LETTER Communicated by David Horn

Spike-Timing-Dependent Plasticity and Relevant Mutual Information Maximization

Gal Chechik
ggal@cs.huji.ac.il
Interdisciplinary Center for Neural Computation, Hebrew University, Jerusalem, Israel

Synaptic plasticity was recently shown to depend on the relative timing of the pre- and postsynaptic spikes. This article analytically derives a spike-dependent learning rule based on the principle of information maximization for a single neuron with spiking inputs. This rule is then transformed into a biologically feasible rule, which is compared to the experimentally observed plasticity. This comparison reveals that the biological rule increases information to a near-optimal level and provides insights into the structure of biological plasticity. It shows that the time dependency of synaptic potentiation should be determined by the synaptic transfer function and membrane leak. Potentiation consists of weight-dependent and weight-independent components whose weights are of the same order of magnitude. It further suggests that synaptic depression should be triggered by rare and relevant inputs but at the same time serves to unlearn the baseline statistics of the network's inputs. The optimal depression curve is uniformly extended in time, but biological constraints that cause the cell to forget past events may lead to a different shape, which is not specified by our current model. The structure of the optimal rule thus suggests a computational account for several temporal characteristics of the biological spike-timing-dependent rules.

1 Introduction

Temporal aspects of synaptic connections have recently become a key research topic. First, it was demonstrated that synapses exhibit short-term facilitation and depression (Abbott, Varela, Sen, & Nelson, 1997; Markram, Lubke, Frotscher, & Sakmann, 1997), and synaptic connections with various dynamics were shown to be expressed following learning (Monyer, Burnashev, Laurie, Sakmann, & Seeburg, 1994; Tang et al., 1999). Moreover, there is now ample evidence that changes in synaptic efficacies depend on the relative timing of pre- and postsynaptic spikes. This conflicts with the concept of learning correlated activity ("those who fire together wire together"), which was the common interpretation of Hebbian learning several years ago. Recent findings show that synapses between two excitatory neurons are potentiated when a presynaptic spike closely precedes a postsynaptic

one, whereas spiking activity in reverse order results in synaptic depression. These effects have been demonstrated in a variety of systems and preparations such as hippocampal studies in vivo (Levy & Steward, 1983), cultures of hippocampal neurons (Debanne, Gahwiler, & Thompson, 1994; Bi & Poo, 1999), retinotectal synapses of xenopus (Zhang, Tao, Holt, Harris, & Poo, 1998), mammalian cortex (Markram et al., 1997; Feldman, 2000; Froemke & Dan, 2002), and others (Magee & Johnston, 1997; Bell, Han, Sugawara, & Grant, 1997; Debanne, Gahwiler, & Thompson, 1998). A recent comparative review of the different spike-dependent learning rules can be found in Roberts and Bell (2002).

This new type of plasticity, sometimes termed spike-timing-dependent plasticity (STDP), has been studied in various theoretical frameworks, and some of its computational properties have been characterized. Importantly, under certain assumptions about the relative magnitude of synaptic potentiation and depression, STDP embodies an inherent competition between incoming inputs, and thus results in normalization of the total incoming synaptic strength (Kempter, Gerstner, & van Hemmen, 1999, 2001), and maintains the irregularity of neuronal spike trains (Abbott & Song, 1999; Song, Miller, & Abbott, 2000). STDP may also play an important role in sequence learning (Mehta, Quirk, & Wilson, 1999; Rao & Sejnowski, 2000) and lead to the emergence of synchronous subpopulation firing in recurrent networks (Horn, Levy, Meilijson, & Ruppin, 2000). The dynamics of synaptic efficacies under the operation of STDP strongly depends on whether STDP is implemented additively (independent of the baseline synaptic value) or multiplicatively (where the change is proportional to the synaptic efficacy) (Rubin, Lee, & Sompolinsky, 2001; Aharonov, Guttig, & Sompolinsky, 2001). Additive learning leads to strong competition between synapses and is intrinsically unstable, while supra-additive learning may involve inherent stabilization and allow a wide range of synaptic values.[1]

This article takes a different approach to the investigation of temporal aspects of learning. Whereas the common approach is to model STDP and study its properties, we start by analytically deriving spike-dependent learning from first principles in a simple rate model, show how it could be approximated by a biologically feasible rule, and compare the latter with the experimentally observed STDP. This comparison provides computational insights into the temporal structure of the experimentally observed STDP and shows that under certain biological constraints, STDP approximates the performance of the optimal rule.

To this end, we apply the generic principle of information maximization (Linsker, 1988), which states that the goal of a neural network's learning procedure is to maximize the mutual information between its output and

[1] Additive and multiplicative learning may be viewed as two instances of a continuum of the form $\Delta W = \eta W^{\alpha}$, where $\alpha = 0$ corresponds to additive learning, $\alpha = 1$ corresponds to multiplicative one, and any $\alpha > 0$ is termed supra additive.

input. This principle, known as Infomax, was applied in Linsker (1992) to a noisy linear network of real-valued ("rate") input neurons with a multivariate gaussian distribution. It yielded a two-phased learning rule: a Hebbian learning phase when a signal is presented to the network and anti-Hebbian learning when only noise is presented. This article extends the Infomax principle and derives a learning rule that maximizes information about some relevant components of the inputs, not about its complete distribution.

A naive interpretation of the Infomax principle asserts that information about the inputs of the network should be maximized. However, this criterion may be inappropriate in the context of information processing in the brain, which is not targeted to reproduce the sensory inputs, but rather to extract their behaviorally relevant components. We therefore consider a variant of the Infomax principle in which information should be maximized about the identity of the input, allowing for the extraction of relevant information.

The article is organized as follows. Section 2 describes the optimization task and the model. Section 3 derives a gradient-ascent learning rule that maximizes input-output relevant mutual information. Section 4 shows how this rule can be approximated by a biologically feasible learning rule and discusses its properties. Section 5 studies possible extensions for the case of a limited supervised signal, and section 6 investigates learning procedures for the synaptic transfer function. The importance of these results and insights into experimentally observed plasticity are discussed in section 7.

2 The Model

We study a generic learning task in a network with N input neurons $S_1, \ldots, S_N$ firing spike trains and a single output (target) neuron Y. At any point in time, the target neuron integrates its inputs with some continuous temporal filter F due to voltage attenuation and a synaptic transfer function,

$$Y(t) = \sum_{i=1}^{N} W_i X_i(t) \tag{2.1}$$

$$X_i(t) \equiv \int_{-\infty}^{t} F_\tau(t - t') S_i(t')\, dt' \quad ; \quad \int_0^\infty F_\tau(t)\, dt = 1,$$

where W_i is the amplitude of synaptic efficacy between the ith input neuron and the target neuron, $S_i(t) = \sum_{t_{spike}} \delta(t - t_{spike})$ and is the ith spike train. The filter F may be used to consider general synaptic transfer functions and voltage decay effects. For example, voltage attenuation due to leak currents in a passive membrane is realized as an exponential filter $F_\tau(t) = \frac{1}{\tau}\exp(-t/\tau), \forall t > 0$ with τ being the membrane time constant.

The learning goal in our model is to set the synaptic weights vector W such that $M + 1$ uncorrelated patterns of input activity $\xi^\eta(t)$ ($\eta = 0..M$) may be discriminated. More formally, the goal is to maximize the mutual

information between the output Y and the identity of the input pattern η. Each pattern $\xi^\eta(t)$ determines the firing rates of all N input neurons as a function of time. The actual input $S_i(t)$ is a stochastic realization of $\xi_i(t)$. The input patterns are presented for short periods of length T. At each period, a pattern ξ^η is randomly chosen for presentation with probability p_η. Most of the patterns are rare ($\sum_{\eta=1}^{M} p_\eta \ll 1$), but ξ^0 is abundant and may be thought of as a background noisy pattern. This assumption is based on the common scenario of a neuron (e.g., a face cell in IT cortex) that most of the time is exposed to irrelevant stimuli and is rarely presented with stimuli it is tuned to.

Two aspects of this model should be emphasized. First, unlike Linsker (1992), information is maximized not about the input values but about the identity of the presented pattern. This follows the idea that the goal of a neural system is not to reproduce the representation of the inputs but to extract relevant information. Second, the input patterns determine the modulating rates that underlie the input spike trains. These rates are not observable, and therefore any learning procedure must depend on the observable input spikes that realize the underlying rates. Therefore, the fact that a learning rule depends on spikes does not necessarily mean that information is coded in spikes instead of in the underlying rates.

3 Mutual Information Maximization

The goal of this section is to derive a learning algorithm that maximizes the input-output mutual information of the above model by changing the synaptic weights. We first focus on the modification of the synaptic magnitudes W, while keeping the temporal filter F fixed. Learning the optimal temporal filter is discussed in section 6.3.

3.1 Deriving a Gradient-Ascent Learning Rule. Let us focus on a single presentation period and look at the value of $Y = Y(T)$ at the end of this period. Assuming that the temporal filter is mostly concentrated in a time period of length T and omitting the notation of t, we obtain from equation 2.1

$$Y = \sum_{i=1}^{N} W_i X_i \quad ; \quad X_i = \int_{-T}^{0} F_\tau(0 - t') S_i(t')\, dt'. \tag{3.1}$$

The mutual information (Shannon, 1948; Cover & Thomas, 1991; see also Linsker, 1992) between the output and pattern identity in this network is defined by

$$I(Y;\eta) = h(Y) - h(Y \mid \eta) \quad ; \quad h(Y) = -\int_{-\infty}^{\infty} f(y) \log(f(y))\, dy, \tag{3.2}$$

where the $h(Y)$ is the differential entropy (Cover & Thomas, 1991) of the distribution of Y, and $h(Y \mid \eta)$ is the differential entropy of the Y distribution

given that the network is presented with a known input pattern η. $f(Y)$ is the probability density function (p.d.f.) of Y. This mutual information measures how easy it is to decide which input pattern η is presented to the network by observing the network's output Y.

Consider the case where the number of input neurons is large and neurons fire independently. Under these conditions, according to the central limit theorem, when the target neuron is presented with the pattern ξ^η, its membrane voltage Y is a weighted sum of many uncorrelated inputs and is thus normally distributed $f(Y \mid \eta) = N(\mu_\eta, {\sigma_\eta}^2)$ with mean $\mu_\eta = \langle WX^\eta \rangle$ and variance ${\sigma_\eta}^2 = \langle (WX^\eta)(WX^\eta)^T \rangle - \langle WX^\eta \rangle^2$. The brackets denote averaging over the possible realizations of the inputs X^η when the network is presented with the pattern ξ^η. Consider now the more complex case where the input neurons fire in a correlated manner according to the covariance matrix $C_{ij} = \mathrm{Cov}(X_i, X_j)$. Y is now a sum of correlated inputs, and its distribution converges to the normal distribution at a rate that is determined not by the number of input neurons but by the number of effectively independent variables. In particular, according to the theory of large samples (Lehman, 1998; Ferguson, 1996), if each neuron is correlated with at most m other neurons, their sum converges to the normal distribution, with a rate that is at the worst case slower by a factor of $2m$ (Ferguson, 1996). Consequently, even with the high levels of synchrony reported in some experiments performed in the mammalian cortex (e.g., Singer & Gray, 1995), the number of effectively independent inputs is still very large. This suggests that the input-dependent weighted sum of synaptic inputs can be well approximated by a normal distribution, with parameters that depend on the activity pattern. Formally, $f(Y \mid \eta) \approx N(\mu_\eta, {\sigma_\eta}^2)$ with mean $\mu_\eta = \langle WX^\eta \rangle$ and variance ${\sigma_\eta}^2 = \langle (WX^\eta)(WX^\eta)^T \rangle - \langle WX^\eta \rangle^2$. Using the expression for the entropy of a gaussian ($h(X) = \frac{1}{2}\log(2\pi e\sigma^2)$ when $X \sim N(\mu, \sigma^2)$), we obtain

$$h(Y \mid \eta) = \sum_{\eta=0}^{M} p_\eta h(Y \mid \xi^\eta) = \sum_{\eta=0}^{M} p_\eta \frac{1}{2}\log(2\pi e{\sigma_\eta}^2). \tag{3.3}$$

To further calculate the mutual information $I(Y;\eta)$, we turn to calculate the second term of equation 3.2, the entropy $h(Y)$. For this purpose, we note that $f(Y)$ is a mixture of gaussians (each resulting from the presentation of a single input pattern) $f(Y) = \sum_\eta f(Y \mid \eta)p_\eta$, and use the assumption that $\sum_{\eta=1}^{M} p_\eta$ is small with respect to p_0. The calculation of $h(Y)$ is still difficult, but fortunately its derivative with respect to W_i is tractable. Using a first-order Taylor approximation, we obtain (see appendix A)

$$\begin{aligned}\frac{\partial I(Y;\eta)}{\partial W_i} = &+\sum_{\eta=1}^{M} p_\eta(\mathrm{Cov}(Y, X_i^\eta)K_\eta^1 + E(X_i^\eta)K_\eta^2) \\ &- \sum_{\eta=1}^{M} p_\eta(\mathrm{Cov}(Y, X_i^0)K_\eta^0 + E(X_i^0)K_\eta^2),\end{aligned}$$

with
$$K_\eta^0 \equiv \frac{(\mu_\eta - \mu_0)^2}{{\sigma_0}^4} + \frac{{\sigma_\eta}^2 - {\sigma_0}^2}{{\sigma_0}^4}$$
$$K_\eta^1 \equiv \frac{1}{{\sigma_0}^2} - \frac{1}{{\sigma_\eta}^2}$$
$$K_\eta^2 \equiv \frac{\mu_\eta - \mu_0}{{\sigma_0}^2}, \tag{3.4}$$

where $E(X_i^\eta)$ is the expected value of X_i^η as averaged over presentations of the ξ^η pattern.

The derived gradient may be used for a gradient-ascent learning rule, by repeatedly calculating the distribution moments μ_η, σ_η that depend on W and updating the weights according to

$$\begin{aligned}\Delta W_i &= \lambda \frac{\partial}{\partial W_i} I(Y;\eta) \\ &= +\lambda \sum_{\eta=1}^{M} p_\eta (\mathrm{Cov}(Y, X_i^\eta) K_\eta^1 + E(X_i^\eta) K_\eta^2) \\ &\quad - \lambda \sum_{\eta=1}^{M} p_\eta (\mathrm{Cov}(Y, X_i^0) K_\eta^0 + E(X_i^0) K_\eta^2).\end{aligned} \tag{3.5}$$

Since this learning rule climbs along the gradient, it is guaranteed to converge to a local maximum of the mutual information. To demonstrate the operation of this learning rule, Figure 1 (left) plots the mutual information during the operation of the learning procedure as a function of time, showing that the network indeed reaches a (possibly local) maximum of the mutual information. When this simulation is repeated for different random initializations of the synaptic weights, the rise time of the mutual information varies across runs, but all runs reach the same mutual information end level.

The above analytical derivation requires the background pattern to be highly abundant $\sum_{\eta=1}^{M} p_\eta \ll p_0$. When this assumption does not hold, performance of the learning rule can be studied numerically, to test whether the obtained learning rule is still beneficial. To this end, we repeated our simulations for the worst-case scenario, where two input patterns are presented with $p_0 = p_1 = 0.5$. In this case too, the simulation yielded a monotonic increase of mutual information until it reached a stable plateau, as shown in Figure 1 (right).

3.2 Dynamics Under the Operation of the Learning Rule. In order to understand the dynamics of the weights W under the operation of the gradient-ascent learning rule, consider now the two sums in equation 3.4. The first sum contains terms that depend on the pattern ξ^η through X_i^η , and

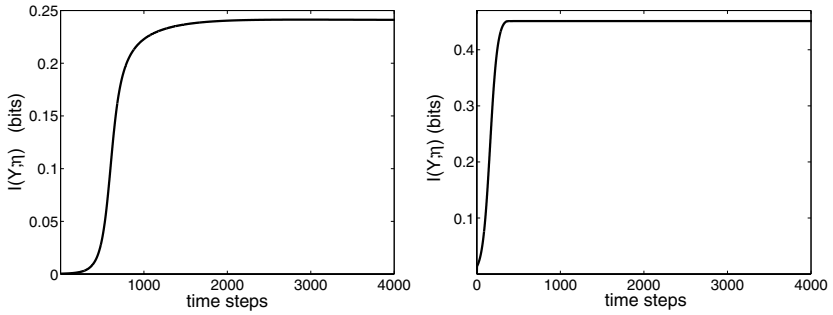

Figure 1: Mutual information during learning with the gradient-ascent learning rule (see equation 3.4). All patterns were constructed by randomly setting 10% of the input neurons to fire Poisson spike trains at 40 Hz, while the remaining input neurons fire at 5 Hz. Poisson spike trains were simulated by discretizing time into 1 millisecond bins. Simulation parameters: $\lambda = 1, N = 1000$. (Left) 100 memory patterns with $p_0 = 0.9$, $p_\eta = 0.001$ for $\eta > 0$. (Right) Two memory patterns with $p_0 = p_\eta = 0.5$.

the second sum contains terms that depend on ξ^0 through X_i^0. When the K terms are positive, these two sums are positive. Thus, the weights that correspond to strong inputs of the background pattern ξ^0 are weakened, while the weights that correspond to strong foreground inputs are strengthened. This result is similar to the case of real-valued inputs investigated in Linsker (1992) where an information gradient-ascent rule alternates between Hebbian and anti-Hebbian learning. The main effect of this process is an increase of the differences $|\mu_\eta - \mu_0|$, while constraining the variances σ_η, thus providing better discrimination of the rare patterns ξ^η from the abundant ξ^0 pattern.

This combination of Hebbian and anti-Hebbian learning characterizes the mutual information target function we use and is not necessarily obtained with other target functions. For example, a maximum output variance criterion, whose optimal learning rule is the traditional correlation Hebbian rule, involves a global decay term that is not input specific. Appendix D compares these two target functions.

The intuition behind the increase in $|\mu_\eta - \mu_0|$ stems from the information maximization criterion: aiming to split the outputs as sparsely as possible, the pattern with the highest prior probability is assigned an opposite sign compared to the other patterns. Figure 2 depicts the distribution of output values during learning as traced in simulation, showing how the output values split into two distinct gaussian-like bumps: one corresponding to the presentation of the background pattern ξ^0 and the other to the presentation of the rest of the patterns. Although in this figure, all rare patterns share a

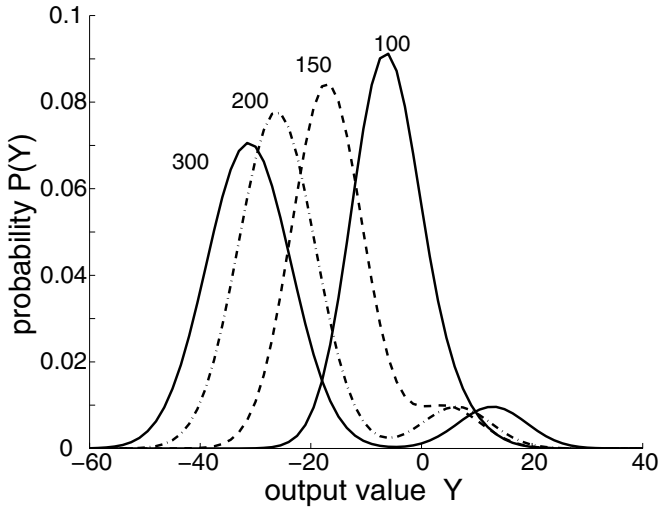

Figure 2: Distribution of the output values Y after 100, 150, 200, and 300 learning steps. Outputs segregate into two distinct bumps: one that corresponds to the presentation of the ξ^0 pattern and a second that corresponds to the rest of the patterns. Simulation parameters are as in Figure 1 (left).

single bump, in the general case, different foreground patterns may have different bumps.

The mutual information manifold is a complex and nonconvex function of the weight vector that does not have a single maximum. In fact, any extremum point has a mirror extermum point with all weights having the opposite signs (see appendix B). A further analytical characterization of the information manifold is difficult, and therefore in order to characterize the steady-state structure of the system, we repeated simulations for different random initializations of synaptic weights,[2] and compared the synaptic efficacies at the steady state. Interestingly, these simulations found only two steady-state solutions for a given set of input patterns, and these were always a pair of mirror solutions (see appendix B). This suggests that the landscape of mutual information as a function of weights values contains only two large basins of attractions, into which the optimization process is drawn.

[2] Simulation was repeated with the same set of memory patterns for 100 times under two conditions. First, weights were initialized as uniform random numbers between 0 and 1. Second, weights were initialized in $(-1, 1)$. This was repeated for 10 different sets of memory patterns.

4 Toward Biological Feasibility

4.1 Deriving an Approximated Rule. With the goal of obtaining a biologically feasible spike-dependent learning rule that maximizes mutual information, we now turn to approximate the analytical learning rule derived above by a rule that can be implemented in biology. We follow four steps to modify the analytically derived rule into a more biologically plausible one.

First, biological synapses are limited to either excitatory or inhibitory regimes. The optimum has two mirror solutions, and for simplicity, we limit the weights W to positive values (but see our discussion on inhibitory synapses in section 7).

Second, we replace the terms $\{K_\eta^1, K_\eta^0, K_\eta^2\}$, which are global functions of W and ξ^η, with constants $\{\lambda^1, \lambda^0, \lambda^2\}$. These constants are optimized but remain fixed during the simulation and therefore cannot provide equal performance as the continuously changing K terms. However, if these constants are close to the K values, the changes in the weight vector have a positive projection on the gradient, and this approximation is sufficient for approaching the optimum. This suggests that high levels of information can be obtained for a large range of values of these constants, allowing cells to operate in a near-optimal regime. To test this idea, we measured the steady-state information for different constants' values. Figure 3 shows that the steady-state level of information is not critically sensitive to the λ values, since high information levels are obtained for a large range of their values.

Third, we implement an on-line learning mode instead of a batch one by replacing explicit summation over patterns with stochastic averaging over the presented patterns. Because summation is performed over the rare patterns alone (see equation 3.4), stochastic averaging is naturally implemented by restricting learning to the presentation of foreground (rare) patterns. Appendix C discusses alternative learning rules in which learning is triggered by the presentation of both the background and foreground patterns, showing that these rules are not robust to fluctuations in p_η and result in lower mutual information. We thus restrict learning to the presentation of foreground patterns only,[3] which yields the following learning rule:

$$\Delta W_i = +(\lambda^1 \operatorname{Cov}(Y, X_i^\eta) + \lambda^2 E[X_i^\eta]) - (\lambda^0 \operatorname{Cov}(Y, X_i^0) + \lambda^2 E[X_i^0])$$
$$\text{when } \xi^\eta \text{ is presented}, \quad \text{for } \eta = 1..M \text{ only.} \tag{4.1}$$

Fourth, we replace the explicit dependence on average quantities $E(X)$ and $\operatorname{Cov}(Y, X)$ by stochastic averaging over spikes. This yields a spike-dependent learning rule. In the case of inhomogeneous Poisson spike trains

[3] This requires a weak form of supervision and is discussed in section 5.

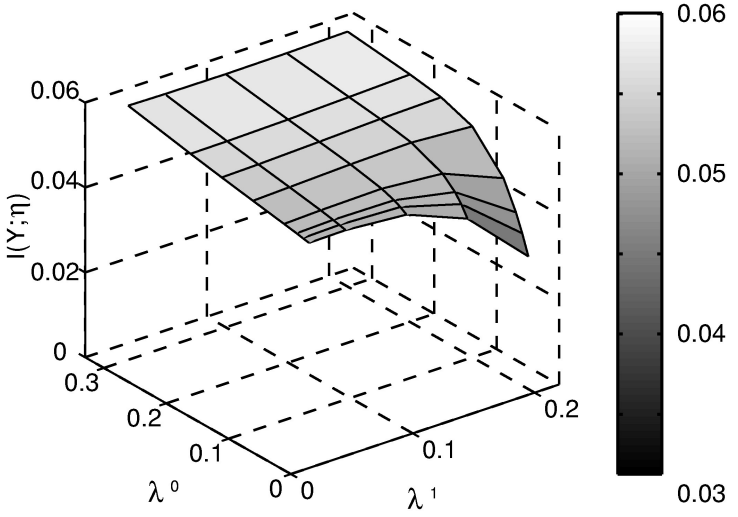

Figure 3: Mutual information as a function of learning parameters λ^1 and λ^0. High information levels are obtained for a wide range of parameter values. Simulation parameters $\lambda^2 = 0.1$. $\lambda = 1.0, M = 20, N = 1000, p_0 = 0.9, p_\eta = 0.001$, $T = 20$ ms.

where input neurons fire independently,[4] the expectation terms obey $E(X_i) = \int_{-T}^{0} F(t')E(S_i(t'))\,dt'$, and the covariance terms obey (Kingman, 1993)

$$\begin{aligned}\mathrm{Cov}(Y, X_i) &= \sum_j W_j \,\mathrm{Cov}(X_j, X_i) = W_i \,\mathrm{Var}\left(\int_{-T}^{0} F(t')S_i(t')\,dt'\right) \\ &= W_i \int_{-T}^{0} \mathrm{Var}(F(t')E(S_i(t')))\,dt' \\ &= W_i \int_{-T}^{0} F^2(t')E(S_i(t'))\,dt'. \qquad (4.2)\end{aligned}$$

The expectations $E(X_i^\eta)$ for $\eta > 0$ may simply be estimated by weighted averages of the observed spikes S_i^η that precede the learning moment. Estimating $E(X_i^0)$ is more difficult because, as discussed in appendix C, learning should be triggered solely by the rare patterns. Namely, ξ^0 spikes should have an effect only when a rare pattern ξ^η is presented. A reasonable approximation can be obtained using the fact that ξ^0 is highly frequent, by

[4] Relaxation of these assumptions and their effect on the learning rule are discussed in section 6.

averaging over all spikes in long time periods. In this case, the highest reliability for estimating ξ^0 spikes is obtained if spikes are weighted uniformly everywhere except during the presentation of ξ^η. To see this, consider a stationary Poisson neuron that fires at a rate $E[S(t)] = p$, which has to be estimated by a weighted average $\hat{p} = \int_a^b f(t)S(t)\,dt$ where $\int_a^b f(t) = 1$. The variance of the estimator is $\int_a^b f^2(t)\,dt\,\mathrm{Var}[S(t)]$, which is minimized for $f(t) = \text{const}$.

This optimal function extends infinitely in time and therefore cannot be realized in biological cells, since these have a limited memory span. When the memory span of a neuron is limited by L, that is, its weighting function is constrained to be zero for $|t| > L$, the optimal weighting function is $f(t) = \frac{1}{(2L-T)}$ for all $t \in (-L, -T) \cup (0, L)$ and zero otherwise. It should be stressed that different constraints on the memory span of the cell, reflecting limitations of the biological hardware, may lead to weighting functions that are different from a uniform f. The exact form of the weighting function for the background spikes depends on such constraints and therefore cannot be predicted by our analysis here.

Formally, the following rule is activated only when one of the foreground patterns (ξ^η, $\eta = 1..M$) is presented:

$$\Delta W_i = +\int_{-T}^{0} [\lambda^1 W_i F^2(t) + \lambda^2 F(t)] S_i(t)\,dt - \int_{-\infty}^{\infty} [\lambda^0 W_i f^2(t) + \lambda^2 f(t)] S_i(t)\,dt. \tag{4.3}$$

This learning rule uses a weak form of supervision signal that activates learning when one of the foreground patterns is presented but does not require an explicit error signal. Its properties are discussed in section 5.

4.2 Comparing Performance. In the previous section, we derived a biologically implementable spike-dependent learning rule that approximates an information maximization learning rule. But how good are these approximations? Does learning with the biologically feasible learning rule increase mutual information, and to what level? The curves in Figure 4 compare the mutual information of the learning rule of equation 3.4 with that of equation 4.3, as traced in a simulation of the learning process. Apparently, the approximated learning rule achieves fairly good performance compared to the optimal rule, and most of the reduction in performance is due to limiting the weights to positive values.

4.3 Interpreting the Learning Rule Structure. In order to obtain intuition into the approximated learning rule, we now examine the components of equation 4.3 and demonstrate them pictorially in Figure 5.

First, synaptic potentiation in equation 4.3 is temporally weighted in a manner that is determined by the same filter F that the neuron applies over

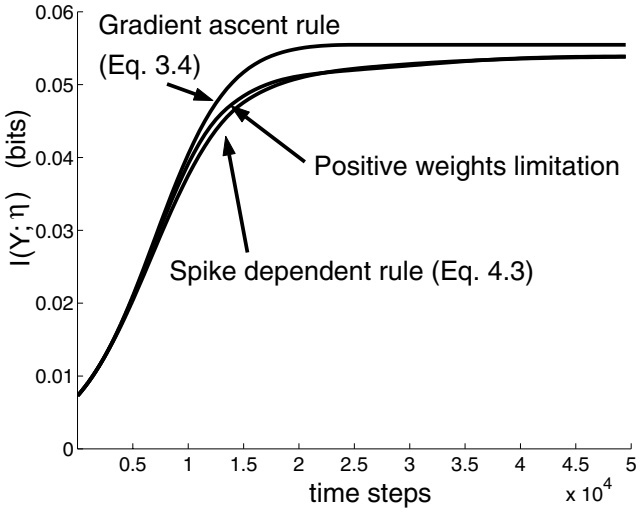

Figure 4: Comparing optimal (see equation 3.4) and approximated (see equation 4.3) learning rules. Poisson spike trains were simulated by discretizing time into 1 millisecond bins. For each pattern, 10% of the input neurons were set to fire at 40 Hz, while the rest fire at 5 Hz. A lower bound of zero was imposed on all weights. Simulation parameters: $\lambda = 1.0, M = 20, N = 2000, \lambda^1 = 0.15, \lambda^0 = 0.05$, $\lambda^2 = 0.1$, $p_0 = 0.9$, $p_\eta = 0.001$, $T = 20$ ms.

its inputs. The learning curve involves an average of F (i.e., $\int^t F(t-t')S(t')\,dt'$) and F^2 (i.e., $\int^t (F(t-t'))^2 S(t')\,dt'$). To demonstrate the shape of the resulting learning rule, we choose the filter F to be an exponential filter $F(t) = \frac{1}{\tau}\exp(-t/\tau)$, which corresponds to the voltage decay in a passive membrane with time constant τ. In this case, the squared filter F^2 is also an exponential curve but with time constant $\tau/2$. Figure 5A presents the weight-dependent (filtered by F^2) and weight-independent (filtered by F) potentiation components for the exponential filter. The relative weighting of these two components (determined by the K terms) was numerically estimated by simulating the optimal rule of equation 3.4 and was found to be on the same order of magnitude (demonstrated in Figure 6). This suggests that learning should contain weight-dependent and -independent components, in agreement with investigations of additive versus multiplicative STDP implementations (Rubin et al., 2001).

Second, equations 3.4 and 4.3 suggest that synaptic depression is activated when a foreground pattern is presented but serves to learn the under-

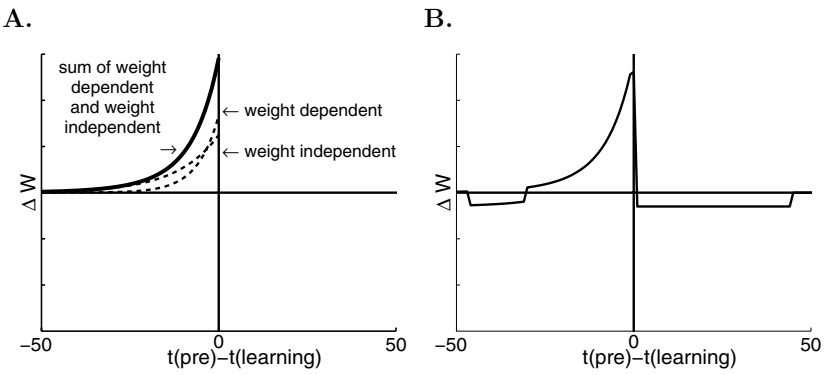

Figure 5: The temporal characteristics of the approximated learning rule of equation 4.3. The changes in synaptic weights ΔW as a function of the temporal difference between the learning time t and the input spike time t_{spike}. (A) The potentiation curve (solid line) is the sum of two exponents with constants τ and $\frac{1}{2}\tau$ (dashed lines). (B) A combination of the potentiation curve of the subplot A and a uniform depression curve with $L = 45$, $T = 30$.

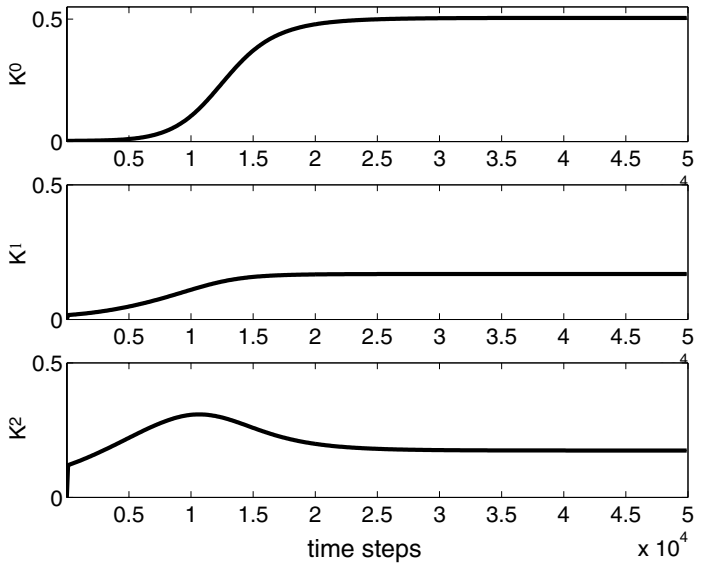

Figure 6: Traces of the values of K_1^0 (up), K_1^1 (middle), and K_1^2 (bottom), as are numerically calculated during the simulation of the gradient-ascent rule, showing that all three values are of the same order of magnitude. Simulation parameters as in Figure 5.

lying structure of the background activity. As discussed in section 4.2, the optimal depression learning curve extends infinitely in time, but this requires long memory that is not biologically feasible. Using a limited memory span of size L allows combining it with the potentiation curve into a single feasible rule. A learning rule that combines both potentiation and depression is presented in Figure 5B.

A major difference between the spike-triggered rule of equation 4.3 and the experimentally observed STDP is the learning trigger. In equation 4.3, learning is triggered by an external learning signal that corresponds to the presentation of rare patterns, while in the experimentally observed rule, it is triggered by the postsynaptic spike. The possible role of the postsynaptic spike is discussed in the following section.

5 Limited Supervised Learning Signal

We have thus far considered a learning scenario in which the system was allowed to use an explicit and continuous external learning signal whenever one of the rare patterns was presented. To take a further step toward biological plausibility, we turn to investigate a more plausible implementation of the supervised learning signal. Such a signal may be noisy, unreliable, and probably available for only a limited period. In biological neural networks, the output spike of an excitatory neuron usually signals the presence of an interesting (or a rare) pattern. The postsynaptic spike is therefore a natural candidate to serve as a learning signal in our model, in which learning should be triggered by the presentation of the rare patterns.[5]

The derivation in the previous sections aimed to maximize the information $I(Y;\eta)$ in the membrane voltage Y. If, however, a neuron is forced to transmit a binary signal, the strategy that minimizes information loss when translating the continuous valued Y into a binary value is to spike when one of the rare foreground patterns $\eta \geq 1$ is presented.[6] We thus turn to investigate whether the postsynaptic spike (signaling the presence of interesting input patterns) may be used as a learning signal. This yields a learning procedure identical to equation 4.3, except this time, learning is triggered by postsynaptic spikes. The resulting learning rule is somewhat similar to

[5] The postsynaptic spike is not the only candidate for providing a learning signal; high but subthreshold depolarization may also testify for the presentation of a foreground pattern. Indeed, there is evidence that learning is triggered by high subthreshold depolarization (Debanne et al., 1994).

[6] Note, however, that as proved in appendix B, two mirror solutions can be achieved by our analytically derived rule. The first one potentiates synapses that correspond to strong inputs and strongly depolarizes the target cell when rare patterns are presented. In this case, output spikes signal the presence of rare patterns. Under the mirror solution, the spikes of target neuron signal the background pattern, leading to higher mean firing rates. This high activity solution is considered in the literature as less desirable in terms of energy consumption.

previous models of the experimentally observed STDP (Kempter et al., 1999; Song et al., 2000; Rubin et al., 2001; Kempter et al., 2001), although we keep the form of the learning rule derived earlier.

To this end, we have simulated an integrate-and-fire neuron that receives additional strong depolarizatory input when foreground patterns are presented. This input usually causes the neuron to fire several milliseconds after the beginning of pattern presentation. As before, the mutual information was traced along time. The main difference between this simulation and our previous ones was that the target neuron was sometimes spontaneously firing even when presented with the background pattern and that the time of the postsynaptic spike was not locked to the end of stimulus presentation. In addition, the integrate-and-fire model embodies a different behavior of the membrane leak. While the model of equation 2.1 accounts for local leaks that influence the voltage from a single input, the membrane leak in the integrate-and-fire model depends on the membrane voltage, which in turn depends on all converging input to the target neuron.

The external signal was supplied for a limited learning period. After this period, learning may still occur, but the postsynaptic spikes that trigger it are no longer supervised but controlled by inputs only. If the network has learned to represent the presented input pattern faithfully, the target neuron would faithfully spike on foreground pattern presentation and serve to keep the network at a stable and informative state. Figure 7 traces the mutual information along and after learning, showing that the system stays in the informative state even after the external supervised learning signal is turned off.

6 Extending the Model

Our analysis derived a biologically feasible learning rule that is near optimal when the statistics of the input spike trains obeys several limitations. This section considers a wider family of input spike train statistics and discusses the changes in the learning rule needed in order to take this statistics into account in an optimal manner. We present three types of extensions: correlated inputs, non-Poisson spike trains, and learnable synaptic transfer functions. Although all three aspects may occur in conjunction, we consider them separately here for the simplicity of exposition.

6.1 Correlated Inputs. Section 4 derived a learning rule for uncorrelated input spike trains. We now turn to consider input cross-correlations and show how these correlations change the covariance term of equation 4.1, but preserve the general structure of the spike-dependent learning rule.

Consider the case where input neurons do not fire independently but in a correlated manner such that their filtered spike trains $X_1^\eta, \ldots, X_N^\eta$ have the covariance matrix $C_{ij}^\eta = \mathrm{Cov}(X_i^\eta, X_j^\eta)$. If these correlations obey the

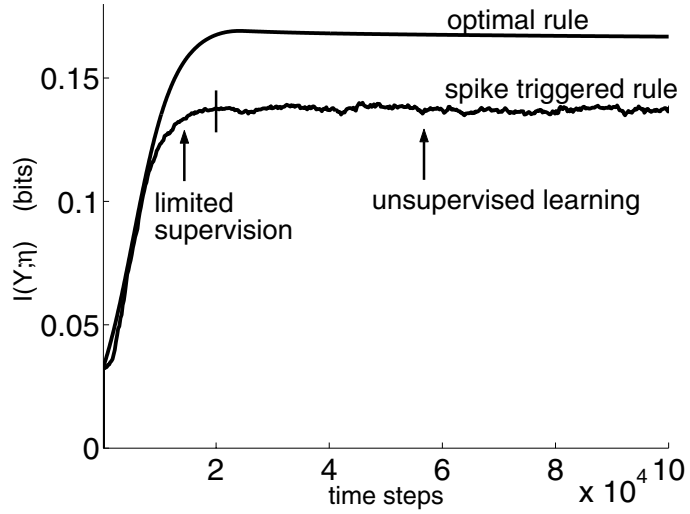

Figure 7: Mutual information along learning with the optimal learning rule (see equation 3.4) and learning with a limited supervised rule in a simulation of an integrate-and-fire neuron (threshold = 55 mV, resting potential = 70 mV, membrane time constant = 10 milliseconds). Learning occurred whenever the target neuron fired a spike. On the first 20,000 simulation steps, a depolarizing input of 5 millivolts was injected into the neuron if one of the foreground patterns was presented (which occurred with probability 5%). This input sometimes resulted in a spike that in turn initiated learning. Following this learning phase, the neuron continues to preserve the discrimination between background and foreground patterns even in the absence of the supervision signal.

conditions discussed in section 3.1, then the weighted sum $Y = \sum_i W_i X_i^\eta$ is normally distributed $f(Y \mid \eta) = N(\mu_\eta, {\sigma_\eta}^2)$, with mean $\mu_\eta = \langle WX^\eta \rangle$ and variance ${\sigma_\eta}^2 = WC_{ij}^\eta W^T - {\mu_\eta}^2$. In this case, the covariance term in equation 4.2 changes into

$$\begin{aligned}
\mathrm{Cov}(Y, X_i) &= \sum_j W_j \,\mathrm{Cov}(X_j, X_i) \\
&= \sum_j W_j \,\mathrm{Cov}\left(\int_{-T}^{0} F(t') S_j(t')\, dt', \int_{-T}^{0} F(t) S_i(t)\, dt \right) \\
&= \sum_j W_j \int_{-T}^{0} F^2(t)\, \mathrm{Cov}[S_j(t), S_i(t)]\, dt. \qquad (6.1)
\end{aligned}$$

In order to take into account these correlations, the learning rule of equation 4.1 becomes

$$\Delta W_i = + \left(\lambda^1 \sum_j W_j \int_{-T}^{0} F^2(t) \operatorname{Cov}[S_j^\eta(t), S_i^\eta(t)]\, dt + \lambda^2 E[X_i^\eta] \right)$$
$$- \left(\lambda^0 \sum_j W_j \int_{-T}^{0} F^2(t) \operatorname{Cov}[S_j^0(t), S_i^0(t)]\, dt + \lambda^2 E[X_i^0] \right)$$
$$\text{when } \xi^\eta \text{ is presented, for } \eta = 1..M \text{ only.} \tag{6.2}$$

When presented with the pattern η, this learning rule increases more strongly the synapses from neurons that are positively correlated with other input neurons. More interesting, while the first term of equation 6.2 always has a positive component $W_i \int_{-T}^{0} F^2(t) \operatorname{Var}(S_i(t))\, dt$, it may be counterbalanced with negative components $\operatorname{Cov}(S_j(t), S_i(t))$ for $j \neq i$ and may operate to weaken the synaptic strength of neurons that are active in foreground patterns. These effects are not unique to Infomax learning, but also characterize maximum variance learning (principal components analysis) and its implementations in traditional covariance-based learning (see appendix D).

As with any other learning rule that takes into account input correlations, implementing the learning rule of equation 6.2 in biology requires the target neuron to measure coactivities of its synaptic inputs. The physiological mechanisms of such correlation-dependent synaptic changes are still largely unknown, although they may in principle be realized to some extent through nonlinear processing in the dendrites (Segev & London, 2000).

6.2 Non-Poisson Spike Trains. The previous section discussed correlations between different inputs. Here, we consider temporal correlations in the spike trains of an individual input neuron. Let $S_i(t)$ be a point process that is no longer Poisson but has covariance $\operatorname{Cov}[S_i(t), S_i(t')]$ for all $t \in [-T, 0]$. Under these conditions, the covariance term in equation 4.2 changes into

$$\operatorname{Cov}(Y, X_i) = W_i \operatorname{Var}(X_i)$$
$$= \int_{-T}^{0} \int_{-T}^{0} F(t) \operatorname{Cov}[S_i(t), S_i(t')] F(t')\, dt'\, dt. \tag{6.3}$$

As in the case of cross-correlations described in section 6.1, the positive term $\operatorname{Cov}[S_i(t), S_i(t)] = \operatorname{Var}(S_i(t))$ may now be counterbalanced by negative covariance terms $\operatorname{Cov}[S_i(t), S_i(t')]$ for $t \neq t'$ if spikes are anticorrelated.

As an example, consider a stationary input spike train with $E(S(t)) = p$ and autocorrelations $\mathrm{Cov}[S(t), S(t')]$. Focus on two points in time, t and $t' = t - \Delta t$, and denote $p(S(t') = 1 \mid S(t) = 1) = p'_{\Delta t}$. When $p'_{\Delta t} < p$, the autocorrelation is negative, as in the case where the two spikes are within the neuronal refractory period. With this notation, the covariance can be written as $\mathrm{Cov}[S(t'), S(t)] = pp'_{\Delta t} - p^2 = pp'_{\Delta t} g(\Delta t)$, with $g(\Delta t) = 1 - p/p'_{\Delta t}$. The spike-triggered learning rule thus includes two components: synaptic changes induced by single-input spikes as described in section 3 and synaptic changes induced by pairs of spikes. Spike pairs occur with probability $pp'_{\Delta t}$ and effect the synapse with a magnitude $g(\Delta t)$. The pairwise component of the learning rule can therefore be implemented by changing the synapse by $\Delta W_i = F(t)g(t - t')F(t')$ whenever a pair of spikes occurs at times t and t'. The function g depends on the inputs' statistics (its firing rate and autocorrelations) and is used by the learning rule. Optimized learning rules can therefore be achieved when g matches the baseline statistics of the input. In the case of negative autocorrelations, the covariance term and the function $g(\Delta t)$ are negative; thus, the first input spike may be regarded as inhibiting the synaptic change induced by the second input spike.

What is the relation between these autocorrelation-sensitive rules and STDP observed in neural tissues? Do neural systems estimate and use autocorrelations in the inputs to optimize their synaptic learning rule? Interestingly, while early work focused on measuring changes in excitatory postsynaptic potential (EPSP) as a function of time differences between single pre- and postsynaptic spikes, the effects of complex and natural spike trains were investigated recently by Froemke and Dan (2002). The structure of patterns of presynaptic spikes was found to have pronounced effects on STDP, suggesting that synaptic plasticity is indeed sensitive to autocorrelations and input structures. Interestingly, an input spike that preceded a pair of input-output spikes that caused an EPSP change was found to suppress the magnitude of STDP. This is in agreement with the above analysis for the case of negative autocorrelation at short time differences.

Sensitivity to complex patterns was observed not only for presynaptic spikes (as in the above analysis) but also for output spikes patterns. An information maximization study of these effects therefore requires more detailed assumptions on the nature of coding with output spikes and exceeds the scope of this article.

6.3 Learning the Synaptic Transfer Function. The above analysis focused on changes in the magnitude of synaptic efficacy. As our model includes the shape of the synaptic transfer function, we can extend the above analysis to derive learning rules that change not only the amplitude of the synaptic value W_i but also the shape of the synaptic transfer function $F_i(t)$. This is achieved by differentiating the mutual information $I(Y; \eta)$ with

respect to $F_i(t)$ in a similar manner to the derivation of $\frac{\partial I(Y;\eta)}{\partial W_i}$ and yields

$$\begin{aligned}\mathrm{Cov}(Y, S_i(t)) &= \mathrm{Cov}\left(\sum_{j=1}^{N} W_j \int F_j(0-t')S_j(t'), S_j(t)\right) \\ &= \mathrm{Cov}(W_i F_i(0-t)S_i(t), S_i(t)) = W_i F_i(0-t)\,\mathrm{Var}(S_i(t)) \\ &\approx W_i F_i(0-t)E(S_i(t)). \qquad (6.4)\end{aligned}$$

Thus, a learning rule that changes the synaptic transfer function contains terms that are proportional to $F_i(t)$. This suggests that in addition to an increase in amplitude, the synaptic transfer function should sharpen because large $F_i(t)$ values would be strengthened more than small $F_i(t)$ values.

In biological neurons, various components affect the shape of the EPSP at the soma. Among these are the dynamics of the receptor, the morphology of the dendrite and spine, and the passive electrical properties of the dendritic tree. Physiological evidence is gathered regarding activity-dependent changes in these components (see Segal & Andersen, 2000), but the experimental results are far from being conclusive.

7 Discussion

The principle of information maximization was originally applied for discriminating gaussian signals from noise (Linsker, 1992). Our analysis extends the Infomax principle to discriminate between spatiotemporal activity patterns. The learning task is thus to maximize the relevant information, which in our model lies in the identity of the presented input, instead of reproducing the inputs. Introducing temporal properties of the target neuron into the analysis (such as its synaptic transfer function) allows us to infer the temporal properties of the learning rule required to maximize this relevant information. Moreover, we show how the analytically derived learning rule is well approximated by a biologically feasible learning rule that operates through spike-dependent plasticity. This rule requires a limited supervisor signal, which activates learning whenever an interesting event occurs but does not require an explicit error signal. The supervision signal is required only for a transient learning period, yielding a network that remains at a stable performance level for long periods. Presumably, such a learning signal may be available in the brain by "novelty detector" or reward prediction circuits (Schultz, Dayan, & Montague, 1997).

What can we learn from the structure of this analytically derived learning rule about spike-timing-dependent plasticity observed experimentally? First, in our analysis, synaptic depression operates to unlearn the statistics of the background activity. It is activated following the presentation of rare patterns, in order to maintain the balance between potentiation and depression, as in traditional Hebbian learning models (Sejnowski, 1977; Dayan &

Willshaw, 1991; Chechik, Meilijson, & Ruppin, 2001). The optimal performance is obtained when the weight of depression is stronger than that of potentiation (see Figures 3 and 6). This constraint is nicely explained by the model of Song et al. (2000), where spike time differences that lead to potentiation occur more often than chance because input spikes contribute to the creation of output spikes. Preventing the divergence of synaptic values requires synaptic potentiation to be balanced by stronger depression or by neuronal-level regulation of synaptic efficacies (Chechik, Horn, & Ruppin, 2002).

Second, the shape of the synaptic potentiation curve is determined by the input-output transfer function of the synapse. This yields direct experimentally testable predictions, suggesting that neurons with fast membrane time constants will exhibit fast synaptic plasticity curves correspondingly. Interestingly, while STDP was often found to be NDMA dependent (Markram et al., 1997; Zhang et al., 1998; Bi & Poo, 1999; Feldman, 2000), the STDP potentiation window lasts only a few tens of milliseconds, while the typical timescale of NMDA receptors extends to hundreds of milliseconds. Our findings suggest a computational reasoning for this apparent mismatch because the major contributors of input current to excitatory cells are AMPA-type channels, whose typical timescales are on the order of a few tens of milliseconds only. The derived learning rules combine both a weight-dependent and a weight-independent component. This is in agreement with experiments that show that the change in excitatory postsynaptic current (EPSC) depends on the initial EPSC amplitude (Bi & Poo, 1999). The correlation between EPSC change and initial EPSC is negative, in that larger initial EPSCs show smaller increases following STDP. This suggests that in this preparation, the weight-dependent depression component is stronger than the weight-dependent potentiation component.

Third, because synaptic depression serves to unlearn the baseline statistics of the inputs, it achieves this goal most effectively for synaptic depression curves that extend in time. The experimental evidence in this regard is mixed: some preparations reveal depression curves that extend to 100 to 200 milliseconds (Debanne et al., 1994; Feldman, 2000), while in others, the time constants of both depression and potentiation are similar and around 40 milliseconds (Zhang et al., 1998; Bi & Poo, 1999; Yao & Dan, 2001). Moreover, the theoretical analysis suggests that depression would be even more effective if it operates when the presynaptic spike occurs long before the postsynaptic one (as in Figure 5B). This effect was indeed observed in some preparations (Nishiyama, Hong, Katsuhiko, Poo, & Kato, 2000; Feldman, 2000), but the evidence is very limited, possibly because longer time differences should be tested to observe this effect in additional preparations.

We have chosen to focus on excitatory synapses in this discussion, but our analytical derivation of Infomax learning is also relevant for inhibitory synapses (see appendix B). Applying the derivation for inhibitory synapses yields an asymmetric learning rule in which synapses become less inhibitory

if their activity is followed by a postsynaptic spike and is strengthened when the opposite order holds. Such a learning rule was observed in synapses of rat cortical interneurons (Holmgren & Zilberter, 2001).

Finally, the analysis suggests that the shape of the synaptic transfer function should also be plastic. EPSP shape is known to change on a short time scale (an effect termed *synaptic facilitation and depression*; Markram et al., 1997; Abbott et al., 1997) due to depletion of synaptic vesicles. Our analysis suggests that long-term changes in EPSP shape should take place following learning, where sharpening of EPSP should accompany synaptic strengthening.

In summary, the main predictions suggested by the model that can be experimentally tested with standard electrophysiological techniques are as follows. First, when comparing neurons with various electrophysiological propoerties, the STDP potentiation curve of each neuron should have a similar time constant as its EPSP's time constants. Second, in addition to synaptic depression that is observed when presynaptic spikes follow the postsynaptic spikes, weakening of synapses is expected to be observed for presynaptic spikes that largely precede the postsynaptic spikes (e.g., by three to four membrane time constants). Third, in parallel with potentiation of synaptic magnitude induced by STDP, the shape of the EPSP at the soma should sharpen.

Our analysis fails to explain several other forms of STDP such as the one observed in Dan and Poo (1998) and Egger, Feldmeyer, and Sakmann (1999), in which symmetric learning rules are observed. It is possible that a model that incorporates high-order temporal correlations in the input and output spike trains may be required to account for these findings.

It was conjectured that STDP supports the idea that information is stored in temporal spike patterns. It is shown here that even in a learning task in which information is coded in the underlying firing rates, spike-timing-dependent learning is required for effective learning. However, natural spike trains that trigger learning in vivo contain complex statistical structures (Paulsen & Sejnowski, 2000; Froemke & Dan, 2002). It thus remains important to explore learning rules that maximize mutual information about naturally structured spike trains.

Appendix A: Gradient-Ascent Learning Rule

This section derives a gradient-ascent learning rule that performs local search on the mutual information manifold. The input-output mutual information we aim to maximize is the mutual information between the output value and the identity of the presented pattern as defined by

$$I(Y;\eta) = h(Y) - h(Y \mid \eta) \quad ; \quad h(Y) = \int f(y)\log[f(y)]\,dy, \tag{A.1}$$

where the first term is the differential entropy of the output Y and the second is the differential entropy of the Y conditioned on input pattern presentation. The conditional entropy is

$$h(Y \mid \eta) = \sum_{\eta=0}^{M} p_\eta h(Y \mid \xi^\eta) = \sum_{\eta=0}^{M} p_\eta \frac{1}{2} \log(2\pi e {\sigma_\eta}^2), \tag{A.2}$$

and its derivative with respect to W_i is

$$\frac{\partial h(Y \mid \eta)}{\partial W_i} = \sum_{\eta=0}^{M} p_\eta \frac{1}{2} \frac{1}{{\sigma_\eta}^2} \frac{\partial {\sigma_\eta}^2}{\partial W_i}. \tag{A.3}$$

To calculate the entropy $H(Y)$, we use the fact that the distribution of Y is a mixture of gaussians, each resulting from the presentation of a different pattern $f(y) = \sum_{\eta=0}^{M} p_\eta \phi_\eta(y)$, and write

$$\begin{aligned} h(Y) &= -\int dy \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) \log\left(p_0 \phi_0(y) + \sum_{\eta=1}^{M} p_\eta \phi_\eta(y) \right) \\ &\approx -\log\left(\frac{p_0}{\sqrt{2\pi}\sigma_0} \right) \int dy \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) - \int dy \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) \frac{(y-\mu_0)^2}{2{\sigma_0}^2} \\ &\quad - \int dy \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) \frac{\sum_{\eta=1}^{M} p_\eta \phi_\eta(y)}{p_0 \phi_0(y)}, \end{aligned} \tag{A.4}$$

where the last equality relies on the fact that signal patterns are only rarely presented ($\sum_{\eta=1}^{M} p_\eta \ll 1$) and using the approximation $\log(x + \epsilon) \approx \log(x) + \frac{\epsilon}{x}$.

We now turn to differentiate each of the above three terms with regard to W_i. For the first term, we use $\int \sum_{\eta=0}^{M} p_\eta \phi_\eta(y)\, dy = 1$ and obtain

$$\begin{aligned} &-\frac{\partial}{\partial W_i} \log\left(\frac{p_0}{\sqrt{2\pi}\sigma_0} \right) \int dy \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) \\ &\quad = \frac{\partial}{\partial W_i} \frac{1}{2} \log({\sigma_0}^2) = \frac{1}{2} \frac{1}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i}. \end{aligned} \tag{A.5}$$

For the second term, we use $\int \phi_\eta(y)(y - \mu_\eta)\, dy = 0$ and obtain

$$\frac{\partial}{\partial W_i} \int dy \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) \frac{{\sigma_\eta}^2}{{\sigma_0}^2} \frac{(y - \mu_\eta + \mu_\eta - \mu_0)^2}{2{\sigma_\eta}^2}$$

$$
\begin{aligned}
&= \sum_{\eta=0}^{M} p_\eta \frac{\partial}{\partial W_i} \frac{1}{2} \frac{{\sigma_\eta}^2}{{\sigma_0}^2} + \sum_{\eta=0}^{M} p_\eta \frac{\partial}{\partial W_i} \frac{(\mu_\eta - \mu_0)^2}{2{\sigma_0}^2} \\
&= \sum_{\eta=0}^{M} \frac{p_\eta}{2{\sigma_0}^4} \left(\frac{\partial {\sigma_\eta}^2}{\partial W_i} {\sigma_0}^2 - {\sigma_\eta}^2 \frac{\partial {\sigma_0}^2}{\partial W_i} \right) \\
&\quad + \sum_{\eta=0}^{M} \frac{p_\eta}{2} \left(\frac{\frac{\partial (\mu_\eta - \mu_0)^2}{\partial W_i} {\sigma_0}^2}{{\sigma_0}^4} - \frac{(\mu_\eta - \mu_0)^2 \frac{\partial {\sigma_0}^2}{\partial W_i}}{{\sigma_0}^4} \right) \\
&= \sum_{\eta=1}^{M} p_\eta \frac{1}{2} \frac{1}{{\sigma_0}^2} \left(\frac{\partial {\sigma_\eta}^2}{\partial W_i} - \frac{{\sigma_\eta}^2}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i} + \frac{\partial (\mu_\eta - \mu_0)^2}{\partial W_i} \right. \\
&\qquad \left. - \frac{(\mu_\eta - \mu_0)^2}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i} \right), \qquad \text{(A.6)}
\end{aligned}
$$

where the last equality results from a vanishing term for $\eta = 0$. For the third term, we use the fact that $\sum_{\eta=1}^{M} p_\eta \ll 1$ and neglect second-order terms of $\sum_{\eta=1}^{M} p_\eta$:

$$
\begin{aligned}
&\frac{\partial}{\partial W_i} \int \sum_{\eta=0}^{M} p_\eta \phi_\eta(y) \frac{\sum_{\eta=1}^{M} p_\eta \phi_\eta(y)}{p_0 \phi_0(y)} \, dy \\
&\quad = \frac{\partial}{\partial W_i} \int p_0 \phi_0(y) \frac{\sum_{\eta=1}^{M} p_\eta \phi_\eta(y)}{p_0 \phi_0(y)} \, dy + \frac{\partial}{\partial W_i} \int \frac{(\sum_{\eta=1}^{M} p_\eta \phi_\eta(y))^2}{p_0 \phi_0(y)} \, dy \\
&\quad \approx \frac{\partial}{\partial W_i} (1 - p_0) = 0. \qquad \text{(A.7)}
\end{aligned}
$$

These three terms together yield

$$
\begin{aligned}
\frac{\partial h(Y)}{\partial W_i} &= \frac{1}{2} \frac{1}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i} + \sum_{\eta=1}^{M} p_\eta \frac{1}{2} \frac{1}{{\sigma_0}^2} \\
&\quad \times \left(\frac{\partial {\sigma_\eta}^2}{\partial W_i} - \frac{{\sigma_\eta}^2}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i} + \frac{\partial (\mu_\eta - \mu_0)^2}{\partial W_i} - \frac{(\mu_\eta - \mu_0)^2 \frac{\partial {\sigma_0}^2}{\partial W_i}}{{\sigma_0}^2} \right). \qquad \text{(A.8)}
\end{aligned}
$$

Combining equations A.2 and A.5 through A.7 while omitting the constant factor $\frac{1}{2}$, we obtain the gradient on the mutual information manifold,

$$
\begin{aligned}
\frac{\partial I(Y;\eta)}{\partial W_i} &= \frac{\partial}{\partial W_i} h(Y) - \frac{\partial}{\partial W_i} h(Y \mid \eta) \\
&\propto - \sum_{\eta=1}^{M} p_\eta \frac{1}{{\sigma_\eta}^2} \frac{\partial {\sigma_\eta}^2}{\partial W_i} - p_0 \frac{1}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i} + \frac{1}{{\sigma_0}^2} \frac{\partial {\sigma_0}^2}{\partial W_i}
\end{aligned}
$$

$$
\begin{aligned}
&+\sum_{\eta=1}^{M} p_\eta \frac{1}{{\sigma_0}^2}\left(\frac{\partial {\sigma_\eta}^2}{\partial W_i}-\frac{{\sigma_\eta}^2}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}\right.\\
&\qquad\left.+\frac{\partial(\mu_\eta-\mu_0)^2}{\partial W_i}-\frac{(\mu_\eta-\mu_0)^2}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}\right)\\
&=+\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}-p_0\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}-\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}\sum_{\eta=1}^{M} p_\eta\frac{{\sigma_\eta}^2}{{\sigma_0}^2}\\
&\quad-\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}\sum_{\eta=1}^{M} p_\eta\frac{(\mu_\eta-\mu_0)^2}{{\sigma_0}^2}\\
&\quad-\sum_{\eta=1}^{M} p_\eta\frac{1}{{\sigma_\eta}^2}\frac{\partial {\sigma_\eta}^2}{\partial W_i}+\sum_{\eta=1}^{M} p_\eta\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_\eta}^2}{\partial W_i}+\sum_{\eta=1}^{M} p_\eta\frac{1}{{\sigma_0}^2}\frac{\partial(\mu_\eta-\mu_0)^2}{\partial W_i}\\
&=\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}\left[(1-p_0)-\sum_{\eta=1}^{M} p_\eta\frac{{\sigma_\eta}^2}{{\sigma_0}^2}-\sum_{\eta=1}^{M} p_\eta\frac{(\mu_\eta-\mu_0)^2}{{\sigma_0}^2}\right]\\
&\quad+\sum_{\eta=1}^{M} p_\eta\frac{\partial {\sigma_\eta}^2}{\partial W_i}\left[\frac{1}{{\sigma_0}^2}-\frac{1}{{\sigma_\eta}^2}\right]+\sum_{\eta=1}^{M} p_\eta\frac{1}{{\sigma_0}^2}\frac{\partial(\mu_\eta-\mu_0)^2}{\partial W_i}\\
&=\frac{1}{{\sigma_0}^2}\frac{\partial {\sigma_0}^2}{\partial W_i}\sum_{\eta=1}^{M} p_\eta\left[1-\frac{{\sigma_\eta}^2}{{\sigma_0}^2}-\frac{(\mu_\eta-\mu_0)^2}{{\sigma_0}^2}\right]\\
&\quad+\sum_{\eta=1}^{M} p_\eta\frac{\partial {\sigma_\eta}^2}{\partial W_i}\left[\frac{1}{{\sigma_0}^2}-\frac{1}{{\sigma_\eta}^2}\right]+\sum_{\eta=1}^{M} p_\eta\frac{1}{{\sigma_0}^2}\frac{\partial(\mu_\eta-\mu_0)^2}{\partial W_i}, \qquad \text{(A.9)}
\end{aligned}
$$

where we used the fact that $1-p_0=\sum_{\eta=1}^{M} p_\eta$. Substituting the derivatives $\frac{1}{2}\frac{\partial {\sigma_\eta}^2}{\partial W_i}=\mathrm{Cov}(Y^\eta, X_i^\eta)$ and $\frac{1}{2}\frac{\partial(\mu_\eta-\mu_0)^2}{\partial W_i}=(\mu_\eta-\mu_0)(\langle X_i^\eta\rangle-\langle X_i^0\rangle)$ into equation A.9, we obtain

$$
\begin{aligned}
\frac{\partial}{\partial W_i}I(Y;\eta)&=\sum_{\eta=1}^{M} p_\eta\,\mathrm{Cov}(Y^\eta, X_i^\eta)K_\eta^1-\sum_{\eta=1}^{M} p_\eta\mathrm{Cov}(Y^0, X_i^0)K_\eta^0\\
&\quad+\sum_{\eta=1}^{M} p_\eta\frac{1}{{\sigma_0}^2}(\mu_\eta-\mu_0)(\langle X_i^\eta\rangle-\langle X_i^0\rangle), \qquad \text{(A.10)}
\end{aligned}
$$

where

$$
\begin{aligned}
K_\eta^0&=\frac{{\sigma_\eta}^2-{\sigma_0}^2}{{\sigma_0}^4}+\frac{(\mu_\eta-\mu_0)^2}{{\sigma_0}^4} \qquad \text{(A.11)}\\
K_\eta^1&=\frac{1}{{\sigma_0}^2}-\frac{1}{{\sigma_\eta}^2}.
\end{aligned}
$$

In the case of inhomogeneous Poisson process, in which spikes in the train are uncorrelated but their underlying rate may vary along time,

$$\begin{aligned}\mathrm{Cov}(Y, X_i) &= \mathrm{Cov}\left(\sum_{i=1}^{N} W_i X_i, X_j\right) = \mathrm{Cov}(W_i X_i, X_i) = W_i \,\mathrm{Var}(X_i) \\ &= W_i V\left[\int_{t'=-T}^{0} F_\tau(0-t')S(t')\right] = W_i \int_{t'=-T}^{0} V[F_\tau(0-t')S(t')] \\ &= W_i \int_{t'=-T}^{0} F_\tau^2(0-t')V[S(t')] \\ &\approx W_i \int_{t'=-T}^{0} F_\tau^2(0-t')E[S(t')], \end{aligned} \tag{A.12}$$

where in the case of exponential filter ($F_\tau(x) = \frac{1}{\tau}\exp(-x/\tau)$), we obtain $\mathrm{Cov}(Y, X_i) = W_i E_{\tau/2}(X_i)$ where $E_\tau(X) = \int \exp(\frac{t-t'}{\tau})S(t')\,dt'$.

Appendix B: Mirror Fixed Points of the Learning Rule

Let W^* be a fixed-point solution of the optimal learning of equation 3.4, that is,

$$\begin{aligned}\left.\frac{\partial I(Y;\eta)}{\partial W_i}\right|_{W^*} &= +\sum_{\eta=1}^{M} p_\eta(\mathrm{Cov}(Y, X_i^\eta)K_\eta^1(W^*) + E(X_i^\eta)K_\eta^2(W^*)) \\ &\quad -\sum_{\eta=1}^{M} p_\eta(\mathrm{Cov}(Y, X_i^0)K_\eta^0(W^*) + E(X_i^0)K_\eta^2(W^*)) = 0. \end{aligned} \tag{B.1}$$

Now note that both K_η^1 and K_η^0 are even functions of W, but K_η^2 and $\mathrm{Cov}(Y, X_i)$ $= \mathrm{Cov}(WX, X)$ are odd functions of W. This yields

$$\begin{aligned}\left.\frac{\partial I(Y;\eta)}{\partial W_i}\right|_{-W^*} &= +\sum_{\eta=1}^{M} p_\eta((-\mathrm{Cov}(Y, X_i^\eta))K_\eta^1(W^*) + E(X_i^\eta)(-K_\eta^2(W^*))) \\ &\quad -\sum_{\eta=1}^{M} p_\eta((-\mathrm{Cov}(Y, X_i^0))K_\eta^0(W^*) + E(X_i^0)(-K_\eta^2(W^*))) \\ &= -\left.\frac{\partial I(Y;\eta)}{\partial W_i}\right|_{W^*} = 0. \end{aligned} \tag{B.2}$$

Therefore, any solution W^* has a mirror solution $-W^*$, which is also a fixed point of the dynamics.

Appendix C: Pattern-Triggered Learning

The batch learning rule of equation 3.4 can be turned into a stochastic on-line version by replacing summation over patterns with a learning rule that modifies the synaptic weights according to the input pattern presented at that moment. There are two fundamentally different alternative implementations: changing synaptic weights triggered by the presentation of all patterns,

$$\Delta W_i = \begin{cases} -(\lambda^0 \operatorname{Cov}[Y, X_i^0] + \lambda^2 E[X_i^0])(\frac{1}{p_0}\sum_{\eta=1}^{M} p_\eta) \\ \quad \text{when } \xi^0 \text{ is presented} \\ -\lambda^1 \operatorname{Cov}[Y, X_i^\eta] + \lambda^2 E[X_i^\eta] \\ \quad \text{when } \xi^\eta \text{ is presented, } \eta > 0, \end{cases} \tag{C.1}$$

or limiting learning to follow the presentation of the foreground patterns only,

$$\Delta W_i = +(\lambda^1 \operatorname{Cov}[Y, X_i^\eta] + \lambda^2 E[X_i^\eta]) - (\lambda^0 \operatorname{Cov}[Y, X_i^0] + \lambda^2 E[X_i^0]) \\ \text{when } \xi^\eta \text{ is presented,} \quad \text{for } \eta = 1..M \text{ only.} \tag{C.2}$$

These two alternatives have different consequences for on-line learning: The latter rule ("learn on rare patterns", equation 4.1) does not explicitly depend on the prior probabilities p_η while the former rule (learn on all patterns, equation C.1) does. The former rule thus suffers from a major drawback when the prior probabilities are unknown or vary in time, because the system has to continuously estimate these probabilities. If the priors p_η are incorrectly estimated, strengthening and weakening of synaptic weights goes out of balance, and synaptic values drift to their limits, causing the neuron to lose all of its discriminative power. This is similar to saturation effects observed in traditional models of Hebbian learning as discussed in Sejnowski (1977), Dayan and Willshaw (1991), and Chechik et al. (2001). To demonstrate this effect, we have conducted the following series of experiments. We set p_η to fluctuate along time, while the learning parameters were set to the average p_η values. Figure 8 compares the above two learning rules and shows that the "learn on all patterns" learning rule is sensitive to fluctuations in p_η, while the "learn on rare patterns" rule is robust to these fluctuations. We conclude that an on-line implementation of the batch rule should take the form of equation 4.1, where learning is triggered by the rare patterns only.

Appendix D: Maximum Output Variance

A clearer intuition into the structure of the Infomax learning rule can be obtained by comparing it with another learning rule, derived according

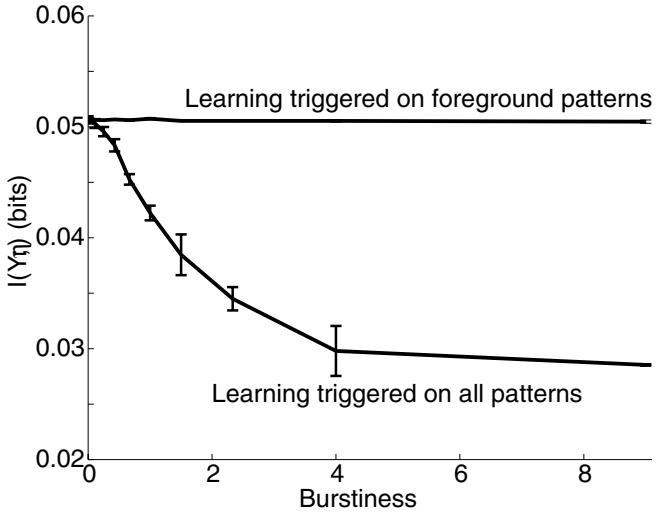

Figure 8: Learning patterns with fluctuating frequencies p_η. The two learning rules of equations 4.1 and C.1 are compared, showing that learning that is triggered on rare patterns is much more robust to such fluctuations. Fluctuating patterns were modeled in the following way. Pattern presentation was divided to five cycles, each containing 2000 presentation steps. Each cycle consisted of a period of length L_0, where only the background pattern was presented, and the remaining period (with length $L_\eta = 2000 - L_0$), where the rest of the patterns are randomly chosen for presentation such that the average probabilities p_η are preserved. The ratio $B = \frac{L_\eta}{L_0}$ thus provides a burstiness measure, where high B values correspond to input statistics with bursts of foreground input patterns. Information value is the average over the last cycle containing both quiet periods and bursts. Simulations were repeated for five different runs, and errors depict variability across these runs.

to the criterion of maximum output variance. Traditional Hebbian learning (learning by correlation) is known to perform gradient ascent on the manifold of this target function. This optimization criterion must be complemented with a constraint over the weights; otherwise, synaptic weights diverge (Sejnowski, 1977; Dayan & Willshaw, 1991). Consider therefore the following target function,

$$\mathrm{Var}(Y) - \lambda \sum_i W_i^\alpha, \tag{D.1}$$

where α is commonly taken to be 1 (additive normalization) or 2 (multiplicative normalization). Differentiating this target function with regard to

W_i yields

$$\frac{\partial}{\partial W_i}\left(\mathrm{Var}(Y) - \lambda \sum_i W_i^\alpha\right) = \mathrm{Cov}(Y, X_i) - \lambda\alpha W_i^{\alpha-1}. \tag{D.2}$$

In this learning rule, all patterns strengthen synapses in a way that corresponds to their input's strength. Synaptic weakening is enforced through a global decay term. The balance between activity-dependent input-output correlation and normalization should yield a nondivergent solution.

In contradistinction, the Infomax learning rule derived in this article tends to strengthen synaptic weights that correspond to strong inputs of some patterns while decreasing the synapses of other pattern. Depression is thus achieved not through a global competition mechanism but is input specific.

Acknowledgments

It is a pleasure to thank N. Tishby, E. Ruppin, D. Horn, and R. Aharonov for helpful discussions and E. Schneidman, A. Globerson, and M. Ben-Shachar for a careful reading of the manuscript. This research was supported by a grant from the Ministry of Science, Israel.

References

Abbott, L., & Song, S. (1999). Temporally asymmetric Hebbian learning, spike timing and neural respons variability. In S. A. Solla & D. A. Cohen (Eds.), *Advances in neural information processing systems, 11* (pp. 69–75). Cambridge, MA: MIT Press.

Abbott, L., Varela, J., Sen, K., & Nelson, S. (1997). Synaptic depression and cortical gain control. *Science, 275*, 220–223.

Aharonov, R., Guttig, R., & Sompolinsky, H. (2001). Generalized synaptic updating in temporally asymmetric Hebbian learning. In J. Bower (Ed.), *Computational neuroscience: Trends in research 2001*. Amsterdam: Elsevier.

Bell, C., Han, V., Sugawara, Y., & Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature, 387*, 278–281.

Bi, Q., & Poo, M. (1999). Precise spike timing determines the direction and extent of synaptic modifications in cultured hippocampal neurons. *J. Neurosci., 18*, 10464–10472.

Chechik, G., Horn, D., & Ruppin, E. (2002). Hebbian learning and neuronal regulation. In M. A. Arbib (Ed.), *Handbook of brain theory and neural networks* (2nd ed.). Cambridge, MA: MIT Press.

Chechik, G., Meilijson, I., & Ruppin, E. (2001). Effective neuronal learning with ineffective Hebbian learning rules. *Neural Computation, 13*(4), 817–840.

Cover, T., & Thomas, J. (1991). *Elements of information theory*. New York: Wiley.
Dan, Y., & Poo, M. (1998). Hebbian depression of isolated neuromascular synapses in vitro. *Science, 256*, 1570–1573.
Dayan, P., & Willshaw, D. (1991). Optimizing synaptic learning rules in linear associative memories. *Biol. Cyber., 65*, 253.
Debanne, D., Gahwiler, B., & Thompson, S. (1994). Asynchronous pre- and post-synaptic activity induces associative long-term depression in area CA1 of the rat hippocampus in vitro. *Proc. Natl. Acad. Sci., 91*, 1148–1152.
Debanne, D., Gahwiler, B., & Thompson, S. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol., 507*(1), 237–247.
Egger, V., Feldmeyer, D., & Sakmann, B. (1999). Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nature Neuroscience, 2*, 1098–1105.
Feldman, D. (2000). Timing based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron, 27*, 45–56.
Ferguson, T. (1996). *A course in large sample theory*. London: Chapman and Hall.
Froemke, R., & Dan, Y. (2002). Spike-timing dependent synaptic modification induced by natural spike trains. *Nature, 416*, 433–438.
Holmgren, C., & Zilberter, Y. (2001). Coincident spiking activity induces long term changes in inhibition of neocortical pyramidal cells. *J. Neuroscience, 21*, 8270–8277.
Horn, D., Levy, N., Meilijson, I., & Ruppin, E. (2000). Distributed synchrony of spiking neurons in a Hebbian cell assembly. In S. Solla, T. Leen, & K. Muller (Eds.), *Advances in neural information processing systems, 12* (pp. 129–135). Cambridge, MA: MIT Press.
Kempter, R., Gerstner, W., & van Hemmen, J. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E, 59*(4), 4498–4514.
Kempter, R., Gerstner, W., & van Hemmen, J. (2001). Intrinsic stabilization of output rates by spike-time dependent Hebbian learning. *Neural Computation, 13*, 2709–2742.
Kingman, J. (1993). *Poisson processes*. New York: Oxford University Press.
Lehman, E. (1998). *Elements of large sample theory*. New York: Springer-Verlag.
Levy, W., & Steward, D. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neurosci., 8*, 791–797.
Linsker, R. (1988). Self-organization in a perceptual network. *Computer, 21*(3), 105–117.
Linsker, R. (1992). Local synaptic learning rules suffice to maximize mutual information in a linear network. *Neural Computation, 4*(5), 691–702.
Magee, J., & Johnston, D. (1997). A synaptic controlled associative signal for Hebbian plasticity in hippocampal neurons. *Science, 275*(5297), 209–213.
Markram, H., Lubke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science, 275*(5297), 213–215.
Mehta, M., Quirk, M., & Wilson, M. (1999). From hippocampus to V1: Effect of LTP on spatio-temporal dynamics of receptive fields. In J. Bower

(Ed.), *Computational neuroscience: Trends in research 1999*. Amsterdam: Elsevier.

Monyer, H., Burnashev, N., Laurie, D., Sakmann, B., & Seeburg, P. (1994). Developmental and regional expression in the rat brain and functional properties of four NMDA receptors. *Neuron, 12*, 529–540.

Nishiyama, M., Hong, K., Katsuhiko, M., Poo, M., & Kato, K. (2000). Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature, 408*, 584–588.

Paulsen, O., & Sejnowski, T. (2000). Natural patterns of activity and long term synaptic plasticity. *Current Opinion in Neurobiology, 10*, 172–179.

Rao, R., & Sejnowski, T. (2000). Predictive sequence learning in recurrent neocortical circuits. In S. Solla, T. Leen, & K. Muller (Eds.), *Advances in neural information processing systems, 12* (pp. 164–170). Cambridge, MA: MIT Press.

Roberts, P. D., & Bell, C. C. (2002). Spike-timing dependent synaptic plasticity in biological systems. *Biol. Cybern., 87*, 392–403.

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

Schultz, W., Dayan, P., & Montague, P. (1997). A neural subtrate of prediction and reward. *Science, 275*, 1593–1599.

Segal, M., & Andersen, P. (2000). Dendritic spines shaped by synaptic activity. *Curr. Opin. Neurobiol., 10*, 582–587.

Segev, I., & London, M. (2000). Untangling dendrites with quantitative models. *Science, 290*, 744–750.

Sejnowski, T. (1977). Statistical constraints on synaptic plasticity. *J. Theo. Biol., 69*, 385–389.

Shannon, C. (1948). A mathematical theory of communication. *Bell Syst. Tech. J., 27*, 379–423.

Singer, W., & Gray, C. (1995). Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci., 18*, 555–586.

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

Tang, Y. P., Shimizu, E., Dube, G. R., Rampon, C., Kerchner, G. A., Zhuo, M., Liu, G., & Tsien, J. Z. (1999). Genetic enhancement of learning and memory in mice. *Nature, 401*, 63–69.

Yao, H., & Dan, Y. (2001). Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron, 32*, 315–323.

Zhang, L., Tao, H. W., Holt, C., Harris, W., & Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature, 395*(3), 37–44.

Received April 29, 2002; accepted January 5, 2003.