Optimality Model of Unsupervised Spike-Timing Dependent Plasticity: Synaptic Memory and Weight Distribution

T. Toyoizumi [1,2], J.-P. Pfister [1] , K. Aihara [3,4] and W. Gerstner [1]

(1) School of Computer and Communication Sciences and Brain-Mind Institute,
Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland
(2) Department of Complexity Science and Engineering, Graduate School of
Frontier Sciences, University of Tokyo, Tokyo 153-8505, Japan
(3) Department of Information and Systems, Institute of Industrial Science,
University of Tokyo, 153-8505 Tokyo, Japan
(4) Aihara Complexity Modelling Project, ERATO, JST
3-23-5 Uehara, Shibuya-ku, Tokyo 151-0064, Japan

Summary

We studied the hypothesis that synaptic dynamics is controlled by three basic principles: (A) Synapses adapt their weights so that neurons can effectively transmit information; (B) homeostatic processes stabilize the mean firing rate of the postsynaptic neuron; and (C) weak synapses adapt more slowly than strong ones, while maintenance of strong synapses is costly. Our results show that a synaptic update rule derived from these principles shares features with spike-timing dependent plasticity, is sensitive to correlations in the input, and is useful for synaptic memory. Moreover, input selectivity (sharply tuned receptive fields) of postsynaptic neurons develops only if stimuli with strong features are presented. Sharply tuned neurons can co-exist with unselective ones and the distribution of synaptic weights can be unimodal or bimodal. The formulation of synaptic dynamics through an optimality criterion provides a simple graphical argument for the stability of synapses, necessary for synaptic memory.

1 Introduction

Synaptic changes are thought to be involved in learning, memory, and cortical plasticity, but the exact relation between microscopic synaptic properties and macroscopic functional consequences remains highly controversial. In experimental preparations, synaptic changes can be induced by specific stimulation conditions defined through pre- and postsynaptic firing rates (Bliss and Lomo 1973; Dudek and Bear 1992), postsynaptic membrane potential (Kelso et al. 1986), calcium entry (Malenka et al. 1988; Lisman 1989), or spike timing (Markram et al. 1997; Bi and Poo 2001). In the theoretical community, conditions for synaptic changes are formulated as 'synaptic update rules' or 'learning rules' (von der Malsburg 1973; Bienenstock et al. 1982; Miller et al. 1989) (for reviews see (Gerstner and Kistler 2002; Dayan and Abbott 2001; Cooper et al. 2004) but the exact features that make a synaptic update rule a suitable candidate for cortical plasticity and memory are unclear.

From a theoretical point of view, synaptic learning rule should be *(i)* sensitive to correlations between pre- and postsynaptic neurons (Hebb 1949) in order to respond to correlations in the input (Oja 1982); they should *(ii)* allow neurons to develop *input selectivity* (e.g., receptive fields) (Bienenstock et al. 1982; Miller et al. 1989), in the presence of strong input features, but *(iii)* distribution of synaptic strength should remain unimodal otherwise (Gütig et al. 2003). Furthermore *(iv)* synaptic memories should show a high degree of stability (Fusi et al. 2005) and nevertheless remain plastic (Grossberg 1987). Moreover, experiments suggest that plasticity rules are *(v)* sensitive to the *presynaptic firing rate* (Dudek and Bear 1992), but *(vi)* depend also on the exact timing of the pre- and postsynaptic spikes (Markram et al. 1997; Bi and Poo 2001).

Many other experimental features could be added to this list, e.g., the role of intracellular calcium, of NMDA receptors, etc., but we will not do so; see (Bliss and Collingridge 1993; Malenka and Nicoll 1993) for reviews.

The items in the above list are not necessarily exclusive, and the relative importance of a given aspect may vary from one subsystem to the next; for example, synaptic memory maintenance might be more important for a long-term memory system than for primary sensory cortices. Nevertheless, all of the above aspects seem to be important features of synaptic plasticity. However, the development of theoretical learning rules that exhibit all of the above properties has posed problems in the past. For example, traditional learning rules that have been proposed as an explanation of receptive field development (Bienenstock et al. 1982; Miller et al. 1989), exhibit a spontaneous separation of synaptic weights into two groups, even if the input shows no or only weak correlations. This is difficult to reconcile with experimental results in visual cortex of young rats where a unimodal distribution was found (Sjöström et al. 2001). Moreover model neurons that specialize early in development on one subset of features cannot readily re-adapt later on. Other learning rules, how-

ever, that exhibit a unimodal distribution of synaptic weights (Gütig et al. 2003) do not lead to a longterm stability of synaptic changes.

In this paper we want to show that all of the above features *(i)* - *(vi)* emerge naturally in a theoretical model where we require only a limited number of objectives that will be formulated as postulates. In particular, we study how the conflicting demands on synaptic memory maintenance, plasticity, and distribution of synaptic synapses could be satisfied by our model. Even though the postulates are rather general and could be adapted to arbitrary neural systems, we had in mind excitatory synapses in neocortex or hippocampus and exclude inhibitory synapses and synapses in specialized systems such as the calix of held in the auditory pathway. Our arguments are based on three postulates:

(**A**) Synapses adapt their weights so as to allow neurons to efficiently transmit information. More precisely, we impose a theoretical postulate that the mutual information I between presynaptic spike trains and postsynaptic firing be optimized. Such a postulate stands in the tradition of earlier theoretical work (Linsker 1989; Bell and Sejnowski 1995), but is formulated here on the level of spike trains rather than rates.

(**B**) Homeostatic processes act on synapses to ensure that the long-term average of the neuronal firing rate becomes close to a target rate that is characteristic for each neuron. Synaptic rescaling and related mechanism could be a biophysical implementation of homeostatis (Turrigiano and Nelson 2004). The theoretical reason for such a postulate is that sustained high firing rates are costly from an energetic point of view (Laughlin et al. 1998; Levy and Baxter 2002).

(**C**). C1: Maintenance of strong synapses is costly in terms of biophysical machinery, in particular in view of continued protein synthesis (Fonseca et al. 2004). C2: Synaptic plasticity is slowed down for very weak synapses in order to avoid a (unplausible) transition from excitatory to inhibitory synapses.

Optimality approaches have a long tradition in the theoretical neurosciences and have been utilized in two different ways. Firstly, optimality approaches allow to derive strict theoretical bounds against which performance of real neural systems can be compared (Barlow 1956; Laughlin 1981; Britten et al. 1992; de Ruyter van Steveninck and Bialek 1995). Secondly, they have been used as a conceptual framework since they allow to connect functional objectives (e.g., 'be reliable!') and constraints (e.g., 'don't use too much energy!') with electrophysiological properties of single neurons and synapses or neuronal populations (Barlow 1961; Linsker 1989; Atick and Redlich 1990; Levy and Baxter 2002; Seung 2003). Our study, i.e., derivation of synaptic update rules from an optimality viewpoint, follows this second, conceptual, approach.

2   The Model

Neuron Model

We simulated a single stochastic point neuron model with $N = 100$ input synapses. Presynaptic spikes at synapse j are denoted by their arrival time t_j^f and evoke EPSPs with time course $\exp[-(t - t_j^f)/\tau_m]$ for $t \geq t_j^f$ where $\tau_m = 20$ms is the membrane time constant. Recent experiments have shown that action potentials propagating back into the dendrite can partially suppress EPSPs measured at the soma (Froemke et al. 2005). Since our model neuron has no spatial structure, we included EPSP suppression by a phenomenological amplitude factor $a(t_j^f - \hat{t})$ that depends on the time difference between presynaptic spike arrival and the spike trigger time $\hat{t}$ of the last (somatic) action potential of the postsynaptic neuron.

In the absence of EPSP suppression, the amplitude of a single EPSP at synapse j is characterized by its weight w_j and its duration by the membrane time constant τ_m. Summation of the EPSPs caused by presynaptic spike arrival at all 100 explicitly modeled synapses gives the total postsynaptic potential

$$u(t) = u_r + \sum_{j=1}^{N} w_j \sum_{t_j^f < t} \exp\left(-\frac{t - t_j^f}{\tau_m}\right) a(t_j^f - \hat{t}) \tag{1}$$

where $u_r = -70$mV is the resting potential and the sum runs over all spike arrival times t_j^f in the recent past, $\hat{t} < t_j^f \leq t$. The EPSP suppression factor takes a value of zero if $t_j^f < \hat{t}$ and is modeled for $t_j^f \geq \hat{t}$ as exponential recovery $a(t_j^f - \hat{t}) = 1 - \exp[-(t_j^f - \hat{t})/\tau_a]$ with time constant $\tau_a = 50$ms (Fig. 1A) unless stated otherwise. The parameters w_j for $1 \leq j \leq N$ denote the synaptic weight of the $N = 100$ synapses and are updated using a learning rule discussed below.

In order to account for unspecific background input that was not modeled explicitly, spikes were generated probabilistically with density

$$\rho(t) = \rho_r + [u(t) - u_r] \cdot g \tag{2}$$

where $\rho_r = 1$Hz is the spontaneous firing rate (in the absence of spike input at the 100 explicitly modeled synapses) and $g = 12.5$Hz/mV is a gain factor. Thus, the instantaneous spike density increases linearly with the total postsynaptic potential $u(t)$. Note, however, that due to EPSP suppression the total postsynaptic potential increases sublinearly with the the number of input spikes and so does the mean firing rate of the postsynaptic neuron (Fig. 1B).

The neuron model is simulated in discrete time with time steps of $\Delta t = 1$ms on a standard personal computer using custom made software written in Matlab.

Objective Function

Postulates **A** and **B** have been used previously (Toyoizumi et al. 2005) and lead to an optimality criterion $\mathcal{L}' = I - \gamma D$ where I is the mutual information between presynaptic input and postsynaptic output and D a measure of the distance of the mean firing rate of the neuron from its target rate. The parameter γ scales the importance of the information term I (postulate **A**) compared
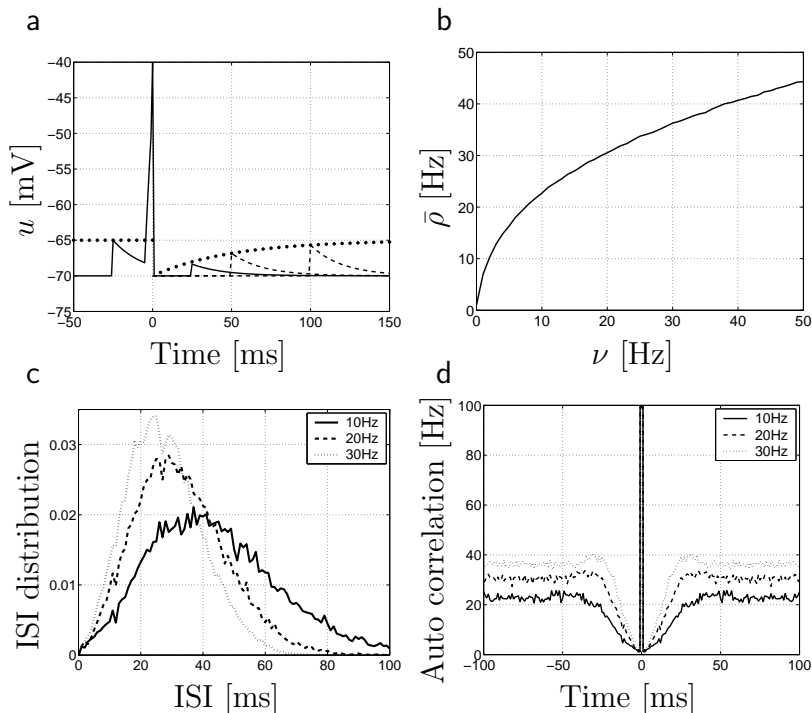
Figure 1: Properties of the stochastically spiking neuron model. (a) Neuron
model and EPSP suppression. EPSPs arriving just after a postsynaptic spike
at $t = 0$ are attenuated by a factor $(1 - e^{-t/\tau_a})$ and recover exponentially to
their maximal amplitude w_j. (b) The output firing rate $\bar{\rho}$ of a neuron that
receives stochastic spike input at rate ν at all 100 synapses. Each presynaptic
spike evokes an EPSPs with maximal amplitude $w_j = 0.4$mV. (c) Interspike
interval (ISI) distribution with input frequency $\nu = 10$Hz (*solid line*), 20Hz
(*dashed line*), and 30Hz (*dotted line*) at all 100 synapses. (d) Autocorrelation
function of postsynaptic action potentials at an input frequency of 10Hz (*solid
line*), 20Hz (*dashed line*), and 30Hz (*dotted line*).

to the homeostatic term D (postulate **B**). It was shown that optimization of
$\mathcal{L}'$ by gradient ascent yields a synaptic update rule which shows sensitivity
to correlations (see point *(i)* above), input selectivity (see point *(ii)* above),
and depends on presynaptic firing rates (see point *(v)* above) (Toyoizumi et al.
2005). However, while the learning rule in (Toyoizumi et al. 2005) showed
some dependence on spike timing, it did not (without additional assumptions)
have the typical features of STDP as measured in vitro (point *(vi)* above); and
exhibited, like earlier models (Bienenstock et al. 1982), *spontaneous* synaptic
specialization, even for very weak input features, which is in contrast to point
(iii) above.

In this earlier theoretical study, synaptic potentiation was artificially stopped
at some upper bound w^{max} (and synaptic depression was stopped at weight
$w = 0$), so as to ensure that weights w stayed in a regime $0 \leq w \leq w^{\mathrm{max}}$. In

the present paper we take the more realistic assumption that strong weights are more likely to show depression than weaker ones but do not impose a hard upper bound. Similarly, we require that adaptation speed is slowed down for very weak synapses, but do not impose a hard bound at zero weight. We will show that with these assumptions the resulting synaptic update rule shows properties of STDP (see point *(vi)* above), is suitable for memory retention (see point *(iv)* above) and leads to synaptic specialization when driven by strong input (see point *(iii)* above), while keeping the properties *(i), (ii)* and *(v)*, that were found in (Toyoizumi et al. 2005).

To avoid hard upper bounds for the synapses, we use postulate **C1** and add a term Ψ to the optimality criterion $\mathcal{L}'$ that is proportional to w^2 (i.e., the square of the synaptic weight) and proportional to the presynaptic firing rate. This term comes with a negative sign, since a *cost* is associated to big weights. Hence from our optimality viewpoint, synapses change so as to maximize a quantity

$$\mathcal{L} = I - \gamma D - \lambda \Psi \tag{3}$$

where I is the information to be maximized, D a measure of the firing rate mismatch to be minimized, and Ψ the cost induced by strong synapses to be minimized. The factors γ and λ control the relative importance of the three different terms. In other words, synapses adjust their weights so as to be able to transmit information while keeping the mean firing rate and synaptic weights at low values. Thus our three postulates **A**, **B**, and **C** give rise to one unified optimality criterion $\mathcal{L}$. We hypothesize that a significant part of findings regarding synaptic potentiation and depression can be conceptually understood as the synapse's attempt to optimize the criterion $\mathcal{L}$.

The learning rule used for the update of the synaptic weights w_j is derived from the objective function (3) i.e., $\mathcal{L} = I - \gamma D - \lambda \Psi$ which contains three terms.

The first term is the mutual information between the ensemble of 100 input spike trains (spanning the interval of a single trial from 0 to T; the ensemble of all presynaptic trains is formally denoted by $X(T) = \{x_j(t) = \sum_{t_j^f} \delta(t - t_j^f) | j = 1, \ldots, 100, \ 0 \le t < T\}$) and the output spike train of the postsynaptic neuron over the same interval (denoted by $Y(T) = \{y(t) = \sum_{t_{post}^f} \delta(t - t_{post}^f) | 0 \le t < T\}$, where t_{post}^f represent output spike-timing), i.e.,

$$I = \left\langle \log \frac{P(Y|X)}{P(Y)} \right\rangle_{Y,X}, \tag{4}$$

where angular brackets $\langle \cdot \rangle_{Y,X}$ denote averaging over all combinations of input and output spike trains[1]. Here $P(Y|X)$ is the conditional probability density

---

[1]From now on, when X or Y are without argument, we take implicitly $X \equiv X(T)$ and $Y \equiv Y(T)$, i.e., the spike trains over the full interval T

of our stochastic neuron model to generate a specific spike train Y with (one or several) spike times $\{t^f_{post}\}$ during a trial of duration T given 100 known input spike trains X. This conditional probability density is given as a product of the instantaneous probabilities $\rho(t^f_{post})$ of firing at the postsynaptic spike times $\{t^f_{post}\}$ and the probability of not firing elsewhere, i.e.,

$$P(Y|X) = \left[\prod_{t^f_{post}} \rho(t^f_{post}) \right] \exp\left[-\int_0^T \rho(t)dt \right]. \tag{5}$$

Similarly, $P(Y)$ is the probability to generate the very same output spike train Y *not* knowing the input. Here 'not knowing the input' implies that we have to average over all possible inputs so as to get the expected instantaneous firing density $\bar{\rho}(t)$ at time t. However, because of the EPSP suppression factor, the expected firing density will also depend on the last output spike before t. We therefore define

$$\bar{\rho}(t) = \langle \rho(t) \rangle_{X(t)|Y(t)}, \tag{6}$$

i.e., we average over the inputs but keep the knowledge of the previous output spikes $t^f_{post} < t$. $P(Y)$ is then given by a formula analogous to (5), but with ρ replaced by $\bar{\rho}$. Hence, given our neuron model, both $P(Y|X)$ and $P(Y)$ in Eq. (4) are well-defined. The information term I of Eq. (4) is the formal instantiation of postulate **A**.

The second term is the homeostatic term

$$D = \left\langle \log \frac{P(Y)}{\tilde{P}(Y)} \right\rangle_Y, \tag{7}$$

which compares the actual distribution of output spike trains $P(Y)$ with that of an ideal distribution $\tilde{P}(Y)$ generated by the same neuron firing at target rate of $\tilde{\rho} = 5$Hz, i.e., formula (5) with ρ replaced by $\tilde{\rho}$. Mathematically speaking, D is the Kullback-Leibler distance between two distributions (Cover and Thomas 1991), but in practice we may think of D simply as a measure of the difference between actual and target firing rates (Toyoizumi et al. 2005). The term D is our mathematical formulation of postulate **B**.

The third term is the cost associated with strong synapses. We assume that the cost increases quadratically with the synaptic weights but that only those synapses that have been activated in the past contribute to the cost. Hence the mathematical formulation of postulate **C1** yields a cost

$$\Psi = \frac{1}{2} \sum_j w_j^2 \langle n_j \rangle_X \tag{8}$$

where n_j is the number of presynaptic spikes that have arrived at synapse j during the duration T of the interval under consideration. Cost terms that

8

are quadratic in the synaptic weights are common in the theoretical literature (Miller and MacKay 1994), but the specific dependence upon presynaptic spiking induced by the factor n_j in Eq. (8) is not. The dependence of Ψ upon presynaptic spike arrival means that, in our model, only activated synapses contribute to the cost. The specific formulation of Ψ is mainly due to theoretical reasons to be discussed below. The intuition is that activation of a synapse in *absence* of any postsynaptic activity can weaken the synapse if the factor λ is sufficiently positive (see also Fig. 3D). The restriction of the cost to previously *activated* synapses is reminiscent of synaptic tagging (Frey and Morris 1997; Fonseca et al. 2004) even though any relation must be seen as purely hypothetical.

The three terms are given a relative importance by choosing $\gamma = 0.1$ (for a discussion of this parameter see (Toyoizumi et al. 2005)) and $\lambda = 0.026$ so as to achieve a baseline of zero in the STDP function (see Appendix A).

Synaptic update rule

We optimize the synaptic weights by gradient ascent

$$\Delta w_j = \alpha(w_j) \frac{\partial \mathcal{L}}{\partial w_j} \tag{9}$$

with a weight-dependent update rate $\alpha(w_j)$. According to postulate C2, plasticity is reduced for very small weights. For the sake of simplicity we chose $\alpha(w_j) = 4 \cdot 10^{-2} \frac{w_j^4}{w_j^4 + w_s^4}$, where $w_s = 0.2\text{mV}$. i.e., learning slows down for weak synapses with EPSP amplitudes around or less than 0.2mV. Note that updates according to Eq. (9) are always uphill; however, because of the w_j dependence of α, the ascent is not necessarily along the steepest gradient.

Using the same mathematical arguments as in (Toyoizumi et al. 2005), we can transform the optimization by gradient ascent into a synaptic update rule. First, differentiating each term, we find

$$\frac{\partial I}{\partial w_j} = \left\langle \frac{1}{P(Y|X)} \frac{\partial P(Y|X)}{\partial w_j} \log \frac{P(Y|X)}{P(Y)} \right\rangle_{Y,X}, \tag{10}$$

$$\frac{\partial D}{\partial w_j} = \left\langle \frac{1}{P(Y|X)} \frac{\partial P(Y|X)}{\partial w_j} \log \frac{P(Y)}{\tilde{P}(Y)} \right\rangle_{Y,X}, \tag{11}$$

$$\frac{\partial \Psi}{\partial w_j} = w_j \left\langle n_j \right\rangle_X. \tag{12}$$

We will rewrite the terms appearing in Eqs. (10) and (11) by introducing the auxiliary variables

$$c_j(t) = \frac{d\rho/du|_{u=u(t)}}{\rho(t)} [y(t) - \rho(t)] \int_0^\infty ds' \epsilon(s') x_j(t - s') \tag{13}$$

and

$$B^{post}(t) = \left[y(t) \log \frac{\rho(t)}{\bar{\rho}(t)} - (\rho(t) - \bar{\rho}(t)) \right] - \gamma \left[y(t) \log \frac{\bar{\rho}(t)}{\tilde{\rho}} - (\bar{\rho}(t) - \tilde{\rho}) \right], \tag{14}$$

Using the definitions in Eqs. (13) and (14), we find the derivative of the conditional probability density that appears in Eqs. (10) and (11)

$$\frac{\partial P(Y|X)}{\partial w_j} = P(Y|X) \int_0^T c_j(t')dt' \tag{15}$$

and

$$\log \frac{P(Y|X)}{P(Y)} - \gamma \log \frac{P(Y)}{\tilde{P}(Y)} = \int_0^T B^{post}(t)dt. \tag{16}$$

As a first interpretation we may say that c_j represents the causal correlation between input and output spikes (corrected for the expected correlation); and B^{post} is a function of postsynaptic quantities, namely, the output spikes y, current firing rate ρ via the membrane potential u, average firing rate $\bar{\rho}$ and the target firing rate $\tilde{\rho}$. More precisely, B^{post} compares the actual output with the expected output and, modulated by a factor γ, the expected output with the target.

Hence, with the results from Eqs. (10) — (16) the derivative of the objective function is written in terms of averaged quantities $\langle \cdot \rangle_{Y,X}$ as

$$\frac{\partial \mathcal{L}}{\partial w_j} = \int_0^T dt \left\langle \left[\int_0^T c_j(t')dt' \right] B^{post}(t) - \lambda w_j x_j(t) \right\rangle_{Y,X}. \tag{17}$$

An important property of c_j is that its average $\langle c_j \rangle_{Y|X}$ vanishes. On the other hand, the correlations between $c_j(t')$ and $B^{post}(t)$ are limited by the time scale τ_{AC} of the auto-correlation function of the output spike train. Hence we can limit the integration to the relevant time scales without loss of generality and introduce an exponential cut-off factor with time constant $\tau_C > \tau_{AC}$

$$C_j(t) = \lim_{\varepsilon \to +0} \int_0^{t+\varepsilon} c_j(t')e^{-(t-t')/\tau_C}dt'. \tag{18}$$

With this factor C_j, we find a batch learning rule (i.e., with expectations over the input and output statistics on the right-hand side) of the form

$$\frac{\partial \mathcal{L}}{\partial w_j} \approx \int_0^T dt \left\langle C_j(t)B^{post}(t) - \lambda w_j x_j(t) \right\rangle_{Y,X}. \tag{19}$$

Finally, for slow learning rate α and stationary input statistics, the system becomes self-averaging (i.e. expectations can be dropped due to automatic temporal averaging (Gerstner and Kistler 2002)) so that we arrive at the online gradient learning rule

$$\frac{dw_j}{dt} = \alpha(w_j) \left[C_j(t)B^{post}(t) - \lambda w_j x_j(t) \right]. \tag{20}$$

See Fig. 2 for an illustration of the dynamics of C_j and B^{post}. The last term has the form of an 'weight decay' term common in artificial neural networks

(Hertz et al. 1991) and arises from the derivative of the weight-dependent cost term Ψ. The parameter λ is set such that $dw_j/dt = 0$ for large enough $|t^{pre} - t^{post}|$ in the STDP in vitro paradigm. A few steps of calculation (see Appendix A) yield $\lambda = 0.026$. In our simulations, we take $\tau_C = 100$ms for the cut-off factor in Eq. (18).

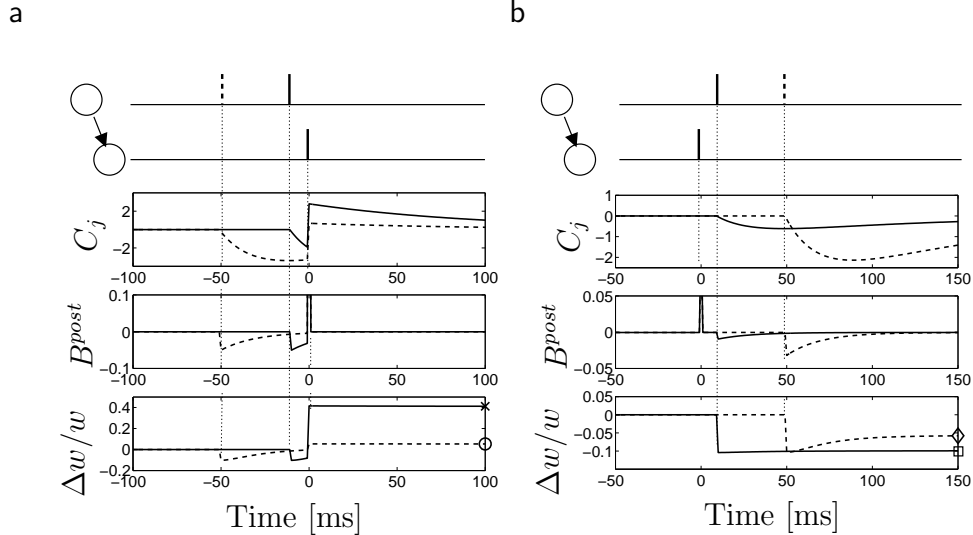Figure 2: Illustration of the dynamics of B^{post} (*Top*), C_j (*Middle*) and $\Delta w/w = (w - w_{init})/w_{init}$ (*Bottom*). We always start from an initial weight $w_{init} = 4$ mV and induce postsynaptic firing at time $t^{post} = 0$. (a) Pre-before-post timing with $t^{pre} = -10$ ms (*solid line*) induces a large potentiation whereas $t^{pre} = -50$ ms (*dashed line*) induces almost no potentiation. (b) Due to the EPSP suppression factor, a post-before-pre timing with $t^{pre} = 10$ ms (*solid line*) induces a large depression whereas $t^{pre} = 50$ ms (*dashed line*) induces a smaller depression. The marks (circle, cross, square and diamond) correspond to the weight change due to a single pair of pre- and postsynaptic spike after weights have converged to their new values. Note that in Fig. 3, the corresponding marks indicate the weight change after 60 pairs of pre- and postsynaptic spikes.

For a better understanding of the learning dynamics defined in Eq. (20), let us look more closely at Fig. 2a. The time course $\Delta w/w$ of the potentiation has three components: first, a negative jump at the moment of the presynaptic spike induced by the weight decay term; second, a slow increase in the interval between pre- and postsynaptic spike times induced by $C_j(t)B^{post}(t) > 0$; third, a positive jump immediately after the postsynaptic spike induced by the singularity in B^{post} combined with a positive C_j.

As it is practically difficult to calculate $\bar{\rho}(t) = \langle \rho(t) \rangle_{X(t)|Y(t)}$, we estimate $\bar{\rho}$ by the running average of the output spikes, i.e.,

$$\tau_{\bar{\rho}} \frac{d\bar{\rho}^{est}}{dt} = -\bar{\rho}^{est}(t) + y(t) \tag{21}$$

11

with $\tau_{\bar{\rho}} = 1$min. This approximation is valid if the characteristics of the stimulus and output spike trains are stationary and uncorrelated. To simulate in vitro STDP experiments, the initial value of $\bar{\rho}^{est}$ is set equal to the injected pulse frequency. Other, more accurate estimates of $\bar{\rho}^{est}$ are possible, but lead to no qualitative change of results (data not shown). In the simulations of Fig. 4-6, the initial synaptic strength is set to be $w_{init} = 0.4 \pm 0.04$mV.

Stimulation paradigm

Simulated STDP in vitro paradigm. For the simulations of Fig. 3, spike timings $t^{pre} = t_j^f$ at synapse j and postsynaptic spike times t^{post} are imposed with a given relative timing $t^{pre} - t^{post}$. For the calculation of the total STDP effect according to a typical in vitro stimulation paradigm, the pairing of pre- and postsynaptic spikes is repeated until a total of 60 spike pairs have been accumulated. Spike pairs are triggered at a frequency of 1Hz except for Fig. 3C where the stimulation frequency was varied.

Simulated stochastic spike arrival. In most simulations presynaptic spike arrival was modeled as Poisson spike input either at a fixed rate (homogeneous Poisson process) or modulated rate (inhomogeneous Poisson process).

For example for the simulations in Fig. 6, with a *Gaussian profile*, spike arrival at synapse j is generated by an inhomogeneou Poisson process with the following characteristics. During a segment of 200 ms, the rate is fixed at $\nu_j = (\nu^{\max} - \nu_0)\exp[-0.01 * d(j-k)^2] + \nu_0$ where $\nu_0 = 1$ Hz is the baseline firing rate and $d(j-k)$ is the difference between index j and k. The value of k denotes the location of the maximum. The value of k was reset every 200ms to a value chosen stochastically between 1 and 100. [As indicated in the main text, presynaptic neurons in Fig. 6 were considered to have a ring topology which has been implemented by evaluating the difference $d(j-k)$ as $d(j-k) = \min\{|j-k|, 100 - |j-k|\}$.]

However, in the simulations for Fig. 5, input spike trains were not independent Poisson, but we included *spike-spike correlations*. A correlation index of $c = 0.2$ implies that between a given pair of synapses 20 percent of spikes have identical timing. More generally, for a given value of c within a group of synaptic inputs, $100\,c$ percent of spike arrival times are identical at an arbitrary pair of synapse within the group.

3   Results

The mathematical formulation of postulates **A**, **B** and **C1** led to an optimality criterion $\mathcal{L}$ which was optimized by changing synaptic weights in uphill direction. In order to include postulate **C2**, the adaptation speed was made to depend on the current value of the synaptic weight so that plasticity was significantly slowed down for synapses with excitatory postsynaptic potentials (EPSPs) of amplitude less than 0.2mV (see Section 2 for details).

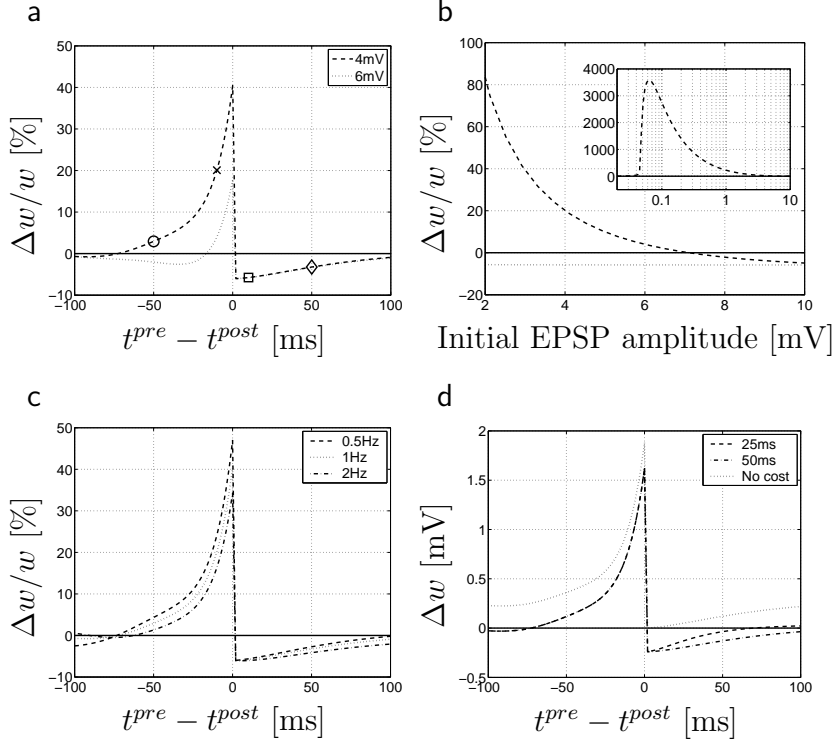As in a previous study based on postulates **A** and **B** (Toyoizumi et al.

Figure 3: The synaptic update rule of the model shares features with STDP. (a) STDP function (percentage change of EPSP amplitude as a function of $t^{pre} - t^{post}$) determined using 60 pre-and-post spike-pairs injected at 1Hz. The initial EPSP amplitudes are 4mV (*dashed line*) and 6mV (*dotted*). Marks (circle, cross, square and diamond) correspond respectively to $t^{pre} = -50$ms, $t^{pre} = -10$ms, $t^{pre} = 10$ms and $t^{pre} = 50$ms also depicted on Fig. 2. (b) The percentage change in EPSP amplitude after 60 pre-and-post spike-pairs injected at 1Hz for pre-before-post timing $((t^{pre} - t^{post}) = -10$ms, *solid line*) and post-before-pre timing $((t^{pre} - t^{post}) = +10$ms, *dashed line*) as a function of initial EPSP amplitude. Our model results qualitatively resemble experimental data (see Fig. 5 in Bi and Poo, 1998).(c) Frequency dependence of the STDP function: spike-pairs are presented at frequencies of 0.5Hz (*dashed line*), 1Hz (*dotted line*), and 2Hz (*dot-dashed line*). The STDP function exhibits only a weak sensitivity to the change in stimulation frequency. (d) STDP function for different choices of model parameters. The extension of the synaptic depression zone for 'pre-after-post' timing ($t^{pre} - t^{post} > 0$) depends on the time scale τ_a of EPSP suppression (*dot-dashed line*, $\tau_a = 50$ms; *dashed line*, $\tau_a = 25$ms). The dotted line shows the STDP function in the absence of a weight-dependent cost term Ψ. The STDP function exhibits a positive offset indicating that without the cost term Ψ unpaired presynaptic spikes would lead to potentiation, i.e., a non-Hebbian form of plasticity.

13

2005), the optimization of synaptic weights can be understood as a synaptic update rule that depends on presynaptic spike arrival, postsynaptic spike firing, the postsynaptic membrane potential and the mean firing rate of the postsynaptic neuron. In addition, the synaptic update rule in the present study included a term that decreases the synaptic weight upon presynaptic spike arrival by a small amount proportional to the EPSP amplitude (see Section 2 for details). This term can be traced back to the additional weight-dependent cost term Ψ in Eq. (3) that accounts for postulate **C1**.

In order to study the consequences of the synaptic update rule derived from postulates **A**, **B**, and **C**, we used computer simulations of a model neuron that received presynaptic spike trains at 100 synapses. Each presynaptic spike evoked an EPSP with exponential time course (time constant $\tau_m = 20$ms). In order to account for dendritic interaction between somatic action potentials and postsynaptic potentials, the amplitude of EPSPs was suppressed immediately after postsynaptic spike firing (Froemke et al. 2005) and recovered with time constant $\tau_a = 50$ms (Fig. 1A). As a measure of the weight of a synapse j we used the EPSP amplitude w_j at this synapse in the absence of EPSP suppression. With all synaptic parameters w_j set to an fixed value, the model neuron fired stochastically with a mean firing rate $\bar{\rho}$ that increases with the presynaptic spike arrival rate (Fig. 1B), has a broad interspike interval distribution (Fig. 1C) and an autocorrelation function with a trough of 10-50 ms that is due to reduced excitability immediately after a spike because of EPSP suppression (Fig. 1D).

The learning rule exhibits STDP.

In a first set of plasticity experiments, we explored the behavior of the model system under a simulated in vitro paradigm as used in typical STDP experiments (Bi and Poo 1998). In order to study the influence of the pre- and postsynaptic activity on the changes of weights as predicted by our on-line learing rule in Eq. 20, we plotted in Fig. 2 the postsynaptic factor B^{post} and the correlation term C_j that both appear on the right-hand side of Eq. 20 together with the induced weight change $\Delta w/w$ as a function of time. Indeed, the learning rule predicts positive weight changes when the presynaptic spike occurs 10 ms before the postsynaptic one and negative weight changes under reversed timing.

For a comparison with experimental results, we used sixty pairs of pre- and postsynaptic spikes applied at a frequency of 1Hz and recorded the total change Δw in EPSP amplitude. The experiment is repeated with different spike timings and the result is plotted as a function of spike timing difference $t^{pre} - t^{post}$. As in experiments (Markram et al. 1997; Zhang et al. 1998; Bi and Poo 1998; Bi and Poo 2001; Sjöström et al. 2001), we find that synapses are potentiated if presynaptic spikes occur about 10 ms before a postsynaptic action potential, but are depressed if the timing is reversed. Compared to synapses with amplitudes in the range of 1 or 2 mV, synapses which are exceptionally strong show a reduced effect of potentiation for pre-before-post

timing, or even depression (Fig. 3A and B), in agreement with experiments on cultured hippocampal neurons (Bi and Poo 1998). The shape of the STDP function depends only weakly on the stimulation frequency (Fig. 3C), even though a significant reduction of the potentiation amplitude with increasing frequency can be observed.

In a recent experimental study (Froemke et al. 2005) a strong correlation between the time scale of EPSP suppression (which was found to depend on dendritic location) and the duration of the LTD part in the STDP function were observed. Since our model neuron had no spatial structure, we artificially changed the time constant of EPSP suppression in the model equations. We found that indeed only the LTD part of the STDP function was affected whereas the LTP part remained unchanged (Fig. 3D).

In order to study the influence of the weight-dependent cost term Ψ in our optimality criterion $\mathcal{L}$, we systematically changed the parameter λ in Eq. (3). For $\lambda = 0$ the weight-dependent cost term has no influence and, because the postsynaptic firing rate is close to the desired rate, synaptic plasticity in our model is mainly controlled by information maximization. In this case, synapses with a resonable EPSP amplitude of one or a few millivolt are always strengthened, even for post-before-pre timing (Fig. 3D, dashed line). This can be intuitively understood since an increase of synaptic weight is always beneficial for information transmission except if spike arrival occurs immediately *after* a postsynaptic spike. In this case, the postsynaptic neuron is insensitive so that no information can be transmitted. Nevertheless, information transmission is maximal in a situation where the presynaptic spike occurs just *before* the postsynaptic one. The weight-dependent cost term derived from postulate **C** is essential to shift the dashed line in Fig. 3D to negative values so as to induce synaptic depression in our STDP paradigm. The optimal value of $\lambda = 0.026$ that ensures that for large spike timing differences $|t^{pre} - t^{post}|$ neither potentiation nor depression occurs, has been estimated from a simple analytical argument (see appendix A).

Both unimodal and bimodal synapse distributions are stable

Under random spike arrival with a rate of 10Hz at all 100 synapses, synaptic weights show little variability with a typical EPSP amplitude in the range of 0.4 mV. This unspecific pattern of synapses stays stable even if 20 out of the 100 synapses are subject to a common rate modulation between 1 and 30Hz (Fig. 4A). However, if modulation of presynaptic firing rates becomes strong, the synapses develop rapidly a specific pattern with large values of weights at synapses with rate-modulated spike input and weak weights at those synapses that received input at fixed rates (*synaptic specialization*, see Fig. 4A and B), making the neuron highly selective to input at one group of synapses (*input selectivity*). Thus, our synaptic update rule is capable of selecting strong features in the input, but does also allow a stable unspecific pattern of synapses in case of weak input. This is in contrast to most other Hebbian learning rules, where unspecific patterns of synapses are unstable so

that synaptic weights move spontaneously towards their upper or lower bounds (Miller et al. 1989; Miller and MacKay 1994; Gerstner et al. 1996; Kempter et al. 1999; Song et al. 2000).

After induction of synaptic specialization by strong modulation of presynaptic input, we reduced the rate modulation back to the value that previously led to an unspecific pattern of synapses. We found that the strong synapses remained strong and weak synapses remained weak, i.e., synaptic specialization was stable against a change in the input (Fig. 4B). This result shows that synaptic dynamics exhibits hysteresis which is an indication of bistability: for the same input, both an unspecific pattern of synapses and synaptic specialization are stable solutions of synaptic plasticity under our learning rule. Indeed under rate-modulation between 1 and 30 Hz for 20 out of the 100 synapses, the objective function $\mathcal{L}$ shows two local maxima (Fig. 4C), a sharp maximum corresponding to synaptic specialization (mean EPSP amplitude about 0.8mV for synapses receiving rate-modulated input and less than 0.1 mV for synapses receiving constant input) and a broader, but slightly lower maximum where both groups of synapses have a mean EPSP amplitude in the range of 0.3-0.5mV; see appendix B for details of method. In additional simulations we confirmed that both the unspecific pattern of synapses and the selective pattern representing synaptic specialization remained stable over several hours of continued stimulation with rate modulated input (data not shown). Bistability of selective and unspecific synapse patterns was consistently observed for rates modulated between 1 and 10Hz or between 1 and 30Hz, but the unspecific synapse pattern was unstable if the rate was modulated between 1 and 50Hz consistent with the weight-dependence of our objective function $\mathcal{L}$ (Fig. 4D).

Retention of synaptic memories

In order to study synaptic memory retention with our learning rule, we induced synaptic specialization by stimulating 20 out of the 100 synapses by correlated spike input (spike-spike correlation index $c = 0.2$, see Section 2 for details). The remaining 80 synapses received uncorrelated Poisson spike input. The mean firing rate (10Hz) was identical at all synapses. After 60 minutes of correlated input at the group of 20 synapses the stimulus was switched to uncorrelated spike input at the same rate. We studied how well synaptic specialization was maintained as a function of time after induction (Fig. 5A and B).

Synaptic specialization was defined by a bimodality index that compared the distribution of EPSP amplitudes at synapses that received correlated input with those receiving uncorrelated input. For each of the two groups of synapses, we calculated the mean $\bar{w} = \langle w \rangle$ and the variance $\sigma^2 = \langle [w_j - \bar{w}]^2 \rangle$ i.e., $\bar{w}_A$ and σ_A^2 for the group of synapses receiving correlated input and , $\bar{w}_B$ and σ_B^2 for those receiving uncorrelated input. We then approximated the two distributions by Gaussian functions. The bimodality index depends on the overlap between the two Gaussians and is given by $b = 0.5 \left[\text{erf}\left(\frac{\bar{w}_A - \hat{s}}{\sqrt{2}\sigma_A} \right) + \text{erf}\left(\frac{\hat{s} - \bar{w}_B}{\sqrt{2}\sigma_B} \right) \right]$ where $\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$ is the error function and $\hat{s}$ is one of the two
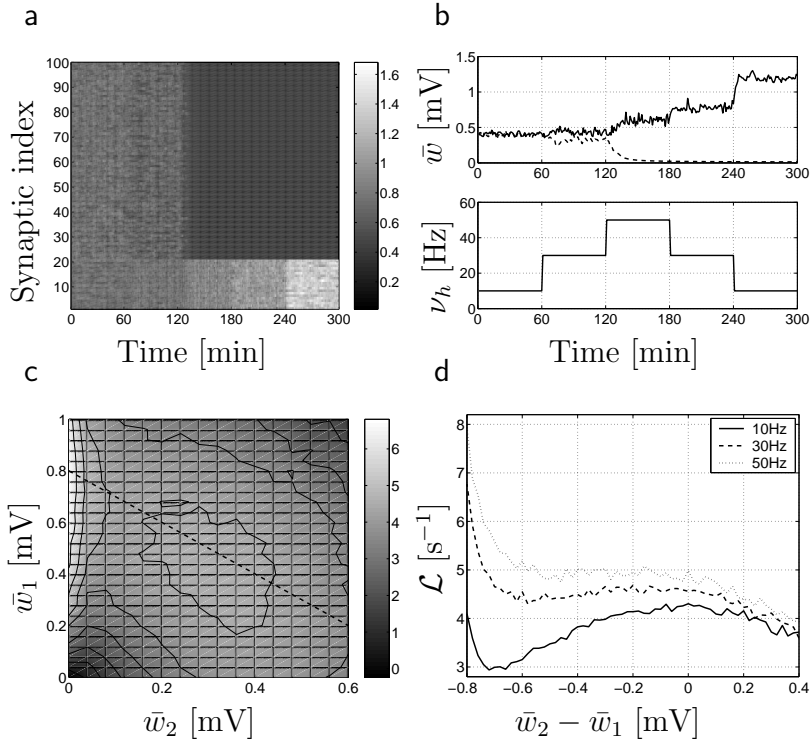
Figure 4: Potentiation and depression depend upon the presynaptic firing rate. 20 synapses of group **1** receive input with common rate modulation while the 80 synapses of group **2** (synapse index 21 - 100) receive Poisson input at a constant rate of 10Hz. The spike arrival rate in group **1** switches stochastically every 200ms between a low rate $\nu_l = 1$Hz and a high rate ν_h taken as a parameter. (a) Evolution of synaptic weights as a function of time for different amplitudes of rate modulation, that is ν_h changes from 10Hz during the first hour to 30 Hz, then to 50 Hz, 30 Hz, and back to 10 Hz. During the first two hours of stimulation, an unspecific distribution of synapses remains stable even though a slight decrease of weights in group **2** can be observed when the stimulus switches to $\nu_h = 30$Hz. A specialization of the synaptic pattern with large weights for synapses in group **1** is induced during the third hour of stimulation and remains stable thereafter. (b) Top: Mean synaptic weights (same data as in A) of group **1** ($\bar{w}_1$, *solid line*) and group **2** ($\bar{w}_2$, *dashed line*). Bottom: The stimulation paradigm, ν_h as a function of time. Note that at $\nu_h = 30$Hz (2nd and 4th hour of stimulation) both an unspecific pattern of synapses with little difference between $\bar{w}_1$ and $\bar{w}_2$ (2nd hour, top) and a highly specialized pattern (4th hour, top, large difference between *solid* and *dashed lines*) are possible. (c) The value of the objective function $\mathcal{L}$ (average value per second of time) in gray code as a function of the mean synaptic weight in group **1** (y-axis, $\bar{w}_1$) and group **2** (x-axis, $\bar{w}_2$) during stimulation with $\nu_h = 30$Hz. Two maxima can be perceived, i.e., a broad maximum for the unspecific synapse pattern ($\bar{w}_2 \approx \bar{w}_1 \approx 0.4$mV) and a pronounced elongated maximum for the specialized synapses pattern ($\bar{w}_1 \approx 0; \bar{w}_2 \approx 0.8$mV). The dashed line indicates a 1-dimensional section (see d) through the two-dimensional surface. (d) Objective function $\mathcal{L}$ as a function of the difference $\bar{w}_2 - \bar{w}_1$ between the mean synaptic weights in groups 2 and 1 along the line indicated in c. For $\nu_h = 30$Hz (*dashed*, same data as in c) two maxima are visible, i.e., a broad maximum at $\bar{w}_2 - \bar{w}_1 \approx 0$ and a narrow, but higher maximum corresponding to the specialized synapse pattern to the very left of the graph. For $\nu_h = 50$Hz (*dotted* line), the broad maximum at $\bar{w}_2 - \bar{w}_1 \approx 0$ disappears and only the specialized synapse pattern remains whereas for $\nu_h = 10$Hz (*solid* line) the broad maximum of the unspecific synapse pattern dominates.
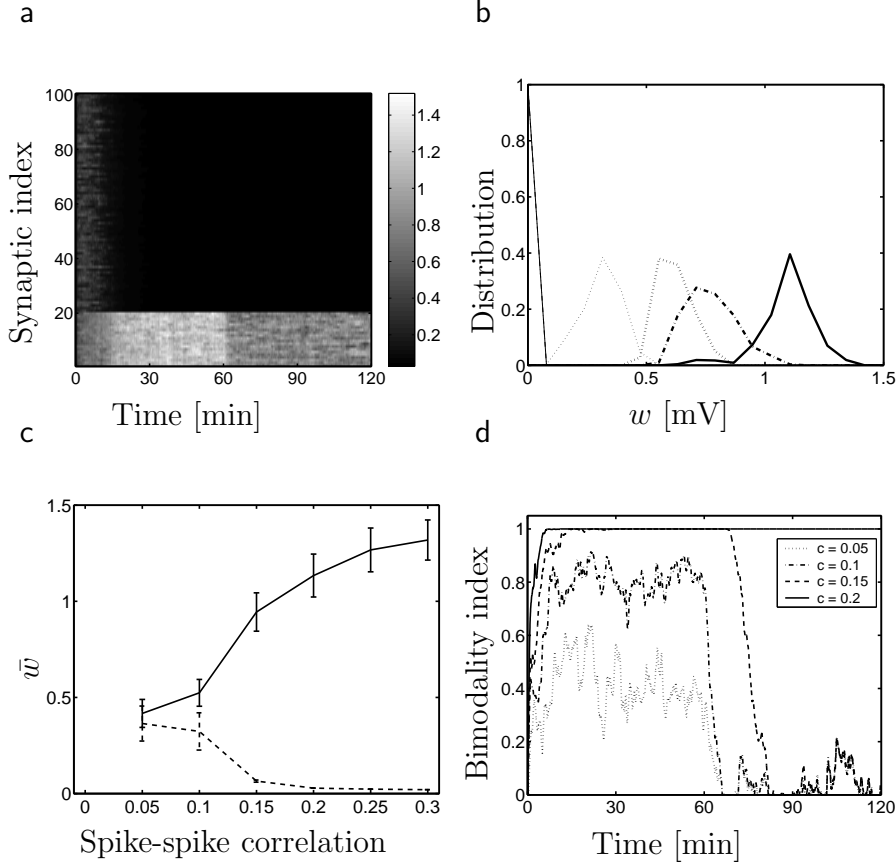
Figure 5: Synaptic memory induced by input with spike-spike correlations. (a) Evolution of 100 synaptic weights (vertical axis) as a function of time. During the first 60 minutes, synapses of group A ($j = 1, \ldots, 20$) receive a Poisson stimulation of 10 Hz with correlated spike input ($c = 0.2$) and those of group B ($j = 21, \ldots, 100$) receive uncorrelated spike input at 10 Hz. (b) Distribution of the EPSP amplitudes across the 100 synapses after $t = 5$ min. *(dotted line)*, $t = 60$ min. *(solid line)* and $t = 120$ min. *(dot-dashed line)*. The thick lines denote group A while the blue ones group B. (c) Mean EPSP amplitude of group A *(solid line)* and B *(dashed line)* at $t = 60$ min. for different values of correlation c of the input applied to group A. (d) Bimodality index b of the two groups of weights as a function of time. The memory induction by correlated spike input to group B stops at $t = 60$min. Memory retention is studied during the following 60 minutes. A bimodality index close to one implies as for the case with $c = 0.2$ implies that synaptic memory is well retained.

18

crossing points of the two Gaussians such that $\bar{w}_B < \hat{s} < \bar{w}_A$.

The two distributions (i.e., strong and weak synapses) started to separate within the first 5 minutes, and remained well separated even after the correlated memory-inducing stimulus was replaced by a random stimulus (Fig. 5B).

In order to study how synaptic memory retention depended on the induction paradigm, the experiment was then repeated with different values of the correlation index c that characterizes the spike-spike correlations during the induction period. For correlations $c < 0.1$ the two synaptic distributions are not well separated at the end of the induction period (bimodality index < 0.9), but well separated for $c \geq 0.15$ (Fig. 5C). Good separation at the end of the induction period alone is, however, not sufficient to guarantee retention of synaptic memory, since for the synaptic distribution induced by stimulation with $c = 0.15$, specialization breaks down at $t = 80$ min., i.e., after only 20 minutes of memory retention (Fig. 5D). On the other hand, for $c = 0.2$ and larger, synaptic memory is retained over several hours (only the first hour is plotted in Fig. 5D).

The long duration of synaptic memory in our model can be explained by the reduced adaptation speed of synapses with weights close to zero (postulate C2). If weak synapses change only slowly because of reduced adaptation speed, strong synapses must stay strong because of homeostatic processes that keep the mean activity of the postsynaptic neuron close to a target value. Moreover, the terms in the online learning rule derived from information maximization favor the bimodal distribution. Reduced adaptation speed of weak synapses could be caused by a cascade of intra-cellular biochemical processing stages with different time constants as suggested by Fusi et al. (2005). Thus our synaptic update rule allows for retention of synaptic memories over time scales that are significantly longer than the memory induction time, as necessary for any memory system. Nevertheless, synaptic memory in our model will eventually decay if random firing of pre- and postsynaptic neurons persists, in agreement with experimental results (Abraham et al. 2002; Zhou et al. 2003). We note that in the absence of presynaptic activity, the weights remain unchanged since the decay of synaptic weights is conditioned on presynaptic spike arrival; see Eq. (20).

Receptive field development

Synaptic plasticity is thought to be involved not only in memory (Hebb 1949), but also in the development of cortical circuits (Hubel and Wiesel 1962; Katz and Shatz 1996; von der Malsburg 1973; Bienenstock et al. 1982; Miller et al. 1989) and, possibly, cortical re-organization (Merzenich et al. 1984; Buonomano and Merzenich 1998). To study how our synaptic update rule would behave during development, we used a standard paradigm of input selectivity (Yeung et al. 2004), which is considered to be a simplified scenario of receptive field development. Our model neuron was stimulated by a Gaussian firing rate profile spanned across the 100 input synapses (Fig. 6A). The center

of the Gaussian was shifted every 200 ms to an arbitrarily chosen presynaptic neuron. In order to avoid border effects, neuron number 100 was considered a neighbor of neuron number 1, i.e., we can visualize the presynaptic neurons as being located on a ring.
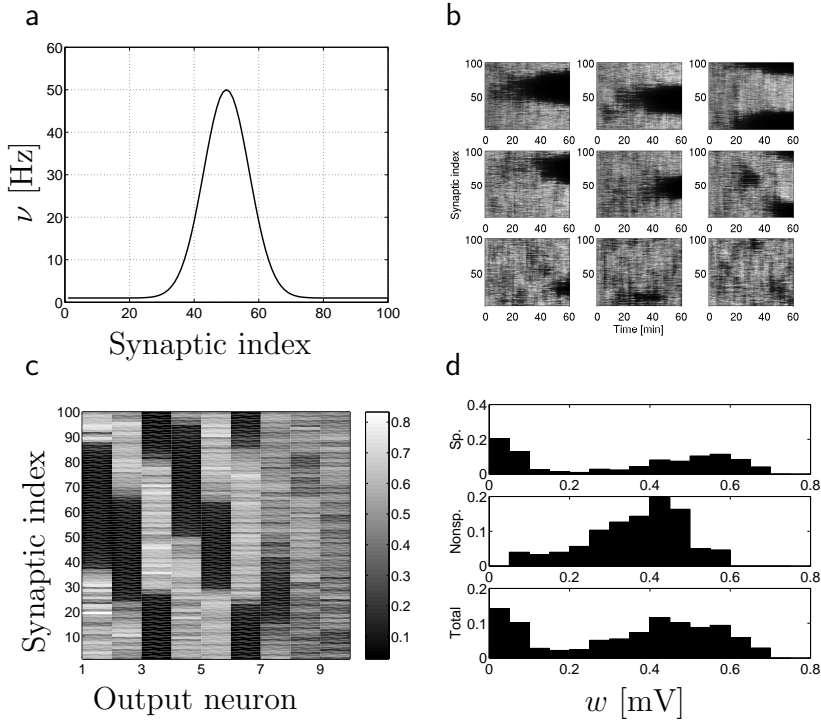

Figure 6: The synaptic update rule leads to input selectivity of the postsynaptic neuron. (a) Gaussian firing rate profile across the 100 presynaptic neurons. The center of the Gaussian is shifted randomly every 200 ms. Presynaptic neurons fire stochastically and send their spikes to nine postsynaptic neurons. (b) Evolution of synaptic weights of the nine postsynaptic neurons. Some neurons become specialized for a certain input pattern at the early phase of learning, others become specialized later, and the last three neurons have not yet become specialized. Since the input spike trains are identical for all the nine neurons, the specialization is due to noise in the spike generator of the postsynaptic neurons. (c) Final synaptic weight values of the nine output neurons after 1 hour of stimulus presentation. (d) The distribution of EPSP amplitudes after 1 hour of stimulation for (top) the specialized output neurons $1, 2, \ldots, 6$; (middle) for non-specialized neurons $7, 8, 9$; (bottom) for all nine output neurons.

Nine postsynaptic neurons with slightly different initial values of synaptic weights received identical input from the same set of 100 presynaptic neurons. During one hour of stimulus presentation, six out of the nine neurons developed synaptic specialization leading to input selectivity. The optimal stimulus for these six neurons varies (Fig. 6B and C), so that any Gaussian stimulus at

an arbitrary location excites at least one of the postsynaptic neurons. In other words, the six postsynaptic neurons have developed input selectivity with different but partially overlapping receptive fields. The distribution of synaptic weights for the selective neurons is bimodal with a first peak for very weak synapses (EPSP amplitudes less than 0.1 mV) and a second peak around EPSP amplitudes of 0.6 mV (Fig. 6D); the amplitude distribution of the unselective neurons is broader with a single peak at around 0.4mV.

The number of postsynaptic neurons showing synaptic specialization depends on the total stimulation time and the strength of the stimulus. If the stimulation time is extended to 3 hours instead of 1 hour, all postsynaptic neurons become selective using the same stimulation parameters as before. However, if the maximal presynaptic firing rate at the center of the Gaussian is reduced to 40 Hz instead of 50 Hz only six out of nine are selective after three hours of stimulation; and with a further reduction of the maximal rate to 30 Hz only a single neuron is selective after 3 hours of stimulation (data not shown). We hypothesize that the coexistence of unselective and selective neurons during development could explain the broad distribution of EPSP amplitudes seen in some experiments (e.g. (Sjöström et al. 2001) in rat visual cortex). For example, if we sample the synaptic distribution across all 9 postsynaptic cells we find the distribution shown at the bottom of Fig. 6D. If the number of unspecific neurons were higher, the relative importance of synapses with EPSP amplitudes of less than 0.1 mV would diminish. If the number of specialized neurons increased, the distribution would turn into a clearcut bimodal one which would be akin to sampling an ensemble of two-state synapses with all-or-none potentiation on a synapse-by-synapse basis ((Petersen et al. 1998) in rat hippocampus).

4   Discussion

What can we and what can we not expect from optimality models?
Optimality models can be used to clarify concepts, but they are unable to make specific predictions about molecular implementations. In fact, the synaptic update rule derived in this paper shares functional features with STDP and classical LTP, but it is blind with respect to interesting questions such as the role of NMDA, Kainate, endocannabinoid, or CaMKII in the induction and maintenance of potentiation and depression (Bliss and Collingridge 1993; Bortolotto et al. 2003; Frey and Morris 1997; Lisman 2003; Malenka and Nicoll 1993; Sjöström et al. 2004). If molecular mechanisms are in the focus of interest, detailed 'mechanistic' models of synaptic plasticity (Senn et al. 2001; Yeung et al. 2004) should be preferred. On the other hand, the mere fact that similar forms of LTP or LTD seem to be implemented across various neural systems by different molecular mechanisms leads us to speculate that common functional roles of synapses are potentially more important for understanding synaptic dynamics than the specific way that these functions are implemented.

Ideally, optimality approaches such as the one developed in this paper should be helpful to put seemingly diverse experimental or theoretical results into a coherent framework. We have listed in the introduction a couple of points, partially linked to experimental results, partially linked to earlier theoretical investigations. Our aim has been to connect these points and trace them back to a small number of basic principles. Let us return to our initial list and discuss the points in light of the results of the preceding section.

Correlations.

Hebb postulated an increase in synaptic coupling in case of repeated coactivation of pre- and postsynaptic neurons as a useful concept for memory storage in recurrent networks (Hebb 1949). In our optimality framework defined by Eq. (3) correlation dependent learning is not imposed explicitly but arises from the maximization of information transmission between pre- and postsynaptic neurons. Indeed, information transmission is only possible if there are correlations between pre- and postsynaptic neurons. Information transmission is maximized if these correlations are increased. Gradient ascent of the information term hence leads to a synaptic update rule that is sensitive to correlations between pre- and postsynaptic neurons; see Eq. (20). An increase of synaptic weights enhances these correlations and maximizes information transmission. We emphasize that, in contrast to Hebb (1949) we do not invoke memory formation and recall as a reason for correlation dependence, but information transmission. Similar to other learning rules (Linsker 1986; Oja 1982), the sensitivity of our update rule to correlations between pre- and postsynaptic neurons gives the synaptic dynamics a sensitivity to correlations in the input as demonstrated in Figs. 4 – 6.

Input selectivity.

During cortical development, cortical neurons develop input selectivity typically quantified as the width of receptive fields (Hubel and Wiesel 1962). As illustrated in the scenario of Fig. 6, our synaptic update rule shows input selectivity and stands hence in the research tradition of many other studies (von der Malsburg 1973; Bienenstock et al. 1982; Miller et al. 1989). Input selectivity in our model arises through the combination of the correlation sensitivity of synapses discussed above with the homeostatic term D in Eq. (3) (Toyoizumi et al. 2005). Since the homeostatic term keeps the mean rate of the postsynaptic neuron close to a target rate, it leads effectively to a normalization of the total synaptic input similar to the sliding threshold mechanism in the Bienenstock-Cooper-Munro rule (Bienenstock et al. 1982). Normalization of the total synaptic input through firing rate stabilization has also been seen in previous STDP models (Song et al. 2000; Kempter et al. 1999; Kempter et al. 2001). Its effect is similar to explicit normalization of synaptic weights which is a well-known mechanism to induce input selectivity (von der Malsburg 1973; Miller and MacKay 1994), but our model does not need an explicit normalization step. So far we studied only a single or a small number of independent postsynaptic neurons, but we expect that, as in many other

studies, e.g., (Erwin et al. 1995; Song and Abbott 2001; Cooper et al. 2004), our synaptic update rule would yield feature maps if applied to a network of many weakly interacting cortical neurons.

Unimodal versus bimodal synapse distributions.

In several previous models of rate-based or spike-timing based synaptic dynamics, synaptic weights evolved *always* towards a bimodal distribution with some synapses close to zero and others close to maximal weight (Miller et al. 1989; Miller and MacKay 1994; Kempter et al. 1999; Song and Abbott 2001; Toyoizumi et al. 2005). Thus synapses specialize on certain features of the input, even if the input has weak or no correlation at all which seems questionable from a functional point of view and which disagrees with experimentally found distributions of EPSP amplitudes in rat visual cortex (Sjöström et al. 2001). As shown in Figs. 4 — 6, an unspecific pattern of synapses with a broad distribution of EPSP amplitudes is stable with our synaptic update rule if correlations in the input are weak. Synaptic specialization (bimodal distribution) develops only if synaptic inputs show a high degree of correlations, either on the level of spikes or firing rates. Thus neurons only specialize on highly significant input features, and not in response to noise. A similar behavior was noted in two recent models on spike-timing dependent plasticity (Gütig et al. 2003; Yeung et al. 2004). Going beyond those studies, we also demonstrated stability of the specialized synapse distribution over some time, even if the amount of correlation through rate-modulation is reduced after induction of synaptic specialization stimulation (Fig. 4). Thus, for the same input characteristics our model can show unimodal or bimodal distribution of synaptic weight, depending on the stimulation history. Moreover, we note that in the state of bimodal distribution, the EPSP amplitude at the depressed synapses are so small, that in an experimental setting they could easily remain undetected or classified as 'silent' synapse (Kullmann 2003). A large proportion of silent synapses has been previously shown to be consistent with optimal memory storage (Brunel et al. 2004). Furthermore, our results show that in a given population of cells, neurons with specialized synapses can coexist with others that have a broad and unspecific synapse distribution. We speculate that these non-specialized neurons could then be recruited later for new stimuli, as hypothesized in earlier models of neural networks (Grossberg 1987).

Synaptic memory.

In the absence of presynaptic input, synapses in our model do not change significantly. Furthermore, our results show that even in the presence of random pre- and postsynaptic firing activity, synaptic memories can be retained over several hours, even though a slow decay occurs. Essential for long-term memory maintenance under random spike arrival is the reduced adaptation speed for small values of synaptic weights as formulated in postulate **C2**. This postulate is similar in spirit to a recent theory of Fusi et al. (2005), with two important differences. Firstly, we work with a continuum of synaptic states

(characterized by the value of w) whereas Fusi et al. assume a cascade of a finite number of discrete internal synaptic states where transitions are uni-directional and characterized by different time constants (Fusi et al. 2005). The scheme of Fusi et al. guarantees slow (i.e., not exponential) decay of memories, whereas in our case decay is always exponential even if the time constant is weight-dependent. Secondly, Fusi et al. use a range of different time constants for both the potentiated and the unpotentiated state, whereas in our model it is sufficient to have a slower time constant for the unpoten-tiated state only. If the change of unpotentiated synapses is slow compared to homeostatic regulation of mean firing rate, then the maintenance of the strong synapses is given by homeostasis (and also supported by information maximization).

The fact that our model has been formulated in terms of a continuum of synaptic weights was for convenience only. Alternatively it is conceivable to define a number of internal synaptic states that give rise to binary, or a small number of, synaptic weight values (Fusi 2002; Fusi et al. 2005; Abarbanel et al. 2005). The actual number of synaptic states is unknown with conflicting evidence (Petersen et al. 1998; Lisman 2003).

Rate dependence.

Our results show that common rate modulation in one group of synapses strengthens these synapses if the modulation amplitude is strong enough. In contrast, an increase of rates to a fixed value of 40 Hz (without modulation) in one group of synapses while another group of synapses receives background firing at 10 Hz does not lead to a synaptic specialization, but only to a minor readjustment of weights (data not shown). For a comparison with experi-mental results it is important to note that rate dependence is typically mea-sured with extracellular stimulation of presynaptic pathways. We assimilate repeated extracelluar stimulation with a strong and common modulation of spike arrival probability at one group of synapses (as opposed to an increased rate of a homogeneous Poisson process). Under this interpretation, our results are qualitatively consistent with experiments.

STDP.

Our results show that our synaptic update rule shares several features with STDP as found in experiments (Markram et al. 1997; Bi and Poo 1998; Bi and Poo 2001; Sjöström et al. 2001). The time scale of the potentiation part of the STDP function depends in our model on the duration of EPSPs. The time scale of the depression part is determined by the duration of EPSP suppression in agreement with experiments (Froemke et al. 2005). Our model shows that the relative importance of LTP and LTD depend on the initial value of the synaptic weight in a way similar to that found in experiments (Bi and Poo 1998). However, for EPSP amplitudes between 0.5 and 2 mV the LTP part clearly dominates over LTD in our model which seems to have less experimental support. Also, the frequency dependence of STDP in our model is less pronounced than in experiments.

Models of STDP have previously been formulated on a phenomenological level (Gerstner et al. 1996; Song et al. 2000; Gerstner and Kistler 2002) or on a molecular level (Lisman 2003; Senn et al. 2001; Yeung et al. 2004). Only recently models derived from optimality concepts have moved into the center of interest (Toyoizumi et al. 2005; Toyoizumi et al. 2005a; Bohte and Mozer 2005; Chechik 2003). There are important differences of the present model to the existing 'optimal' models. Chechik (2003) used information maximization, but limited his approach to static input patterns, while we consider arbitrary inputs. Bell and Parra (2005) minimize output entropy and Bohte and Mozer (2005) maximize spike reliability whereas we maximize the information between full input and output spike trains. None of these studies considered optimization under homeostatic and maintenance cost constraints.

After we had introduced the homeostatic constraint in a previous study which gave rise to a learning rule with several interesting properties (Toyoizumi et al. 2005), we realized that this model did not exhibit properties of STDP in an in vitro situation without some additional assumptions. Indeed, as shown in Fig. 3D, the STDP function derived from information maximization alone exhibits no depression in the absence of an additional weight-dependent cost term in the optimality function. The weight-dependent cost term introduced in this paper plays hence a crucial role in STDP since it shifts the STDP function to more negative values.

How realistic is a weight-dependent cost term?

The weight-dependent cost term Ψ in the optimality criterion $\mathcal{L}$ depends quadratically on the value of the weights of all synapses converging onto the same postsynaptic neuron. This turns out to be equivalent to a 'decay' term in the synaptic update; see last term on the right-hand side of Eq. (20). Such decay terms are common in the theoretical literature (Bienenstock et al. 1982; Oja 1982), but the question is whether such a decay term (leading to a slow depression of synapses) is realistic from a biological point of view.

We emphasize that the decay term in our synaptic update rule is proportional to presynaptic activity. Thus, in contrast to existing models in the theoretical literature (Bienenstock et al. 1982; Oja 1982), a synapse which receives no input is protected against slow decay. The specific form of the 'decay' term considered in this paper was such that synaptic weights decreased with each spike arrival, but presynaptic activity could also be represented in the decay term by the mean firing rate rather than spike arrival, with no qualitative changes to the results.

An important aspect of our cost term is that only synapses that have recently been activiated are at risk regarding weight decay. We speculate that the weight-dependent cost term could, in a loose sense, be related to the amount of plasticity factors that synapses require and compete for during the first hours of synaptic maintenance (Fonseca et al. 2004). According to the synaptic tagging hypothesis (Frey and Morris 1997), only those synapses that have been activated in the recent past compete for plasticity factors, while unpotentiated

synapses do not suffer from decay (Fonseca et al. 2004). We emphasize that such a link of our cost term to the competetion for plasticity factors is purely hypothetical. Many relevant details of tagging and competition for synaptic maintenance are omitted in our approach.

Predictions and experimental tests.

In order to achieve synaptic memory that is stable over several hours, the reduced adaptation speed for weak synapses formulated in postulate **C2** turns out to be essential. Thus an essential assumption of our model is testable: for synapses with extremely small EPSP amplitudes, in particular 'silent synapses', the induction of both LTP and LTD should require stronger stimulation or stimulation sustained over longer times, compared to synapses that are of average strength. This aspect is distinct from other models (Gütig et al. 2003) which postulate for weak synapses a reduced adaptation speed for depression only, but maximal effect of potentiation. Thus, comparison of LTP induction for silent synapses (Kullmann 2003) with that for average synapses should allow to differentiate between the two models. In an alternative formulation, synaptic memory could also be achieved by making strong synapses resist to further changes. As an aside we note that the model of Fusi et al. (2005) assumes a reduced speed of transition between several *internal* synaptic states so that transition would not necessarily be visible as a change in synaptic weight.

A second test of our model concerns the pattern of synaptic weights converging on the same postsynaptic neuron. Our results suggest that early in development most neurons would show an unspecific synapse pattern, i.e., a distribution of EPSP amplitudes with a single, but broad peak whereas later a sizeable fraction of neurons would show a pattern of synaptic specialization with some strong synapses and many silent ones, i.e., a bimodal distribution of EPSP amplitudes. Ideally the effect would be seen by scanning all the synapses of individual postsynaptic neurons; it remains to be seen if modern imaging and staining methods will allow to do this. Alternatively, by electrophysiological methods, distributions of synaptic strengths could be built up by averaging over many synapses on different neurons (Sjöström et al. 2001). In this case, our model would predict that during development the histogram of EPSP amplitudes would change in two ways (Fig. 6D): (i) the number of silent synapses increases so that the amplitude of the sharp peak at small EPSP amplitude grows; and (ii) the location of the second, broader, peak shifts to larger values of the EPSP amplitudes. Furthermore, and in contrast to other models where the unimodal distribution is unstable, the transition to a bimodal distribution depends in our model on the stimulation paradigm.

Limitations and extensions

We would like to emphasize that properties of our synaptic update rules have so far only been tested for single neurons in an unsupervised learning paradigm. Extensions are possible in several directions. Firstly, instead of single neurons, a large recurrent network could be considered. This could on

one side further our understanding of the model properties in the context of cortical map development (Erwin et al. 1995), on the other side scrutinize the properties of the synaptic update rule as a functional memory in recurrent networks (Amit and Fusi 1994). Secondly, instead of unsupervised learning where the synaptic update rule treats all stimuli alike whether they are behaviorally relevant or not, a reward-based learning scheme could be considered (Dayan and Abbott 2001; Seung 2003). Behaviorally relevant situations can be taken into account by optimizing reward instead of information transmission (Schultz et al. 1997).

Acknowledgments

A   Determination of the Parameter λ

The parameter λ is set to give $\Delta w_j = 0$ for large enough $|t^{pre} - t^{post}|$ in an simulated STDP in vitro paradigm. In order to find the appropriate value of λ, we separately consider the effects of a presynaptic spike and a postsynaptic one - which is possible since they are assumed to occur at a large temporal distance. Since a postsynaptic spike alone does not change synaptic strength ($C_j(t) = 0$, always), we choose a λ that gives no synaptic change when a presynaptic spike alone is induced. For a given presynaptic spike at t_j^f, we have

$$
\begin{aligned}
C_j(t) &= -g \int_0^t \epsilon(t' - t_j^f) e^{-(t-t')/\tau_m} dt' \\
&= -g \frac{\tau_C \tau_m}{\tau_C - \tau_m} \left[e^{-(t-t_j^f)/\tau_C} - e^{-(t-t_j^f)/\tau_m} \right].
\end{aligned} \tag{22}
$$

Since $\tilde{\rho} \approx \rho_r$, and $\bar{\rho} \approx \rho_r$ in this in vitro setting, the factor B^{post} in Eq. (14) is approximated as

$$
B^{post}(t) \approx -w_j g e^{-(t-t_j^f)/\tau_m}. \tag{23}
$$

Hence, we find the effect of a presynaptic spike as

$$
\Delta w = \int_0^T \frac{dw_j}{dt} dt = w_j g^2 \frac{\tau_C \tau_m}{\tau_C - \tau_m} \left[\frac{\tau_C \tau_m}{\tau_C + \tau_m} - \frac{\tau_m}{2} \right] - \lambda w_j. \tag{24}
$$

The condition of no synaptic change gives $\lambda = g^2 \frac{\tau_m \tau_C}{\tau_C - \tau_m} \left(\frac{\tau_m \tau_C}{\tau_m + \tau_C} - \frac{\tau_m}{2} \right)$. We used this λ in the numerical code.

B   Weight Dependent Evaluation of the Optimality Criterion

In Fig. 4 C and D the optimality criterion has been evaluated as a function of some artifical weight distribution. Specifically, values of synaptic weights have been chosen stochastically from two Gaussian distributions with mean $\bar{w}_1$ and standard deviation σ_1 for group 1 and $\bar{w}_2$ and σ_2 for group 2. In order to account for differences in standard deviations due to the weight-dependent update rate $\alpha(w)$, we chose $\sigma(\bar{w}) = 0.1\text{mV} \cdot \bar{w}^4/(w_s^4 + \bar{w}^4)$ which gives a variance of synaptic weights in both groups which is consistent with the variance seen in Fig. 4A.

For a fixed set of synaptic weight values, the network is simulated during a trial time of 30 minutes while the synaptic updated rule has been turned off and the objective function $\mathcal{L}$ defined in Eq. (3) is evaluated using $\bar{\rho}^{est}$ from Eq. (21). The result is divided by the trial time t and plotted in Fig. 4 C and D in units of s^{-1}. The mesh size of mean synaptic strength is 0.04 mV. The one dimensional plot in Fig. 4 D is taken along the direction $\bar{w}_2 = 0.8 \text{ mV} - \bar{w}_1$.

References

Abarbanel, H., R. Huerta, and M. Rabinovich (2002). Dynamical model of long-term synaptic plasticity. *Proc. Natl. Academy of Sci. USA 59*(10137-10143).

Abarbanel, H., S. Talathi, L. Gibb, and M. Rabinovich (2005). Synaptic plasticity with discrete state synapses. *to appear xx*(xx).

Abraham, W., B. Logan, J. Greenwood, and M. Dragunow (2002). Induction and experience-dependent consolidation of stable long-term potentiation lasting months in the hippocampus. *J. Neuroscience 22*, 9626 – 9634.

Amit, D. and S. Fusi (1994). Learning in neural networks with material synapses. *Neural Computation 6*, 957–982.

Artola, A., S. Bröcher, and W. Singer (1990). Different voltage dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature 347*, 69–72.

Atick, J. and A. Redlich (1990). Towards a theory of early visual processing. *Neural Computation 4*, 559–572.

Barlow, H. (1956). Retinal noise and absolute threshold. *J. Opt. Soc. Am. 46*, 634–639.

Barlow, H. (1980). The absolute efficiency of perceptual decisions. *Phil. Trans. R. Soc. Lond. Ser. B 290*, 71–82.

Barlow, H. and W. Levick (1969). 3 factors limiting reliable detection of light by retinal ganglion cells of cat. *J. Physiology (London) 200*, 1–24.

Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenbluth (Ed.), *Sensory Communication*, pp. 217–234. MIT Press.

Barlow, H. B. (1989). Unsupervised learning. *Neural. Comp. 1*, 295–311.

Bell, A. and T. Sejnowski (1995). An information maximization approach to blind separation and blind deconvolution. *Neural Computation 7*, 1129–1159.

Bell, A. J. and L. C. Parra (2005). Maximising sensitivity in a spiking network. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 121–128. Cambridge, MA: MIT Press.

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

Bi, G. and M. Poo (2001). Synaptic modification of correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci. 24*, 139–166.

Bienenstock, E., L. Cooper, and P. Munroe (1982). Theory of the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience 2*, 32–48. reprinted in Anderson and Rosenfeld, 1990.

Bliss, T. and T. Lomo (1973). Long-lasting potentiation of synaptic transmission in the dendate area of anaesthetized rabbit following stimulation of the perforant path. *J. Physiol. 232*, 351–356.

Bliss, T. V. P. and G. L. Collingridge (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature 361*, 31–39.

Bohte, S. M. and M. C. Mozer (2005). Reducing spike train variability: A computational theory of spike-timing dependent plasticity. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 201–208. Cambridge, MA: MIT Press.

Bortolotto, Z. A., S. Lauri, J. T. R. Isaac, and G. L. Collingridge (2003). Kainate receptors and the induction of mossy fibre long-term potentiation. *Phil. Trans. R. Soc. Lond B: Biological Sciences 358*(657 - 666).

Britten, K., M. Shadlen, W. Newsome, and J. Movshon (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *J. Neuroscience 12*, 4745–4765.

Brunel, N., V. Hakim, P. Isope, J.-P. Nadal, and B. Barbour (2004). Optimal information storage and the distribution of synaptic weights: Perceptron versus purkinje cell. *Neuron 43*, 745–757.

Buonomano, D. and M. Merzenich (1998). Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience 21*, 149–186.

Castellani, G., E. Quinlan, L. Shouval, and H. Cooper (2001). A biophysical model of bidirectional synaptic plasticity: dependence on ampa and nmda receptors. *Proc. Natl. Acad. Sci. USA 98*, 12772–12777.

Chechik, G. (2003). Spike-timing-dependent plasticity and relevant mututal information maximization. *Neural Computation 15*, 1481–1510.

Cooper, L., N. Intrator, B. Blais, and H. Z. Shouval (2004). *Theory of cortical plasticity.* Singapore: World Scientific.

Cover, T. and J. Thomas (1991). *Elements of Information Theory.* New York: Wiley.

D. Liao, N.A. Hessler, R. M. (1995). Activation of postsynaptically silent synapses during pairing-induced ltp in ca1 region of hippocampal slice. *Nature 375*, 400–404.

Dan, Y. and M. Poo (2004). Spike timing-dependent plasticity of neural circuits. *Neuron 44*, 23–30.

Dayan, P. and L. F. Abbott (2001). *Theoretical Neuroscience.* Cambridge: MIT Press.

de Ruyter van Steveninck, R. R. and W. Bialek (1995). Reliability and statistical efficiency of a blowfly movement-sensitive neuron. *Phil. Trans. R. Soc. Lond. Ser. B 348*, 321–340.

Dudek, S. M. and M. F. Bear (1992). Homosynaptic long-term depression in area ca1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. *Proc. Natl. Acad. Sci. USA 89*, 4363–4367.

Erwin, E., K. Obermayer, and K. Schulten (1995). Models of orientation and ocular dominance columns in the visual cortex: a critcal comparison. *Neural Comput. 7*, 425–468.

Fairhall, A. L., G. Lewen, W. Bialek, and R. van Steveninck (2001). Efficiency and ambiguity in an adaptive neural code. *Nature 412*, 787–792.

Fonseca, R., U. Nägerl, R. Morris, and T. Bonhoeffer (2004). Competition for memory: Hippocampal ltp under the regimes of reduced protein synthesis. *Neuron 44*, 1011–1020.

Frey, U. and R. Morris (1997). Synaptic tagging and long-term potentiation. *Nature 385*, 533 – 536.

Froemke, R. C., M.-M. Poo, and Y. Dan (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature 434*, 221–225.

Fusi, S. (2002). Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol. Cybern. 87*, 459–470.

Fusi, S., M. Annunziato, D. Badoni, A. Salamon, and D.J.Amit (2000). Spike-driven synaptic plasticity: theory, simulation, vlsi implementation. *Neural Computation 12*, 2227–2258.

Fusi, S., P. Drew, and L. Abbott (2005). Cascade models of synaptically stored memories. *Neuron 45*, 599–611.

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

Gerstner, W. and W. K. Kistler (2002). *Spiking Neuron Models*. Cambridge UK: Cambridge University Press.

Grossberg, S. (1987). *The Adaptive Brain I*. Elsevier.

Gütig, R., R. Aharonov, S. Rotter, and H. Sompolinsky (2003). Learning input correlations through non-linear temporally asymmetry hebbian plasticity. *J. Neuroscience 23*, 3697–3714.

Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley.

Hertz, J., A. Krogh, and R. G. Palmer (1991). *Introduction to the Theory of Neural Computation*. Redwood City CA: Addison-Wesley.

Holmes, W. R. and W. B. Levy (1990). Insights into associative long-term potentiation from computational models of nmda receptor-mediated calcium influx and intracellular calcium concentration changes. *J. Neurophysiol. 63*, 1148–1168.

Hubel, D. H. and T. N. Wiesel (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (London) 160*, 106–154.

Isaac, J. T., R. A. Nicoll, and R. C. Malenka (1995). Evidence for silent synapses: Implications for the expression of ltp. *Neuron 15*, 427–434.

Karmarkar, U. and D. Buonomano (2002). A model of spike-timing dependent plasticity: one or two coincidence detectors. *J. Neurophysiology 88*, 507–513.

Katz, L. C. and C. J. Shatz (1996). Synaptic activity and the construction of cortical circuits. *Science 274*, 1133–1138.

Kelso, S. R., A. H. Ganong, and T. H. Brown (1986). Hebbian synapses in hippocampus. *Proc. Natl. Acad. Sci. USA 83*, 5326–5330.

Kempter, R., W. Gerstner, and J. L. van Hemmen (1999). Hebbian learning and spiking neurons. *Phys. Rev. E 59*, 4498–4514.

Kempter, R., W. Gerstner, and J. L. van Hemmen (2001). Intrinsic stabilization of output rates by spike-based hebbian learning. *Neural Computation 13*, 2709–2741.

Kempter, R., C. Leibold, H. Wagner, , and J. L. van Hemmen (2001). Formation of temporal-feature maps by axonal propagation of synaptic learning. *Proc. Natl. Acad. of Sci. USA 98*, 4166–4171.

Kistler, W. M. and J. L. van Hemmen (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput. 12*, 385–405.

Kullmann, D. M. (2003). Silent synapses: what are they telling us about long-term potentiation? *Phil. Trans. R. Soc. Lond B: Biological Sciences 358*, 727 – 733.

Laughlin, S. (1981). A simple coding procedure enhances a neurons information capacity. *Z. Naturforschung 36*, 910–912.

Laughlin, S., R. R. deRuyter van Steveninck, and J. Anderson (1998). The metabolic cost of neural information. *Nature Neuroscience 1*(36-41).

Levy, W. and R. Baxter (2002). Energy-efficient neuronal computation via quantal synaptic failures. *J. Neuroscience 22*, 4746–4755.

Linsker, R. (1986). From basic network principles to neural architecture: emergence of spatial-opponent cells. *Proc. Natl. Acad. Sci. USA 83*, 7508–7512.

Linsker, R. (1989). How to generate ordered maps by maximizing the mutual information between input and output signals. *Neural Computation 1*(3), 402–411.

Lisman, J. (1989). A mechanism for hebb and anti-hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA 86*, 9574–9578.

Lisman, J. (2003). Long-term potentiation: outstanding questions and attempted synthesis. *Phil. Trans. R. Soc. Lond B: Biological Sciences 358*, 829 – 842.

Malenka, R. C., J. Kauer, R. Zucker, and R. A. Nicoll (1988). Postsynaptic calcium is sufficient for potentiation of hippocampal synaptic transmission. *Science 242*, 81–84.

Malenka, R. C., J. A. Kauer, D. J. Perkel, M. D. Mauk, and P. T. Kelly (1989). An essential role for postsynaptic calmodulin and protein kinase activity in long-term potentiation. *Nature 340*, 554–557.

Malenka, R. C. and R. A. Nicoll (1993). Nmda-receptor-dependent plasticity: multiple forms and mechanisms. *Trends Neurosci. 16*, 480–487.

Markram, H., J. Lübke, M. Frotscher, and B. Sakmann (1997). Regulation of synaptic efficacy by coincidence of postysnaptic AP and EPSP. *Science 275*, 213–215.

Merzenich, M., R. Nelson, M. Stryker, M. Cynader, A. Schoppmann, and J. Zook (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *J. Comparative Neurology 224*, 591–605.

Miller, K., J. B. Keller, and M. P. Stryker (1989). Ocular dominance column development: analysis and simulation. *Science 245*, 605–615.

Miller, K. D. and D. J. C. MacKay (1994). The role of constraints in hebbian learning. *Neural Computation 6*, 100–126.

Nadal, J.-P. and N. Parga (1997). Redundancy reduction and independent component analysis: Conditions on cumulants and adaptive approaches. *Neural Computation 9*, 1421–1456.

Oja, E. (1982). A simplified neuron model as a principal component analyzer. *J. Mathematical Biology 15*, 267–273.

Petersen, C., R. Malenka, R. Nicoll, and J. Hopfield (1998). All-or-none potentiation of ca3-ca1 synapses. *Proc. Natl. Acad. Sci. USA 95*, 4732–4737.

Rubin, J., D. D. Lee, and H. Sompolinsky (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Physical Review Letters 86*, 364–367.

Schultz, W., P. Dayan, and R. Montague (1997). A neural substrate for prediction and reward. *Science 275*, 1593–1599.

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

Seung, H. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron 40*, 1063–1073.

Sjöström, P. and S. Nelson (2002). Spike timing, calcium signals and synaptic plasticity. *Current Opinion in Neurobiology 12*, 305–314.

Sjöström, P., G. Turrigiano, and S. Nelson (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron 32*, 1149–1164.

Sjöström, P., G. Turrigiano, and S. Nelson (2004). Endocannabinoid-dependent neocortical layer-5 ltd in the absence of postsynaptic spiking. *J. Neurophysiol. 92*, 3338–3343.

Song, S. and L. Abbott (2001). Column and map development and cortical re-mapping through spike-timing dependent plasticity. *Neuron 32*(339-350).

Song, S., K. Miller, and L. Abbott (2000). Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience 3*, 919–926.

Sutton, R. and A. Barto (1998). *Reinforcement learning*. MIT Press, Cambridge.

Toyoizumi, T., J.-P. Pfister, K. Aihara, and W. Gerstner (2005). Generalized bienenstock-cooper-munro rule for spiking neurons that maximizes information transmission. *Proc. National Academy Sciences (USA) 102*, 5239–5244.

Toyoizumi, T., J.-P. Pfister, K. Aihara, and W. Gerstner (2005a). Spike-timing dependent plasticity and mutual information maximization for a spiking neuron model. In L. K.

Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 1409–1416. Cambridge, MA: MIT Press.

Turrigiano, G. and S. Nelson (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience 5*, 97–107.

van Rossum, M. C. W., G. Q. Bi, and G. G. Turrigiano (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience 20*, 8812–8821.

von der Malsburg, C. (1973). Self-organization of orientation selective cells in the striate cortex. *Kybernetik 14*, 85–100.

Yeung, L. C., H. Shouval, B. Blais, and L. Cooper (2004). Synaptic homeostasis and input selectivity follow from a model of calcium dependent plasticity. *Proc. Nat. Acad. Sci. USA, 101*, 14943–14948.

Zador, A., C. Koch, and T. H. Brown (1990). Biophysical model of a hebbian synapse. *Proc. Natl. Acad. Sci. 87*, 6718–6722.

Zhang, L., H. Tao, C. Holt, W.A.Harris, and M.-M. Poo (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature 395*, 37–44.

Zhou, Q., H. Tao, and M. Poo (2003). Reversal and stabilization of synaptic modifications in the developping visual system. *Science 300*, 1953–1957.