrix an extensive number of random, uncorrelated patterns of mean firing rates. Each stimulus was assumed to drive to elevated spike frequencies a randomly selected subset of cells. The spike activity of different cells was generated by injecting statistically independent currents into the pre and post-synaptic neurons, and hence the spike activity was asynchronous. After the removal of the stimulus, a randomly selected, low fraction of synapses underwent long term changes to encode the mean frequencies of the pre and post-synaptic neurons.

The random selection of a subset of synapses is needed to avoid too fast forgetting when the synaptic dynamics are realistic, i.e. when the variables describing the long term synaptic dynamics are bounded and the long term changes induced by the neural activity cannot be arbitrarily small [18]. Indeed, networks of neurons connected by this kind of realistic synapses show the palimpsest property, i.e. they forget automatically older stimuli whenever a new stimulus is learnt. If the fraction of modified synapses per presentation is large, then both learning and forgetting are fast and only a small number of patterns can be stored in memory. Otherwise, for slow learning, the optimal storage capacity can be recovered [6].

In [1] the stochastic selection of synapses was achieved by exploiting the irregularity of the spike-trains. The mean spike rates were encoded as follows: the pre-synaptic spikes were triggering the synaptic changes, and the post-synaptic sub-threshold depolarization was determining the direction of the changes, i.e. whether the synapse had to be potentiated or depressed. Reading the instantaneous sub-threshold depolarization of the post-synaptic neuron turned out to be a simple way, local in time, to encode the mean spike frequency. The updating rule resulting from this dynamics was good enough to store an extensive number of random patterns generated as explained above.

However, going beyond these simplistic patterns requires a regulatory mechanism that ensures a balanced redistribution of the memory resources among the different

stimuli to be stored[‡]. This is usually achieved by introducing a mechanism that guarantees a global normalization of the synaptic weights, either directly by evaluating the total synaptic strength, or indirectly, through the postsynaptic activity (see e.g. [8, 10]). One possible local mechanism is based on the idea of the perceptron learning rule [28] and has been more recently studied in [11, 12, 31]: the internal state of the synapse is changed according to a Hebbian paradigm only if it is necessary, i.e. if the pattern of activity to be encoded is not yet correctly memorized. For instance, when a stimulus to be learnt imposes a pattern of elevated activity to a subset of neurons, the synapses connecting these neurons are potentiated only if the activity of the post-synaptic (output) neuron is below some threshold value. Otherwise the synapse is left unchanged because the pattern of activity imposed by the external stimulus is already stable and the memory resources can be exploited in a more efficient way by other stimuli. Such a mechanism would prevent the system from over-learning and would decrease the forgetting rate by selecting a small fraction of those synapses that really need to be changed. Learning prescriptions based on this principle permit to memorize complex correlated patterns which are linearly separable [7, 11, 31], and to classify real data, like handwritten digits [12], provided that learning is slow, i.e. that every pattern has to be presented several times. These non-monotonic learning rules are achieved here by introducing in the spike-driven synaptic dynamics of [1], which depends solely on the subthreshold depolarization of the post-synaptic neuron, an additional dependence on the effects of the post-synaptic action potential. In particular an extra depression occurs when the pre-synaptic spike follows the post-synaptic action potential within some time interval. This combination of depolarization and spike-timing dependence makes the model more realistic and captures the observations of recent experiments on long term synaptic modifications [14]. The implementation of a non-monotonic learning rule in terms of detailed spike-drive synaptic dynamics is the first step to develop a full network which is capable of storing and classifying complex correlated patterns of mean rates.

2. The synaptic dynamics

The synaptic dynamics can be fully described in terms of a single internal variable X which represents the state of the synapse and determines the synaptic efficacy. In the absence of any pre or post-synaptic spike the synapse is made bistable by a recall force that drives the internal state to one of the two stable states: to $X = 1$, that corresponds to the upper bound, if X is above some threshold θ_X , or to the lower bound $X = 0$ if $X < \theta_X$. This bistability would be compatible with the experimental data of [13], which shows that synaptic potentiation is all-or-none, and with the models of molecular processes underlying LTP and LTD induction presented in [16].

Above θ_X the synapse is in a potentiated state and the synaptic efficacy on the

[‡] For instance when a group of binary synapses are modified to memorize a pattern, they become resources which are allocated for that specific pattern and cannot hold any information about other previously learned patterns.

post-synaptic neuron is elevated, whilst below θ_X the synapse is in a depressed state (low efficacy). These dynamics can be expressed by:

$$\frac{dX}{dt} = -\alpha\Theta(-x + \theta_X) + \beta\Theta(x - \theta_X)$$

where α and β represent the two refresh currents. Pre- and post-synaptic spikes trigger temporary modifications of the internal variable. They depend on the depolarization V of the post-synaptic neuron [1] and on the time that passed since the last occurrence of a pre or a post-synaptic spike [15, 2]. When the depolarization is high ($V > V_H$) every pre-synaptic spike pushes X upwards, weakly if the last post-synaptic spike has been emitted within the previous time interval T_- , strongly otherwise. In particular, if the pre-synaptic spike arrive at time t and k is the number of post-synaptic spikes emitted within the time interval $[t - T_-, t]$, then $X \rightarrow X + a - kb'$ upon the arrival of each pre-synaptic spike. Note that $a - kb'$ can become negative, and the synapse can be depressed even if the post-synaptic depolarization is above the threshold V_H . In the example that we will study in Section 3 we imposed an upper bound on k ($k \leq k_{max} = 2$) in order to have that when the depolarization is high, the temporary synaptic modifications are always positive ($a > k_{max}b'$). This upper bound might correspond to the saturation of the mechanism which is responsible for the extra negative jumps or to the existence of an absolute refractory period of the post-synaptic neuron which limits the maximal number of action potentials which can fall within T_- . In case k is larger than the maximal allowed value k_{max} , the number of effective negative jumps is set to k_{max} .

If the depolarization is low ($V < V_L$), the pre-synaptic spikes induce a depression which is either weak, if $k = 0$, or strong otherwise: each pre-synaptic spike induces a jump $X \rightarrow X - b - kb'$. Pre-synaptic spikes preceding post-synaptic action potentials within a time interval T_+ provoke small upwards jumps: $X \rightarrow X + a'$. The upper and the lower bound act as a rigid barrier and do not let the internal variable X leave the interval $[0, 1]$.

Table 1 summarizes the effects of the spikes and lists the corresponding probabilities of occurrence when the pre- and post-synaptic neurons fire asynchronously, with a Poisson distribution, at a mean frequency ν_{pre} and ν_{post} respectively. Q_a (Q_b) is the probability that $V > V_H$ ($V < V_L$) upon the arrival of a pre-synaptic spike. The modifications induced by the spikes leave a long term memory trace only if X crosses the threshold θ_X during the stimulation. Otherwise the state prior to the stimulation is restored by the refresh currents. If the activity of the pre and the post-synaptic neurons is highly irregular as in vivo cortical recordings, then the transitions from one stable state to another are stochastic and, at parity of mean pre and post-synaptic spike rates, occur with some probability. This provides a selection mechanism that chooses automatically only a fraction of synapses that are to be changed upon the presentation of a stimulus [1].

Table 1. Synaptic dynamics: the temporary modifications (upwards and downwards jumps) of the internal synaptic variable X induced by pre and post-synaptic spikes.

Effects of a pre-synaptic spike that arrives at time t_{pre}		
Condition	Effect	Prob. of occurrence
$V > V_H$ and no post-syn. spike in $[t_{pre} - T_-, t_{pre}]$	$X \rightarrow X + a$	$\nu_{pre} Q_a e^{-\nu_{post} T_-}$
$V > V_H$ and k post-syn. spikes in $[t_{pre} - T_-, t_{pre}]$	$X \rightarrow X + a - kb'$	$\nu_{pre} Q_a e^{-\nu_{post} T_-} \frac{(\nu_{post} T_-)^k}{k!}$
$V < V_L$ and no post-syn. spike in $[t_{pre} - T_-, t_{pre}]$	$X \rightarrow X - b$	$\nu_{pre} Q_b e^{-\nu_{post} T_-}$
$V < V_L$ and k post-syn. spikes in $[t_{pre} - T_-, t_{pre}]$	$X \rightarrow X - b - kb'$	$\nu_{pre} Q_b e^{-\nu_{post} T_-} \frac{(\nu_{post} T_-)^k}{k!}$
Effects of a post-synaptic spike that arrives at time t_{post}		
k pre-syn. spikes in $[t_{post} - T_+, t_{post}]$	$X \rightarrow X + ka'$	$\nu_{post} e^{-\nu_{pre} T_+} \frac{(\nu_{pre} T_+)^k}{k!}$

2.1. Mean spike rates vs the distribution of the subthreshold depolarization

The probabilities of occurrence of the temporary modifications a, b, a', b' control the direction in which the synapse is modified by the activity of the pre- and post-synaptic neurons. Q_a and Q_b depend on the statistics of the depolarization of the post-synaptic neuron under stimulation and can be calculated analytically when the model of the neuron is simple and the pre and post-synaptic activities are statistically independent. As for the stimulus we assume that each stimulus imposes a pattern of asynchronous activity by changing the statistics of the input current of a subset of cells (see below). For the neuron we will adopt the linear integrate-and-fire (LIF) model introduced in [17]. The neuron integrates linearly the total synaptic current until the depolarization V crosses a threshold θ and a spike is emitted. Following a refractory period τ_r , the neuron is reset to $V = H$, from where it starts again integrating the current. The depolarization is limited from below by a rigid barrier that does not allow the depolarization to go below the resting potential (here arbitrarily set to $V = 0$). When such a neuron is embedded in a large network and is subject to the heavy bombardment of statistically independent synaptic inputs, the total somatic current can be emulated by a stochastic process known as Ornstein-Uhlenbeck process [19]. Such an *in vivo*-like current is Gauss distributed and fully characterized by its time correlation length τ , and its mean μ and its variance σ^2 per unit time. For simplicity we assume that the current is delta-correlated in time ($\tau = 0$). When injected such a current, the LIF neuron responds as real pyramidal cells [20].

The stationary distribution $p(v)$ can be computed as in [17, 1] and is given by:

$$p(v) = \frac{\nu}{\mu} \left[\Theta(v - H) \left(1 - e^{-\frac{2\mu}{\sigma^2}(\theta - v)} \right) + \Theta(H - v) \left(e^{-\frac{2\mu}{\sigma^2}H} - e^{-\frac{2\mu}{\sigma^2}\theta} \right) e^{\frac{2\mu}{\sigma^2}v} \right]$$

where ν is the mean firing frequency:

$$\nu = \left[\tau_r + \frac{\sigma^2}{2\mu^2} \left(e^{-\frac{2\mu\theta}{\sigma^2}} - e^{-\frac{2\mu H}{\sigma^2}} \right) + \frac{\theta - H}{\mu} \right]^{-1}$$

τ_r is the absolute refractory period. Q_a and Q_b are given by the integral of $p(v)$ in the interval $[V_H, \theta]$ and $[0, V_L]$ respectively. It is straightforward to compute these integrals analytically. They are reported in Appendix A for convenience.

μ and σ characterize the synaptic input and depend on the network interactions. Each pair μ, σ determines unequivocally a mean spike rate ν . However the opposite is not true: for each ν there exists an entire curve in the μ, σ space along which the mean firing rate is constant. As a consequence, we need to specify a protocol for driving the post-synaptic neuron to a frequency ν_{post} and determine unequivocally the statistics of the distribution of the depolarization for every ν_{post} . Here we assume that ν_{post} is changed by increasing or decreasing the average spike frequency ν_{ext} of a subpopulation of pre-synaptic neurons that represent the external (sensory) input to the network (similarly to [1]). If the recurrent feedback of the post-synaptic neurons does not affect much the network activity, then the parameters of the input current move along a linear trajectory in the (μ, σ^2) space:

$$\begin{aligned} \mu &= J_E N_E \nu_{ext} + \mu_0^E - \mu^I \\ \sigma^2 &= J_E^2 N_E \nu_{ext} + \sigma_{E0}^2 + \sigma_{I0}^2 \end{aligned}$$

where J_E is the mean synaptic efficacy of the spikes emitted by the N_E external neurons, and μ_0^E, μ_0^I are the mean synaptic currents due to the spontaneous activity of excitatory and inhibitory cells respectively. σ_{E0}^2 and σ_{I0}^2 are the corresponding variances per unit time of the stochastic process that model the total input current. Instead of ν_{ext} , we can choose μ as an independent parameter, and $\sigma^2 = J\mu + K$ (in the following analysis: $J = 0.02\theta, K = 0.01\theta^2/ms$). If in a spontaneous activity state the recurrent input is as large as the external input, we have that $J = J_E$ (the average coupling between all the excitatory neurons) and $K = \nu_0^I N_I J_I (J_I + J_E)$, where ν_0^I is the spontaneous activity of the N_I inhibitory neurons that are projecting onto the post-synaptic cell (mean coupling J_I). For each desired output frequency ν_{post} , μ is tuned to determine the suitable pair (μ, σ) that drives the post-synaptic neuron at a firing rate ν_{post} .

The distribution of the depolarization and the probabilities Q_a and Q_b are plotted in Fig. 1. As the external stimulus increases the output frequency ν_{post} , the distribution of the depolarization V changes in such a way that Q_b decreases and Q_a increases. The introduction of a reset potential H which is distinct from the resting potential ($H = 0.7\theta \neq 0$) moves the maximum of the distribution of V to H for high rates and




Figure 1. Mean spike rate vs the distribution of the depolarization V . Left: Distribution $p(v)$ as a function of the post-synaptic frequency ν_{post} . For each ν_{post} the parameters μ, σ characterizing the input current are calculated as explained in the text. μ, σ determine the subthreshold distribution of the depolarization that is plotted here. The white lines over the surface are drawn in correspondence of $V = V_L$ and $V = V_H$, the thresholds that determine the direction of the temporary synaptic modifications. Note that here the reset potential H coincides with V_H ($H = V_H = 0.7\theta$). $\tau_r = 2ms$. Right: Probability of occurrence of upwards (Q_a) and downwards (Q_b) jumps upon the arrival of a pre-synaptic spike for different post-synaptic activities. Q_a is the integral of $p(v)$ between the white line corresponding to V_H and the threshold. Analogously, Q_b is the integral of $p(v)$ between the resting potential ($V = 0$) and $V_L = 0.35\theta$. As the post-synaptic activity increases, the peak of the distribution $p(v)$ moves from the resting potential to the reset potential H . As a consequence, Q_a increases and Q_b decreases.

therefore makes the distribution more suitable for distinguishing low and high post-synaptic spike rates. Finding the proper parameters of the synaptic dynamics is much easier than in [1] where H coincided with the resting potential. When the pre and post-synaptic neurons are simultaneously activated, most of the numerous temporary modifications of the synapse are upwards and the synapse tends to potentiate. On the contrary, when the pre-synaptic neuron is active and the post-synaptic neuron has low spontaneous activity, most of the jumps are downwards and the synapse is eventually depressed. When ν_{pre} is low, the few jumps per unit time are usually not enough to drive the internal state across the threshold θ_x , and no transition occurs. It has been proved that this scheme of modifications (a Hebbian paradigm) maximizes the signal of the stimuli that can be extracted by the network dynamics from the synaptic matrix [6, 18] and actually implements a classical covariance updating rule (see e.g. [22, 10]). Hence the dependence on the post-synaptic depolarization expressed by Q_a and Q_b contain most of the relevant information to encode the patterns of mean spike rates imposed by

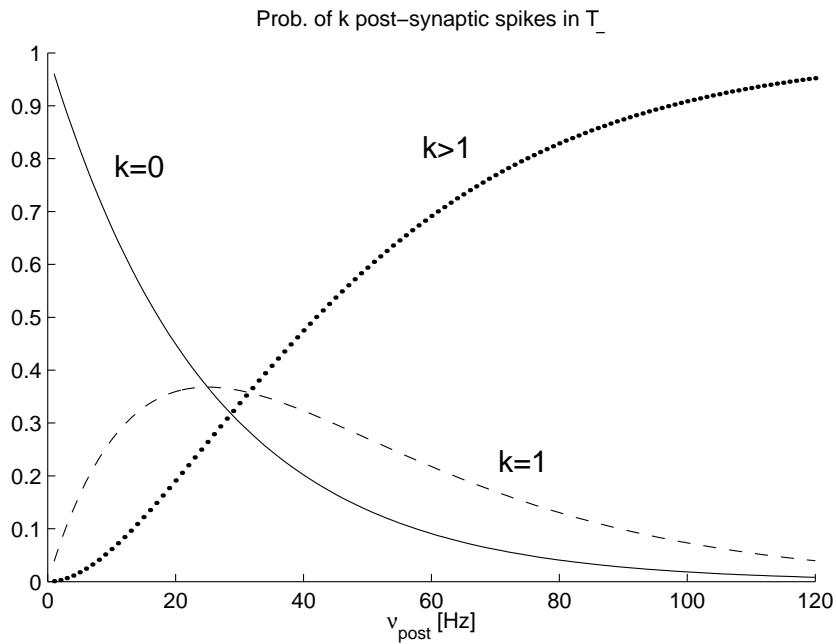


Figure 2. Probability that k post-synaptic spikes are within a time window T_- . It determines the depressing and regulatory effects of the action potential on the synaptic dynamics. The curve $k = 0$ gives the fraction of cases in which the temporary synaptic modifications are entirely controlled by the depolarization of the post-synaptic neurons. As the post-synaptic frequency increases, the additional effect of action potentials becomes predominant. Were it only for Q_b , the maximal LTD would be always achieved at $\nu_{post} = 0$. The effect of the extra negative jumps is twofold: it stop LTP when the post-synaptic frequency is too high, and it moves the maximal LTD to higher frequencies in order to reduce when depression when the neuron is already firing at low frequencies.

the stimulus.

2.2. The additional effect of the action potentials

The extra jumps (a' , b') implement the part of the dynamics that depend on the relative spike-timing. When the pre and post-synaptic spike activities are asynchronous the probability of occurrence of these jumps depends essentially on the product of the mean frequencies of the pre and the post-synaptic neuron. The ratio between the number of potentiating events and the number of depressing events is roughly constant when the pre- and the post-synaptic frequencies are multiplied by the same factor. This activity dependence does not contribute much to the signal (see e.g. [18] for the exact definition) left by the mean spike rates, but has an important regulatory role. Indeed, if $T_- \gg T_+$, as the frequency of the post-synaptic neuron increases, the fraction of cases in which the direction of the temporary modifications is controlled by the depolarization V decreases and depression becomes dominant, especially if one considers the possibility of multiple spikes within the time interval T_- (see Figure 2). In what follows we will assume that $T_+ = 0$. The extra jumps a' actually contribute neither to the signal nor to

the regulatory mechanism. Indeed for the signal the dependence on the sub-threshold depolarization is sufficient, and for implementing the non-monotonic behavior the effect of the extra jumps b' is enough. Hence, we will ignore the temporary jumps a' in the following analysis.

3. The non-monotonic learning rule

The next step is to compute the probability that a stimulus induces a stable modification in the synaptic couplings, i.e. that during the stimulation the synapse makes a transition from one stable state to another stable state. The transition probability is defined as the fraction of cases, out of a large number of repetitions of the same stimulation conditions, in which at the end of the stimulation, the synapse made a transition to a state different from its original state. These probabilities control the learning process. Low transition probabilities correspond to a slow learning scenario in which the stimuli need to be presented a large number of times to be learnt. A high number of repetitions is the price to be paid if one wants an optimal redistribution of the memory synaptic resources among the stimuli to be learnt [6, 11, 12]. In fact, high transition probabilities allow fast acquisition of information about the stimuli, but limit dramatically the storage capacity of the network because past stimuli are forgotten very quickly [6].

For uncorrelated patterns of activities the optimal storage capacity is achieved when the coding level is low (a small fraction of neurons is activated by each stimulus) and the LTD probability is small enough, compared to the LTP probability, to preserve the balance between potentiations and depressions [6]. For correlated patterns a control mechanism is needed to stop learning in case the stimulus is already correctly encoded in the synaptic matrix [11, 12]. This would require a non-monotonic learning rule, i.e. a U-shaped dependence of LTP and LTD probabilities on the post-synaptic mean-frequency.

3.1. Solving the long term synaptic dynamics

The transition probabilities can be approximately calculated by using the density approach introduced in [1]. Similar approaches can be found in [29, 30]. Since the pre and the post-synaptic neuron fire asynchronously, the Takács equations that govern the density function $p(X, t)$ of the synaptic internal variable can be easily generalized to incorporate the additional effects of a' and b' §. The equations are reported in Appendix B. For each ν_{post} the statistics of the input current (μ, σ) that produces the desired mean spike frequency is determined as explained in section 2.1. The distribution of the depolarization is then unequivocally determined, and Q_a and Q_b can be computed as in Appendix A. The initial probability density function of the internal synaptic variable X is set to a $\delta(X)$ when the synapse starts from the depressed state $X = 0$, and to

§ The equations can also be generalized to incorporate the contribution of synchrony between the pre and the post-synaptic spike trains as in [30]

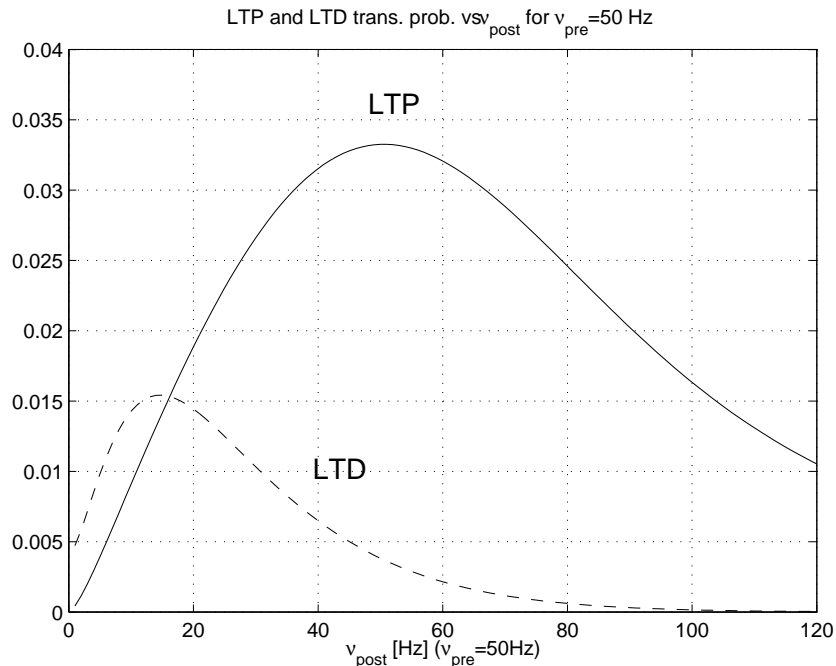


Figure 3. LTP and LTD transition probabilities as a function of post-synaptic frequency when the pre-synaptic neuron is active ($\nu_{pre} = 50\text{Hz}$). To each post-synaptic frequency corresponds a pair μ, σ characterizing the input current. This pair determines the distribution of the depolarization as explained in section 2.1. The synaptic parameters are: the temporary synaptic modifications $a = 0.26$, $b = 0.085$, $a' = 0$, $b' = 0.09$; the time window for the additional effect of the post-synaptic action potential $T_- = 40\text{ms}$; the synaptic threshold $\theta_x = 0.5$; the refresh currents $\alpha = 0.003\text{ms}^{-1}$, $\beta = 0.008\text{ms}^{-1}$.

$\delta(X - 1)$ when the initial condition is a potentiated state $X = 1$. The Takács equations are integrated numerically throughout the stimulation interval (in our case 250ms). The LTP (LTD) transition probability is given by the final probability that the synapse is above (below) the internal threshold θ_X when the synapse started from the depressed (potentiated) state before the arrival of the stimulus.

3.2. The transition probabilities

The numerical solution of the density equations gives the LTP and LTD transition probabilities for any pair of ν_{pre}, ν_{post} (see Fig. 3). The parameters of the synaptic dynamics can be easily tuned to achieve: 1) LTP and LTD transition probabilities that maximize the signal embedded in the synaptic matrix by each pattern of activity (LTP for high ν_{post} and LTD for low ν_{post}); 2) a ratio between LTP and LTD that ensures the balance of total number of potentiations and depressions when the mean number of neurons activated by each stimulus is low (low coding level); 3) the preservation of memory on time scales that are much longer than the inherent time constants of the synapse (the transitions in case of spontaneous activity of both pre- and post-synaptic

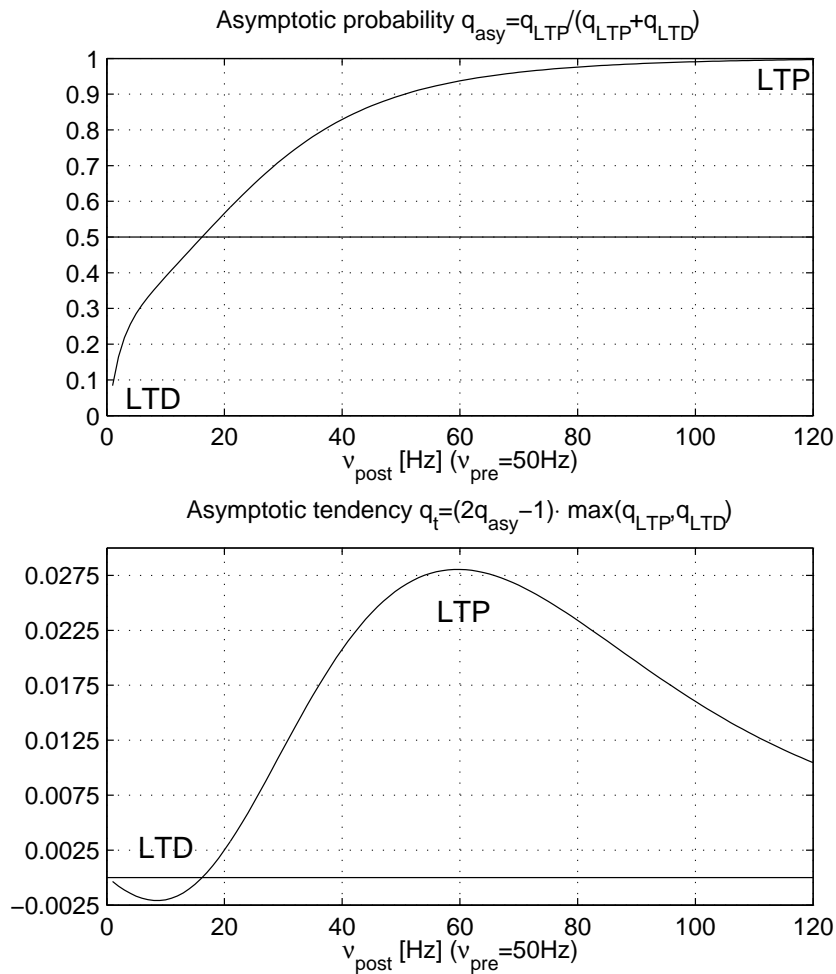


Figure 4. The asymptotic probability that the synapse is potentiated (top) and the asymptotic tendency of having either LTP or LTD (bottom) as a function of the post-synaptic firing frequency when the pre-synaptic neurons is active ($\nu_{pre} = 50$ Hz). The asymptotic probability $q_{LTP} / (q_{LTP} + q_{LTD})$ is the probability that a synapse ends up in the potentiated state after an infinite number of repetitions of the same stimulus. This quantity expresses whether LTP ($q_{asy} > 1/2$) or LTD dominates ($q_{asy} < 1/2$). For post-synaptic frequencies above $\nu_{post} \simeq 15$ Hz, LTP always dominates and the LTP transition probabilities are systematically much larger than the LTD probabilities (vice versa below 15 Hz). However, if the transition probabilities are very small, q_{asy} is not a good indicator because the synapse does not change anyway and neither LTP nor LTD actually dominate. For this reason we plotted the asymptotic tendency, i.e. the q_{asy} minus 1/2 (the baseline corresponding to a perfect balance between LTP and LTD) weighted by the maximum between q_{LTP} and q_{LTD} . From this second plot it is clear that for high post-synaptic frequencies, although LTP clearly dominates, the transition probabilities become smaller and smaller. Analogously, for low ν_{post} , the LTD probabilities become negligible. Note that, up to ~ 60 Hz, the behavior resembles the one of the BCM rule [26].

neurons are negligible); 4) a regulatory mechanism that blocks LTP when ν_{post} is too high, and stops LTD when ν_{post} is too low. This is a possible implementation of the mechanism required to store correlated patterns [11, 12, 31]. It also resembles the one in the BCM rule [26] and it is essentially based on the additional effect of the action potential that signals when the activity of the post-synaptic neuron is too elevated or too low. To illustrate all these properties we plot the LTP and LTD transition probabilities in three different ways (see Figs. 3,4,5). Figure 3 shows the raw transition probabilities when the pre-synaptic neuron is active. Interestingly, both LTP and LTD curves are non-monotonic. The peak of LTD probabilities is around $\nu_{post} \simeq 10Hz$, while the maximum LTP occurs at much higher frequencies: $\nu_{post} \simeq 50Hz$. The peak LTD would be at $\nu_{post} = 0$ if only the depolarization of the post-synaptic neuron is read. The extra effects of the action potentials are negligible for very low post-synaptic frequencies, when the probability that the neuron is hyperpolarized is maximal. However they increase quickly with ν_{post} , and hence the maximum of LTD is moved towards higher post-synaptic frequencies. As ν_{post} increases, the probability that more than a post-synaptic spike falls into the time window T_- becomes not negligible and it eventually dominates the synaptic dynamics. In such a regime, every pre-synaptic spike triggers multiple downwards jumps, and the probability of LTP decreases. The fact that for low ν_{post} LTD dominates, and for high ν_{post} LTP prevails, is what is required to implement a Hebbian paradigm, or a covariance-rule for updating the synapse (see e.g. [22, 8, 10]). The non-monotonicity provides the regulatory mechanism that stops potentiation when too many synapses are already contributing in the same way in activating the post-synaptic neuron. This behavior is shown also in Figure 4 where we plotted the asymptotic probability of ending up in a potentiated state and the asymptotic tendency expressing whether LTP or LTD dominates. The asymptotic probability is defined as the probability that the synapse is in a potentiated state following an infinite number of presentations of the same stimulus. It is plotted in Figure 4(top). When the transition probabilities are too small (e.g. in the presence of spontaneous activity) there is no real tendency of having either LTP or LTD since the synapses remain unchanged even after a very large number of presentations. For this reason the asymptotic probability, minus the baseline $1/2$ that corresponds to a perfect balance between LTP and LTD, is weighted by the maximum between LTP and LTD transition probabilities. This quantity, defined as the asymptotic tendency, is plotted in Figure 4(bottom) as a function of the post-synaptic firing frequency. The minimum, corresponding to a dominating LTD, is around $\nu_{post} = 10Hz$ while the maximum probability of potentiating a synapse is around $\nu_{post} = 50Hz$. For both too low or too high post-synaptic frequencies the asymptotic tendency curve goes to 0, indicating that learning is blocked.

Finally we show in Figure 5 the effect of the strong non-linearities due to the synaptic threshold θ_X . When the pre-synaptic neuron fires at spontaneous rates, the LTP and LTD transitions probabilities are several orders of magnitude smaller than the corresponding probabilities under stimulation ($\nu_{pre} = 50Hz$). This allows to preserve the memory on time scales ($\sim 10^2$ days) that are much longer than the longest inherent

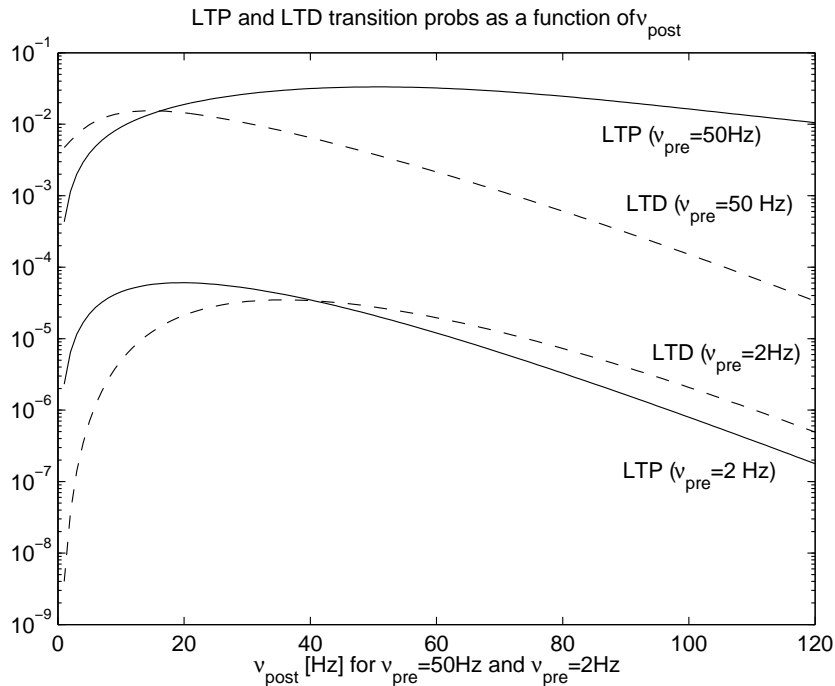


Figure 5. Memory preservation in the presence of spontaneous activity: LTP and LTD transition probabilities as a function of the post-synaptic frequency in two cases: when the pre-synaptic neuron is active ($\nu_{pre} = 50\text{Hz}$, the two topmost curves) and when it has only spontaneous activity (2Hz, the two curves at the bottom). The transition probabilities drop sharply when the activity of the pre-synaptic neuron is low (note that the scale is logarithmic). The synaptic threshold θ_X protects memory against noisy, spontaneous fluctuations of the neural activity.

time constant of the synapse ($1/\alpha \simeq 300\text{ms}$).

4. Discussion

The synaptic dynamics introduced here implements the updating rule required for storing complex, correlated patterns of asynchronous activity [11, 12, 31]. Synaptic potentiation becomes unlikely when the post-synaptic activity is too high, and synaptic depressions are blocked if the post-synaptic spike rate is too low. The pre-synaptic neuron triggers temporary modifications, while the post-synaptic activity determines the direction of the long term change.

The mean spike rate of the post-synaptic neuron is encoded by reading the instantaneous subthreshold depolarization of the post-synaptic cell. A single reading of the instantaneous depolarization does not contain much information about the post-synaptic mean rate. However the required information is distributed across several neurons that are driven to the same activity by the stimulus. This means that even a single instantaneous reading of the depolarization of a population of cells contains all the information about the mean spike rate and many other statistical properties (e.g.

the statistics of the input current) of the activity of these neurons. The regulatory mechanism is provided by the effects of the post-synaptic action potential, that, for instance, signals when the post-synaptic neuron is too active. One might wonder whether: 1) it is possible to have the regulatory mechanism without making recourse to a dynamics based on spike-timing; 2) it is possible to encode the post-synaptic spike frequency without reading the depolarization. The answer is positive to both questions.

As for the first issue, a regulatory mechanism without explicit spike timing that makes the LTP and LTD curves as a function of ν_{post} non-monotonic is actually possible. For instance it can be achieved by choosing a depolarization range for upwards jumps that starts from V_H and ends at a value V'_H below the reset potential H . As the mean spike rate of the post-synaptic neuron increases, the interval $[V_H, V'_H]$ is first filled in with more and more realizations, but eventually it becomes empty because the post-synaptic neuron spends most of its time between the reset potential and the threshold when it is driven by the mean current to a supra-threshold regime. As a consequence the LTP probability at some point starts to decrease. Unfortunately this non-monotonicity cannot be easily controlled and fine tuning of the parameters defining the depolarization intervals is required. These difficulties convinced us that it is necessary to introduce some additional mechanism to provide a robust regulation of the synaptic weights. The effect of the post-synaptic action potential appeared as a natural candidate, since it is known to have a dramatic effect on many internal variables, such as calcium concentration, that control the expression of LTP and LTD.

The second issue concerns the possibility of using a synaptic dynamics entirely based on spike-timing. Provided that the time windows are long enough [4] or that strong non-linearities are introduced in the synaptic dynamics [32], it is possible to encode mean rates by implementing a covariance rule. However reading the depolarization seems to be a much simpler solution that has probably not been overlooked by biology. Indeed recent experiments [14] show that the protocols to induce LTP and LTD, even if entirely based on spike-timing, are very much affected by the depolarization of the post-synaptic neuron. For instance, by hyperpolarizing the post-synaptic neuron in the interval between successive action potentials one can induce LTD or block the induction of LTP, even if the pre-synaptic spikes always precede the post-synaptic potentials as required by the spike-timing protocol introduced in [15].

Interestingly the regulatory mechanism studied here was already present in the model of [1], provided that during the refractory period the post-synaptic neuron is considered to be in a state that leads only to downwards jumps (e.g. the depolarization is set to a value below V_L during the refractoriness) \parallel . The refractory period would be equivalent to the time interval T_- that always follows the emission of a post-synaptic action potential. Unfortunately the refractory period affects also the neural dynamics, and cannot be tuned as an independent parameter. This makes the search for good synaptic parameters more difficult.

\parallel I am grateful to D.J. Amit and G. Mongillo for this observation

The spike-driven synaptic dynamics introduced here is a simple implementation of schematic learning rules which guarantee the convergence of learning in the case of patterns of mean rates which are linearly separable [4, 31]. Whether it is actually possible to store and classify these patterns in a full network of integrate-and-fire neurons will be the next step and will be studied elsewhere [24].

Appendix A

The distribution $p(v)$ of the depolarization of a LIF neuron when a delta-correlated Gauss current is injected into the cell can be derived from the stationary solution of the density equation (see e.g.[17]). The expression for $p(v)$ is already present in the text, here it is reported again for convenience:

$$p(v) = \frac{\nu}{\mu} \left[\Theta(v - H) \left(1 - e^{-\frac{2\mu}{\sigma^2}(\theta - v)} \right) + \Theta(H - v) \left(e^{-\frac{2\mu}{\sigma^2}H} - e^{-\frac{2\mu}{\sigma^2}\theta} \right) e^{\frac{2\mu}{\sigma^2}v} \right]$$

The probability Q_a is the integral of $p(v)$ between V_H and the threshold for emitting a spike θ . Q_b is the integral between the resting potential and V_L . The integral over a generic interval V_1, V_2 is given by:

$$\int_{V_1}^{V_2} p(v) dv = \begin{cases} \frac{\nu\sigma^2}{2\mu^2} (e^{-sH} - e^{-s\theta}) (e^{sV_2} - e^{sV_1}) & \text{for } V_2 < H \\ \frac{\nu}{\mu} \left[V_2 - V_1 - \frac{\sigma^2}{2\mu} \left(e^{s(V_2 - \theta)} - e^{s(V_1 - \theta)} \right) \right] & \text{for } V_1 > H \\ \frac{\nu\sigma^2}{2\mu^2} (e^{-sH} - e^{-s\theta}) (e^{sH} - e^{sV_1}) + \\ + \frac{\nu}{\mu} \left[V_2 - H - \frac{\sigma^2}{2\mu} \left(e^{s(V_2 - \theta)} - e^{s(H - \theta)} \right) \right] & \text{elsewhere} \end{cases} \quad (1)$$

where $V_2 > V_1$ and $s = 2\mu/\sigma^2$.

Appendix B

The distribution of the synaptic internal variable X is described by the density $p(X, t)$ for the interval $]0, 1[$ and by the two discrete probabilities $P_0(t)$ and $P_1(t)$ of being at the reflecting barriers $X = 0$ and $X = 1$ respectively. The equations that determine the time development of the distribution are a generalization of those introduced in [1]. For example, if the jump a' is ignored, the equation for $p(X, t)$ is:

$$\begin{aligned} \frac{\partial p(X, t)}{\partial t} &= \alpha \frac{\partial p(X, t)}{\partial^{(+)} X} + \nu_{pre} [A(X, t) - p(X, t)] \quad \text{if } X \in]0, \theta_X[\\ \frac{\partial p(X, t)}{\partial t} &= -\beta \frac{\partial p(X, t)}{\partial^{(-)} X} + \nu_{pre} [B(X, t) - p(X, t)] \quad \text{if } X \in]\theta_X, 1[\end{aligned}$$

where the terms with α and β account for the refresh currents that preserve the memory in the absence of stimuli and A and B contain the effects of the temporary changes

induced by the pre-synaptic spikes. The expressions of A and B should contain now all the possible jumps ($a, b, a - b', a - 2b', \dots, b - b', b - 2b', \dots$). The final equations are:

$$\begin{aligned}
A(X, t) &= Q_a P_0(t) \left[\delta(X - a) Q'_b(0) + \sum_{k=1}^{\infty} \delta(X - a + kb') Q'_b(k) \right] \\
&\quad + Q_a \left[p(X - a, t) Q'_b(0) + \sum_{k=1}^{\infty} p(X - a + kb') Q'_b(k) \right] \\
&\quad + Q_b \left[p(X + b, t) Q'_b(0) + \sum_{k=1}^{\infty} p(X + b + kb', t) Q'_b(k) \right] \\
B(X, t) &= Q_a \left[p(X - a, t) Q'_b(0) + \sum_{k=1}^{\infty} p(X - a + kb') Q'_b(k) \right] \\
&\quad + Q_b [p(X + b, t) Q'_b(0) + \sum_{k=1}^{\infty} p(X + b + kb', t) Q'_b(k)] \\
&\quad + Q_b P_1(t) [\delta(X - (1 - b)) Q'_b(0) + \sum_{k=1}^{\infty} \delta(X - (1 - b - kb')) Q'_b(k)]
\end{aligned}$$

where $Q'_b(k)$ is the probability that k post-synaptic spikes occurred in the time interval $[t - T_-, t]$ (t is the time of the arrival of a pre-synaptic spike). For a Poisson process it is given by:

$$Q'_b(k) = e^{-\nu_{post} T_-} \frac{(\nu_{post} T_-)^k}{k!}$$

We assume that $p(x, t) = 0$ for $x \notin [0, 1]$. The sums in A and B over k go from 1 to infinity, but it is reasonable to keep only a few terms (in our case we kept only 2 terms because $k_{max} = 2$). When a finite number of terms is kept, then the contribution of the $k = k_{max}$ term is weighted by the probability of having k_{max} spikes or more (e.g. for b' jumps: $\sum_k = k_{max}^{\infty} Q'_b(k)$). In fact for high frequencies the intervals between successive spikes of the post-synaptic neurons are not Poisson distributed and the refractory period imposes an upper bound on the number of post-synaptic action potentials that can fall in the time window T_- . Moreover the upper bound on k might be due to some inherent saturation mechanism. The dynamics of the probabilities P_0 and P_1 are modified in a similar way. All the equations are approximated because they do not consider the temporal auto-correlations of the post-synaptic depolarization (see [1] for more details).

Acknowledgements

We are grateful to M. Mascaró, P. Del Giudice, M. Mattia, D.J. Amit, and G. Mongillo for many useful discussions. We would like to thank G. La Camera and W. Senn for a careful reading of the manuscript and W. Gerstner for a recent discussion which helped to improve the revised version of the paper. J.M. Brader is currently simulating a network of integrate-and-fire neurons connected by the dynamic synapses studied here and helped to fix several mistakes in the reported parameters. This work was partially supported by the EU grant IST-2001-38099 ALAVLSI.

References

- [1] S. Fusi, M. Annunziato, D. Badoni, A. Salamon, D.J.Amit (2000) Spike-driven synaptic plasticity: theory, simulation, VLSI implementation, *Neural Comp.*, **12**, 2227–2258
- [2] W. Senn, H. Markram and M. Tsodyks (2001) An algorithm for modifying neurotransmitter release probability based on pre- and post-synaptic spike timing, *Neural Comp.* **13**, 35–68
- [3] L. Abbott and S. Nelson, Synaptic plasticity: taming the beast, *Nature Neuroscience* **3** Supp. 1178 - 1183
- [4] P. Del Giudice and M. Mattia, (2001) Long and short-term synaptic plasticity and the formation of working memory: A case study, *Neurocomputing*, 38 - 40, 1175-1180
- [5] D.J. Amit and G. Mongillo (2002) Spike-driven Synaptic Dynamics Generating Working Memory States, submitted to *Neural Computation*
- [6] Amit, D. J., and Fusi, S. (1994). Dynamic learning in neural networks with material synapses *Neural Comp.*, **6**, 957.
- [7] Bloch, H. (1962), The perceptron: a model for brain functioning, *Reviews of modern physics*, **34**, 123-135
- [8] P. Dayan, L.F. Abbott (2001), *Theoretical neuroscience*, MIT Press
- [9] P. Del Giudice, S. Fusi and M. Mattia, Modeling the formation (2003) of working memory with networks of integrate-and-fire neurons connected by plastic synapses, *J. Physiology Paris*, in press
- [10] W. Gerstner and W. Kistler (2002) *Spiking neuron models: single Neurons, populations, plasticity*, Cambridge University Press
- [11] Diederich, S., and Oppen, M. (1987) Learning correlated patterns in spin-glass networks by local learning rules. *Phys. Rev. Lett.*, **58**, 949–952
- [12] Y. Amit and M. Mascaró (2001) Attractor Networks for Shape Recognition, *Neural Comp.* **13** 1415–1442
- [13] Petersen C. C. H., Malenka, R. C., Nicoll, R. A. and Hopfield, J. J. (1998) All-or-none potentiation at CA3-CA1 synapses, *Proc. Natl. Acad. Sci.* **95**, 4732
- [14] P.J. Sjöström, G.G. Turrigiano and S.B. Nelson, (2001) Rate, Timing and Cooperativity Jointly Determine Cortical Synaptic Plasticity, *Neuron* **32**, 1149-1164
- [15] Markram H., Lubke J., Frotscher M., Sakmann B. (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs *Science*, **275**, 213-215
- [16] M. Migliore and F. Alicata (1995) A model for long term potentiation and depression, *J. of Comp. Neurosci.*, **2**, 335-343
- [17] Fusi, S., and Mattia, M. (1999) Collective behavior of networks with linear (VLSI) integrate-and-fire neurons *Neural Comp.*, **11**, 633
- [18] S. Fusi, Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates, *Biological Cybernetics*, **87**, 459-470, (2002)
- [19] H. C. Tuckwell, 1988. *Introduction to Theoretical Neurobiology* Vol. 2 (Cambridge: Cambridge University Press)
- [20] A. Rauch, G. La Camera G., H-R. Lüscher, W. Senn, S. Fusi, Neocortical pyramidal cells respond to *in vivo*-like input currents like integrate-and-fire neurons, submitted
- [21] H. D. I. Abarbanel, R. Huerta, and M. I. Rabinovich, Dynamical model of long-term synaptic plasticity, *PNAS* **99**, 10132-10137 (2002)
- [22] Sejnowski T.J. (1977) Storing covariance with nonlinearly interacting neurons *J. Math. Biol.* **4** 303
- [23] H. Z. Shouval, M. F. Bear, and L. N. Cooper, A unified model of NMDA receptor-dependent bidirectional synaptic plasticity, *PNAS* **99**, 10831-10836 (2002)
- [24] J.M. Brader, M. Mascaró, W. Senn, S. Fusi, Learning to classify complex patterns of mean rates with networks of integrate-and-fire neurons connected by bistable plastic synapses, in preparation
- [25] Brunel N., Carusi F., Fusi S. (1998) Slow stochastic Hebbian learning of classes of stimuli in a recurrent neural network, *Network*, **9**, 123-152

- [26] Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982) Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, **2**, 32
- [27] R. Kempter and W. Gerstner and J.L. van Hemmen (2001), Intrinsic stabilization of output firing rates by spike-based Hebbian learning, *Neural Computation*, **13**, 2709-2741
- [28] Rosenblatt, F. The perceptron: a probabilistic model for information storage and organization in the brain, *Psychological Review* **65**, 386-408 (1958)
- [29] M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano (2000) Stable Hebbian Learning from Spike Timing-Dependent Plasticity, *J. Neurosci.* **20**, 8812-8821
- [30] Rubin J, Lee D, and Sompolinsky H (2001) The Equilibrium Properties of Temporally Asymmetric Hebbian Plasticity, *Physical Review Letters* **86** 364.
- [31] W. Senn and S. Fusi (2003), Slow stochastic online learning with global inhibition solves the binary perceptron problem, submitted
- [32] W. Senn, Beyond spike-timing: the role of nonlinear plasticity and unreliable synapses, Hebb in Perspective, special issue of *Biological Cybernetics*, 87:344-355