FACETS

FP6-2004-IST-FETPI 15879

Fast Analog Computing with Emergent Transient States

`Functional consequences of Plasticity models'

Report Version: 1.0
Classification: Public
Type: Report

Due date:
Date issued: end of Month 48/ End of August 2009
Report Preparation: Wulfram Gerstner

Contract Start Date: 1 September 2005 Duration: 5 Years
Project Coordinator: Karlheinz Meier (Heidelberg)
Partners: ENSEIR Bordeaux, CNRS (Gif-sur-Yvette, Marseille), U Debrecen, TU Dresden, U Freiburg, TU Graz, U Heidelberg, EPFL Lausanne, U London, U Plymouth, INRIA, KTH Stockholm, SCCH

Project funded by the European Community under the "Information Society Technologies" Programme

DELIVERABLE SUMMARY SHEET

Project Number: FP6-2004-IST-FETPI 15879

Project Acronym: FACETS

Title: Fast Analog Computing with Emergent Transient States

Deliverable N°: D 4-6

Due date: end of Month 48/ End of August 2009

Delivery Date: end of Month 48/ End of August 2009

Short description:

In August 2007 with the milestone M4-1 'Availability of plasticity model concepts and discussion of their relevance to software and hardware modelling approaches', the Facets consortium has already agreed on a framework of how to treat plasticity in the hardware as well as in the software simulations. This plasticity framework offers a certain amount of freedom regarding the choice of parameters and allows for the instantiation of different specific versions of plasticity, e.g. different weight-dependence, different voltage dependence, specific choices of spike-timing dependence.

Within the framework, the FACETS partners have studied the functional properties of several models.

The work can be summarized as follows:

(i) Voltage and Frequency Dependence of STPD and LTP and their functional consequences

Partner EPFL-LCN has developped a model where presynaptic activity is combined with the voltage of the postsynaptic membrane. The essential model ideas were developed while Lars Buesing from FACETS partner TU Graz was spending several weeks at the EPFL-LCN. The voltage enters into the model both as instantaneous voltage and as filtered voltage averaged over the recent past (in the range of some tens of millisecond). For example, LTP occurs if a presynaptic spike arrive AND shortly afterwards the instantaneous voltage is above firing threshold AND the average voltage is significantly above rest. LTD occurs if the average voltage is slightly above rest at the moment when a presynaptic spike arrives. This model accounts for

(a) the voltage dependence found by Artola-Broecher-Singer;

(b) STDP found by Markram or Bi and Poo;

(c) the frequency dependence of STDP found by Markram and Sjoestroem;

(d) the intricate interaction of voltage dependence and spike-timing dependence reported by Sjoestroem.

Moreover, this model has been implemented in small to medium-sized recurrent networks driven by different stimulation protocols. It was found that for inputs that show slow rate modulation, the network develops a structure with strong bidirectional connections, leading to clusters of strongly connected neurons. However, the same learning rule and the same network showed strong unilateral connections and no lateral connectivity when driven with a stimulus showing systematic and rapid temporal modulation. These results suggest an intriguing relation between network connectivity and neural coding.

The model has been presented in a manuscript submitted to Nature Neuroscience and available as Nature precedings

Claudia Clopath, Lars Büsing, Eleni Vasilaki, Wulfram Gerstner:

Connectivity reflects coding: A model of voltage-based spike-timing-dependent-plasticity with homeostasis

Nature Neuroscience (submitted)

http://precedings.nature.com/documents/3362/version/1

This manuscript is included as appendix 1.

(ii) Frequency dependence of STDP

In a parallel study, Christian Mayr from the TU Dresden tested whether the same scenarios (a) - (d) could also be explained by a slightly different variant of the FACETS model.

C. Mayr, J. Partzsch, and R. Schüffny, BCM and

Membrane Potential: Alternative Ways to Timing Dependent Plasticity,

15th International Conference on Neural Information Processing (ICONIP

2008), Springer LNCS, vol. 5506, pp. 137-144, 2009

This paper is added as appendix 2.

(iii) The voltage based plasticity rule discussed under point (i) has been used by partner EPFL-LCN as the plasticity induction step in a model that combines induction of LTP with consolidation during a late phase of LTP. After the induction, a dopaminergic (or other neuromodulatory signal) is necessary to trigger protein synthesis and to transform the early trace of LTP into stable synaptic weights.

This model is related to the paradigm of tagging introduced by Frey and Morris more than 10 years ago and explains a large body of tagging, cross-tagging, and consolidation experiments.

The study has appeared in December 2008 in PLOS Computational Biology

C. Clopath, L. Ziegler, E. Vasilaki, L. Büsing and W. Gerstner,

Tag-Trigger-Consolidation: A Model of Early and Late

ong-Term-Potentiation and Depression,

PLoS Comput Biol, Vol. 4, Nr. 12, 2008.

This study is included as appendix 3.

(iv) Functional consequences of Spike-Timing Dependent Plasticity modulated by reward have been analyzed by the FACETS partner TU-Graz. Legenstein et al. derived conditions under which modulation of synaptic plasticity can learn a given task. Moreover, they applied the learning rule to a model of biofeedback of Fetz and Bakerso as to illustrate some of the functional consequences.

The work of Legenstein et al. has appeared in PLOS Computational Biology

A Learning Theory for Reward-Modulated Spike-Timing-Dependent Plasticity with Application to Biofeedback. Robert Legenstein, Dejan Pecevski, and Wolfgang Maass

PLoS Comput Biol. 2008 October; 4(10): e1000180.

This paper is included as appendix 4.

(v) Moritz Helias and colleages from the FACETS partner ALUF in Freiburg showed how a

simplified plasticity model, similar the one used in the FACETS consortium can be derived from basic biophysical signals such as the calcium concentration in spines and correlation detection by Hebbian synaptic plasticity. Moreover, they could relate this model to structural plasticity, i.e., changes in spines and connection patterns rather than only changes in synaptic weights. This work has appeared in Frontiers in Computational Neuroscience:

M. Helias, S. Rotter, M.-O. Gewaltig, and Markus Diesmann,

Structural Plasticity controlled by calcium based correlation detection

Frontiers in Computational Neuroscience December 2008 2: 7

This paper is included as appendix 5.

Partners owning: Wulfram Gerstner, (Partner 8 EPFL-LCN)

Partners contributed: Institute for Theoretical Computer Science, Graz; Hochparallele VLSI-Systeme und Neuromikroelektronik, Dresden, ALUF Freiburg.

Made available to: Public

2 Appendix

2.1 Appendix 1

Claudia Clopath, Lars Bsing, Eleni Vasilaki, Wulfram Gerstner: Connectivity reflects coding: A model of voltage-based spike-timing-dependent-plasticity with homeostasis Nature Neuroscience (submitted)

http://precedings.nature.com/documents/3362/version/1

Connectivity reflects Coding: A Model of Voltage-based Spike-Timing-Dependent-Plasticity with Homeostasis

Claudia Clopath, Lars Büsing*, Eleni Vasilaki, Wulfram Gerstner

Laboratory of Computational Neuroscience

Brain-Mind Institute and School of Computer and Communication Sciences

Ecole Polytechnique Fédérale de Lausanne

1015 Lausanne EPFL, Switzerland

* permanent address: Institut für Grundlagen der Informationsverarbeitung, TU Graz, Austria

June 3, 2009

Abstract

Electrophysiological connectivity patterns in cortex often show a few strong connections in a sea of weak connections. In some brain areas a large fraction of strong connections are bidirectional, in others they are mainly unidirectional. In order to explain these connectivity patterns, we use a model of Spike-Timing-Dependent Plasticity where synaptic changes depend on presynaptic spike arrival and the postsynaptic membrane potential, filtered with two different time constants. The model describes several nonlinear effects in STDP experiments, as well as the voltage dependence of plasticity under voltage clamp and classical paradigms of LTP/LTD induction. We show that in a simulated recurrent network of spiking neurons our plasticity rule leads not only to receptive field development, but also to connectivity patterns that reflect the neural code: for temporal coding paradigms strong connections are predominantly unidirectional, whereas they are bidirectional under rate coding. Thus variable connectivity patterns in the brain could reflect different coding principles across brain areas; moreover our simulations suggest that rewiring the network can be surprisingly fast.

1 Introduction

Experience-dependent changes in receptive fields [1, 2, 3] or in learned behavior [4] may occur through changes in synaptic strength. Thus, electrophysiological measurements of functional connectivity patterns in slices of

neural tissue [5, 6] or anatomical connectivity measures [7] can only present a snapshot of the momentary connectivity – which may change with the next set of stimuli. Indeed, modern imaging methods show that spine motility can lead to a rapid rewiring of the connectivity pattern [8, 9] by formation of new synapses or by strengthening or weakening of existing synapses. The question then arises whether the connectivity patterns and changes that are found in experiments can be connected to basic rules of synaptic plasticity, in particular to modern or traditional forms of Hebbian plasticity [10] such as Long-Term Potentiation and Depression [11].

Long-term potentiation LTP and depression LTD of synapses depends on the exact timing of pre- and postsynaptic action potentials [12, 13], but also on postsynaptic voltage [14, 15], and presynaptic stimulation frequency [16]. Spike-Timing-Dependent Plasticity (STDP) has attracted particular interest in recent years, since temporal coding schemes where information is contained in the exact timing of spikes rather than mean frequency could be learned by a neural system using STDP [17, 18, 19, 20, 21]. The question, however, whether STDP is more fundamental than frequency dependent plasticity or voltage dependent plasticity rules has not been resolved, despite an intense debate [22]. Moreover it is unclear how the interplay of coding and plasticity yield the functional connectivity patterns seen in experiments. In particular, the presence or absence of bidirectional connectivity between cortical pyramidal neurons seems to be contradictory across experimental preparations in visual [5] or somatosensory cortex [6].

Recent experiments have shown that STDP is strongly influenced by postsynaptic voltage before action potential firing [23], but could not answer the question whether spike timing dependence is a direct consequence of voltage dependence, or the manifestation of an independent process. In addition, STDP depends on stimulation frequency [23] suggesting an interaction between timing and frequency dependent processes — or this interaction could be the manifestation of a single process in different experimental paradigms. We show that a simple Hebbian plasticity rule that pairs presynaptic spike arrival with the postsynaptic membrane potential is sufficient to explain STDP and the dependence of plasticity upon presynaptic stimulation frequency. Moreover, the intricate interplay of voltage and spike-timing dependence seen in experiments [23] as well as the frequency dependence of STDP can be explained in our model from one single principle. In contrast of earlier attempts towards a unified description of synaptic plasticity rule that focused on detailed biophysical descriptions [24, 25], our model is a mechanistic one (phenomenological model). It does not give an explicit interpretation in terms of biophysical quantities such a Calcium concentration [24], CaMKII [25], glutamate binding, NMDA receptors etc. Rather it aims at a minimal description of the major phenomena observed in electrophysiology experiments.

The advantage of such a minimal model is that it allows us to discuss functional consequences in small [26, 27], and possibly even large [28, 29], networks. We show that in small networks of up to 10 neurons the learning rule leads to input specificity, necessary for receptive field development - similar to earlier models of STDP [17, 26] or rate-based plasticity rules [30, 31]. Going significantly beyond earlier studies we explicitly address the question of whether functional connectivity patterns of cortical pyramidal neurons measured in recent electrophysiological studies [5, 6] could be the result of plasticity during continued stimulation of neuronal model networks. We found that connectivity patterns strongly depend on the underlying coding hypothesis: With a temporal coding hypothesis, where input spikes arrive in a fixed temporal order, the recurrent network develops a connectivity

pattern with a few strong unidirecitional connections. However, under a rate coding paradigm, where stimuli are stationary during a few hundred milliseconds the same network exhibits sustained and strong bidirectional connections. This is in striking contrast to standard STDP rules where bidirectional connections are impossible [26].

The mathematical simplicity of the model enables us to identify conditions under which it becomes equivalent to the well-known Bienenstock-Cooper-Munro model [30] used in classical rate-based descriptions of developmental learning; and equivalent to some earlier models of STDP [32] — and why our model is fundamentally different from classical STDP models [17, 26, 21], widely used for temporal coding.

2 Results

In order to study how connectivity patterns in cortex can emerge from an interplay of plasticity rules and coding, we need a plasticity rule that is consistent with a large body of experiments, not just a single paradigm such as STDP. Since synaptic depression and potentiation take place through different pathways [33] our model uses separate additive contributions to the plasticity rule, one for LTD and another one for LTP (see Fig. 1 and methods).

2.1 Fitting the Plasticity Model to Experimental Data

Consistent with voltage clamp [15] and stationary depolarization experiments [14] LTD is triggered in our model if presynaptic spike arrival occurs while the membrane potential of the postsynaptic neuron is slightly depolarized (above a threshold θ_-) whereas LTP occurs if depolarization is big (above a second threshold θ_+ (see Fig. 1). The mathematical formulation of the plasticity rule makes a distinction between the momentary voltage u and the low-pass filtered voltage variables $\bar{u}_-$ or $\bar{u}_+$ which denote temporal averages of the voltage over the recent past (the symbols $\bar{u}_-$ and $\bar{u}_+$ indicate filtering of u with two different time constants). Similarly, the event x of presynaptic spike arrival needs to be distinguished from the trace $\bar{x}(t)$ that is left at the synapse after stimulation by neurotransmitter. Potentiation occurs only if the momentary voltage is above θ_+ (this condition is fulfilled during action potential firing) AND the average voltage $\bar{u}_+$ above θ_- (this is fulfilled if there has a been a depolarization in the recent past) AND the trace $\bar{x}$ left by a previous presynaptic spike event is nonzero (this condition holds if a presynaptic spike arrived a few milliseconds earlier at the synapse); these conditions for plasticity are illustrated in Fig. 1B. LTD occurs if the average voltage $\bar{u}_-$ is above rest at the moment of a presynaptic spike arrival (see Fig. 1A). The amount of LTD in our model depends on homeostatic process on a slower time scale [34]. Low-pass filtering of the voltage by the variable ($\bar{u}_-$ or $\bar{u}_+$) refers to some unidentified intracellular processes triggered by depolarization, e.g., increase in calcium concentration or second messengers messenger chains. Similarly, the biophysical nature of the trace $\bar{x}$ is irrelevant for the functionality of the model, but a good candidate process is the fraction of glutamate bound to postsynaptic receptors.

We checked the performance of the model on a simulated STDP protocol, where presynaptic spikes arrive

a few milliseconds before or after a postsynaptic spike that is triggered by a strong depolarizing current pulse. If a post-pre pairing with a timing difference of 10 millisecond is repeated 60 times at frequencies below 35Hz, LTD occurs in our model (Fig. 2 A, B), consistent with experiments [23]. Repeated pre-post pairings (with 10 millisecond timing difference) at frequencies above 10Hz yield LTP, but pairings at 0.1Hz do not show any significant change in the model or in experiments [23]. In the model these results can be explained by the fact that at 0.1Hz repetition frequency, the low-pass filtered voltage $\bar{u}_+$ which increases abruptly during postsynaptic spiking decays back to zero before the next impulse arrives, so that LTP can not be triggered. However, since LTD in the model requires only a weak depolarization of $\bar{u}_-$ at the moment of presynaptic spike arrival, post-pre pairings give rise to depression, even at very low frequency. At repetition frequencies of 50Hz, the post-pre paradigm is nearly indistinguishable from a pre-post timing, and LTP dominates.

Since spike-timing dependence in our model is induced only indirectly via voltage dependence of the model, we wondered whether our model would also be able to account for the intricate interactions of voltage and spike timing found by Sjöström et al. [23]. If a pre-post protocol at 0.1Hz, that normally does not induce LTP, is combined with a depolarizing current pulse (lasting from 50ms before to 50ms after the postsynaptic firing event), then potentiation is observed in the experiments [23], as well as in our model (Fig. 2 C, F, I). Due to the injected current, the low-pass filtered voltage variable $\bar{u}_+$ is depolarized before the pairing. Thus at the the moment of the postsynaptic spike, the average voltage $\bar{u}_+$ is above the threshold θ_- leading to potentiation. Similarly, a pre-post protocol that normally leads to LTP can be blocked if the postsynaptic spikes are triggered on the background of a hyperpolarizing current (Fig. 2 E, H, I).

In order to study some nonlinear aspects of STDP, we simulate a protocol of burst-timing-dependent plasticity where presynaptic spikes are paired with 1, 2 or 3 postsynaptic spikes [35] (see Methods). We observe that 60 pre-post pairs at 0.1Hz do not change the synaptic weight, as discussed above. However, repeated triplets pre-post-post generate potentiation in our model because the first postsynaptic spike induces a depolarizing spike after potential so that $\bar{u}_+$ is depolarized. Adding a third postsynaptic spike to the protocol (i.e., quadruplets pre-post-post-post) does not lead to stronger LTP (Fig. 3A). Our model also describes the dependence of LTP upon the intra-burst frequency (Fig. 3B). At an intra-burst frequency of 20Hz, no LTP occurs, because the second spike in the burst comes so late that the presynaptic trace $\bar{x}$ has decayed back to zero. At higher intra-burst frequencies, the three conditions for LTP ($u(t) > \theta_+$ and $\bar{u}_+ > \theta_-$ and $\bar{x} > 0$) are fulfilled. The burst timing dependence (Fig. 3C) is qualitatively similar to that found in experiments [35], but only four of the six experimental data points are quantitatively reproduced by the model.

2.2 Functional implications

Connectivity patterns in a local cortical circuit have been shown to be non-random, i.e. the majority of connections are weak and the rare strong ones have a high probability of being bidirectional [5]. However, standard models of STDP do not exhibit stable bidirectional connections [36]. Intuitively, if the cell A fires before the cell B, a pre-post pairing for the 'AB' connection is formed so that the connection is strengthened. The post-pre

pairing occurring at the same time in the 'BA' connection leads to depression. Therefore it is impossible to strengthen both connections at the same time. Moreover, in order to assure long-term stability of firing rates parameters in standard STDP rules are typically chosen such that inhibition slightly dominates excitation [17] which implies that under purely random spike firing connections decrease, rather than increase. However, the non-linearity aspects of plasticity in our model change such a simple picture. If we simulate two neurons with bidirectional connections at low firing rates, the plasticity model behaves like standard STDP and only unidirectional connections emerge. However, from Fig. 3B we expect that at higher neuronal firing rates, our model could develop a stable bidirectional connection, in striking contrast to standard STDP rules [21].

Since bidirectional connections require neurons to fire at a high rate, we wondered how coding and connectivity relate to each other. We hypothesized that bidirectional connections are supported by rate-coding as opposed to temporal-coding. To test this idea we first simulated a small network of 10 all-to-all connected neurons in a simplified rate-coding scheme where each neuron fires at a fixed frequency, but the frequency varies across neurons. We find that bidirectional connections are formed only between those neurons that both fire at a high rate, but not if one or both of the neurons fire at low frequencies (Fig. 4A). In a second paradigm, the neurons in the same network are stimulated such that they are firing in a distinct order (1, 2 , 3,..) mimicking an extreme form of temporal coding [37]. In that case, the weights form a loop where strong connections from 1 to 2, 2 to 3, ... develop, but no bidirectional connections (Fig. 4B). These results are in striking contrast to simulation experiment with a standard STDP rule, where connections are always unidirectional, independently of coding (Fig. 4C, D).

We wondered whether the same results would emerge in a more realistic network of excitatory and inhibitory neurons driven by feedforward input. We simulated a network of 10 excitatory neurons and 3 inhibitory neurons. Each inhibitory neuron receives input from 8 randomly selected excitatory neurons and randomly projects back to 6 excitatory neurons. In addition to the recurrent input, each excitatory neuron receives feedforward spike input from 500 presynaptic neurons j that generate stochastic Poisson input at a rate ν_j. The rates of neighboring input neurons are correlated, mimicking the presence or absence of spatially extended objects. In a rate-coding scheme, the location of the stimulus is switched every 100ms to a new random position. In case of retinal input, this would correspond to a situation where the subject fixates every 100ms on a new stationary stimulus. In a temporal-coding paradigm, the model input is shifted every 20ms to a neighboring location, mimicking movement of an object across an array of sensory receptors. For both scenarios the network is identical. Feedforward connections and lateral connections between model pyramidal neurons are plastic whereas connections to and from inhibitory neurons are fixed.

After 1000s of stimulation with the rate-coding paradigm, the excitatory neurons developed localized receptive fields and a structured pattern of synaptic connections (Fig. 5B). While the labeling of the excitatory neurons at the beginning of the experiment was randomly assigned, we can relabel the neurons after the formation of lateral connectivity patterns so that neurons with strong reciprocal connections have similar indices, reflecting the neighborhood relation of the network topology. After reordering we can clearly distinguish that three groups of neurons have been formed, characterized by similar receptive fields and strong bidirectional

connectivity within the group, and different receptive fields and no lateral connectivity between groups (Fig. 5C). If the overall amplitude of plastic changes is small (compared to that found in the experiments) the pattern of lateral connectivity is stable and shows a few strong bidirectional connections in a sea of weak lateral connectivity. Unidirectional strong connections are nearly absent. If the amplitude and rate of plasticity is more realistic and in agreement with the data of Fig. 2, then the pattern of lateral connectivity changes between one snapshot and another one 5 seconds later, but the overall pattern is stable when averaged over 100s. In each snapshot, about half of the strong connections are bidirectional (Fig. 5H).

This is in striking contrast with the temporal coding paradigm. Neurons develop receptive fields similar to those seen with the rate-coding paradigm. As expected for temporal Hebbian learning rate [21] the receptive field slowly shifts over time. More importantly, amongst the lateral connections, strong reciprocal links are completely absent (Fig. 6). This suggests that temporal coding paradigms are reflected in the functional connectivity pattern by strong uni-directional connections whereas rate coding leads to strong bidirectional connections.

3 Discussion

Plasticity models over the last decades have primarily focused on questions of development of receptive fields and cortical maps [30], or memory formation [38]. Because traditional plasticity rules are rate models, the relation between coding and connectivity could not be studied. Our plasticity rule is formulated on the level of postsynaptic voltage. Since action potentials present large and narrow voltage peaks, they act as singular events in a voltage rule so that in the presence of spike our rule turns automatically into spike-timing dependent rule. Indeed, for spike coding (and in the absence of significant subthreshold voltage manipulations) our plasticity rule behaves like a STDP rule where triplets of spikes with pre-post-post or post-pre-post timing evoke LTP, whereas pairs with post-pre timing evoke LTD. Moreover, for rate coding where pre- and postsynaptic neurons fire with Poisson firing statistics, our plasticity rule presents structural similarities to the model of Bienenstock, Cooper, and Munro (BCM-model, [30]). Both our spiking rule and the rate-based BCM model require presynaptic activity in order to induce a change. Furthermore for our rule as well as for the simplest BCM rule (see [30]), the depression terms are linear and the potentiation terms are quadratic in the postsynaptic variables (i.e., the postsynaptic potential or the postsynaptic firing rate). Beyond these qualitative similarities, an approximate quantitative relation between the BCM model and our model can be constructed under appropriate assumptions. In this case the total weight change Δw in our model is proportional to $\nu^{\rm pre}\nu^{\rm post}(\nu^{\rm post}-\vartheta)$ where $\nu^{\rm pre}$ and $\nu^{\rm post}$ denotes the firing rate of a pre- and postsynaptic neurons, respectively and ϑ is a sliding threshold related to the ratio between the LTP and LTD inducing processes (see methods).

Due to its similarities to BCM, it is not surprising that our spike-based learning rule with sliding threshold is able to support independent component analysis (ICA) that has been hypothesized to underly receptive field development [30, 39]. In our experiments, the input consists of small patches of natural images using standard preprocessing [40]. Image patches are selected randomly and presented to the neuron for $T = 200$ms, which

is on the order of a fixation time between saccades [41]. Pixel intensities above an average grey value are converted to spike trains of ON-cells and and those below reference intensity to spikes in OFF-cells, using the relative intensity as the rate of a Poisson process. The spike trains from ON- and OFF-cells are the input to a cortical neuron. The synaptic weights undergo plasticity following our learning rule (Eq. 3). After learning, the weights exhibit a spatial structure that can be interpreted as a receptive field (Fig. 7). In contrast to the principal component analysis of the image patches (as for example implemented by Hebbian learning in linear neurons [42]), the receptive fields are *localized* (i.e. the region with significant weights does not stretch across the whole image patch). Development of localized receptive fields can be interpreted as a signature of ICA [40]. In contrast to most other ICA algorithms [43] our rule is biologically more plausible since it is consistent with a large body of plasticity experiments.

For a comparison of our model with experiments we have mainly focused on experiments in slices of visual cortex, but some of the results can also be related to work in hippocampus. First, as the model explicitly takes into account the postsynaptic membrane potential it can successfully reproduce the voltage dependence of LTP/LTD seen in experiments under depolarization of the postsynaptic membrane [14, 15]. Second, for classical STDP experiments such as [13, 23, 44], which have a stimulation protocol unambiguously defined in terms of pre- and postsynaptic spike times, the model gives a timing dependence reminiscent of the typical STDP function [13]. Moreover in contrast to standard STDP rules [21], more complicated effects such as the pairing frequency dependence [23] and burst-timing dependence plasticity [35] are qualitatively described. In addition the rule is expected to reproduce the triplet and quadruplet experiments in hippocampal slices [44] (data not shown), because for all STDP protocols the plasticity rule in this paper is similar to an earlier nonlinear STDP rule [32]. Deriving STDP rules from voltage dependence has been attempted before [45, 46]. However, since these earlier models use the momentary voltage [46] or its derivative [45], rather than a combination of momentary and averaged voltage as in our model, these earlier models cannot account for the broad range of nonlinear effects in STDP experiments or interaction of voltage and spike-timing. Our model shows similarities with LTP induction in the TagTriC model [47], but the TagTriC model focuses on the long-term stability of synapses, rather than spike timing dependence of the induction mechanism.

Our plasticity rule allows to explain experiments from two different laboratories by one single principle. Both the "potentiation is rescued by depolarization" [23] scenario (Fig. 2F) and that of burst-timing dependent LTP [35] (Fig. 3) show that LTP at low frequency is induced when the membrane is depolarized before the pre-post pairing. This depolarization can be due to a previous spike during a postsynaptic burst [35] or to a depolarization current. Our model is also consistent with results that LTP can be induced in distal synapses only if additional cooperative input or dendritic depolarization prevent failure of backpropagating action potentials [48]. A further unexpected result is that, with the set of parameters derived from visual cortex slice experiments, synapses fluctuate between strong and weak weights. This aspect is interesting in view of synapse mobility reported in imaging experiment [8].

There are, however, certain limitations to our plasticity rule. First, we did not address the problem of weight dependence of synaptic plasticity and simply assumed that weights can grow to a hard upper bound. Neverthe-

less, the rule can be easily changed to soft bounds [21] by changing the prefactors $A_{\rm LTP}, A_{\rm LTD}$ accordingly [47]. Second, short term plasticity [49] could be added for a better description of the plasticity phenomena occurring especially during high frequency protocols. Third, our plasticity rule describes only induction of potentiation or depression during the early phase of LTP/LTD [50]. Additional mechanisms need to be implemented in the model to describe the transition from early to late LTP/LTD [47, 51]. Finally, in modeling voltage-clamp experiments, we assume in our model a unique voltage throughout the whole neuron. In particular the dendrite is assumed to be equipotential to the soma. Yet, experiments controlling the voltage at the soma do not guarantee an equal or even fixed voltage at the synapse with respect to the soma. An obvious and promising improvement would be to use a multi-compartment neuron model (e.g. distinct compartments for the soma and dendrites). In the presented work we did not use a more sophisticated multi-compartment model as this would introduce a considerable number of new parameters making overfitting more likely to occur.

Our plasticity model leads to several predictions that could be tested in slice experiments. First, under the assumption of voltage clamp, our rule is linear in the presynaptic activities (see Methods). Thus the model predicts that in voltage clamp experiments the weight change is only dependent on the voltage and the *number* of presynaptic spikes but not on their exact timing (e.g., low frequency, tetanus, burst input should give the same result). Second, in the scenario where potentiation is rescued by depolarization, the amount of weight change should be the same whether a depolarizing current of amplitude B stops precisely when the postsynaptic spike is triggered or whether a current of slightly bigger amplitude B' stops a few milliseconds earlier. Third, multiple STDP experiments have shown that pre-post pairing (with 10 millisecond timing difference) repeated at 10Hz leads to potentiation [23]. In our plasticity model, LTP occurs in that case because the depolarizing spike-afterpotential of the last postsynaptic spike leads to an increase of the filtered membrane voltage just before the next postsynaptic spike. If this interpretation is correct, a hyperpolarizing current sufficient to cancel the spike afterpotential during 40 milliseconds should block LTP (note that this is different from blocking LTP by a hyperpolarizing current a few milliseconds *before* the next spike [23]). Alternatively cutting dendrites, i.e. dendrotomy [52] would sharpen the spike after potential.

The influence of STDP on temporal coding has been studied in the past primarily with respect to changes in the feedforward connections [21]. The effect of STDP on lateral connectivity has been studied much less [28, 29, 27]. We have shown in this paper that, because of STDP, coding influences the network topology, because different codes give different patterns of lateral connectivity. Our results are in contrast to standard STDP rules which always suppress short loops, and in particular bidirectional connections [36]. Our more realistic plasticity model shows that under a rate coding paradigm bidirectional connectivity and highly connected clusters with multiple loops are not only possible, but even dominant. It is only for temporal coding, that our biologically plausible rule leads to dominant unilateral directions. Our model also predicts that for a code consisting of synchronous firing events at low frequencies synapses decrease, consistent with earlier findings [27]. We speculate that the differences in coding between different brain areas could lead, even if the learning rule were exactly the same, to different network topologies. Our model predicts that experiments where cells in a recurrent network are repeatedly stimulated in a fixed order would decrease the fraction of strong bidirectional connections, whereas

a stimulation pattern where clusters of neuron fire at high rate during episodes of a few hundred milliseconds would increase this fraction. In this views it is tempting to connect the low degree of bidirectional connectivity in barrel cortex [6] to the bigger importance of temporal structure in whisker input [37], compared to visual input.

4 Acknowledgments

This work has been supported by the European projects FACETS as well as by the Swiss National Science Foundation.

5 Figure Captions

Figure 1: Illustration of the model. Synaptic weights react to presynaptic events (top) and postsynaptic membrane potential (bottom) A. The synaptic weight is decreased if a presynaptic spike x (green) arrives when the low pass filtered value $\bar{u}_-$ (magenta) of the membrane potential is above θ_- (dashed horizontal line) B. The synaptic weight is increased if the membrane potential u (black) is above a threshold θ^+ and the low pass filtered value of the membrane potential $\bar{u}_+$ (blue) higher than a threshold θ^- as well as the presynaptic low pass filter $\bar{x}$ (orange) non zero. C. Step current injection makes the postsynaptic neuron fire at 50Hz in the absence of presynaptic stimulation (membrane potential u in black). No weight change is observed. Note the depolarizing spike-afterpotential consistent with experimental data D., reproduced from [23]. E-H. Voltage clamp experiment. A neuron receives weak presynaptic stimulation of 2Hz during 50s while the postsynaptic voltage is clamped to values between -60mV and 0mV. E-G. Schematic drawing of the trace $\bar{x}$ (orange) of the presynaptic spike train (green) as well as the voltage (black) and the synaptic weight (blue) for the experimental conditions E. Hyperpolarization F. Slight depolarization and G. Large depolarization. H. The weight change as a function of clamped voltage using the standard set of parameters for visual cortex data (blue line, voltage paired with 25 spikes at the synapse). With a different set of parameters the model fits experimental data (red circles) in hippocampal slices [15], see methods for details.

Figure 2: A-B. Simulated STDP experiments. A. Spike-timing dependent learning window. The change of the synaptic weight is shown for different time intervals T between the presynaptic and the postsynaptic spike using 60 presynaptic/postsynaptic spike pairs at 20Hz. B. Weight change as a function of repetition frequency for 5 spike pairs at frequency ρ with a time delay of +10ms (pre-post, blue) and -10ms (post-pre, red), repeated 15 times at 0.1Hz (only 10 times for frequency of ρ=0.1Hz). Weight changes are shown as a function of the frequency, dots represent the data taken from Sjöström et al. [23] and lines the plasticity model simulation. C-I. Interaction of voltage and STDP. C-E. Schematic induction protocols (green: presynaptic input, black: postsynaptic current, blue: evolution of synaptic weight). C. Low-Frequency Potentiation is rescued by depolarization [23]. Low frequency (0.1Hz) pre-post spike pairs yield LTP if a 100ms-long depolarized current is injected around the pairing. D. LTP fails in the previous scenario if an additional brief hyperpolarized pulse is applied 14-ms before postsynaptic spike so that voltage is brought to rest. E. Hyperpolarization preceding action potential prevents potentiation. Sjöström et al. [23] show that high frequency (40Hz) pairing leads to LTP. However, when a constant hyperpolarizing current is applied on top of the short pulses inducing the spikes, no weight change is measured. F. The simulated postsynaptic voltage u (black) following protocol A. is shown as well as the temporal averages $\bar{u}_-$ (magenta) and $\bar{u}_+$ (blue). The presynaptic spike time is indicated by the green arrow. Using the model Eq. 3 this setting results in potentiation. G. Same as F, but following protocol D. No weight change is measured. H. Same as F., but following protocol E. No weight change is measured. I. Histogram summarizing the normalized synaptic weight of the simulation (bar) and the experimental data [23] (dot, blue bar=variance) 0.1Hz pairing (control 1); 0.1Hz pairing with the depolarization (protocol C.); 0.1Hz pairing with the depolarization and brief hyperpolarization (protocol D.); 40Hz pairing (control 2); 40Hz pairing with the constant hyperpolarization (protocol E.). The parameters are summarized in Table 1B.

Figure 3: Burst-timing-dependent plasticity. One presynaptic spike is paired with a burst of postsynaptic spikes. This pairing is repeated 60 times at 0.1Hz. A. Normalized weight is shown as a function of the number of postsynaptic spikes (1,2,3) at 50Hz. (dots: data from [35], crosses: simulation). The presynaptic spike is paired +10ms before the first postsynaptic spike (blue) or -10ms after (red). B. Normalized weight as a function of the frequency between the three postsynaptic action potentials (dot: data, line: simulation; blue: pre-post, red: post-pre). C. Normalized weight as a function of the timing between the presynaptic spike and the first postsynaptic spike of a 3-spike burst at 50Hz (dot: data, line: simulation). A hard upper bound has been set to 250% normalized weight.

Figure 4: Weight evolution in a all-to-all connected network of 10 neurons. A. Rate code: Neurons fire at different frequencies, neuron 1 at 2Hz, neuron 2 at 4Hz... neuron 10 at 20Hz. The weights (bottom) averaged over 100s show that neurons with high firing rates develop strong bidirectional connections (light blue: weak connections (under 2/3 of the maximal value); yellow: strong unidirectional connections (above 2/3 of the maximal value); brown: strong bidirectional connections). The cluster is schematically represented on top ("after"). B. Temporal code: Neurons fire successively every 20ms (neuron 1 then 20ms later neuron 2, then 3..). Connections (bottom) are unidirectional with strong connections from presynaptic neuron with index n (vertical axis) to postsynaptic neuron with index n+1, n+2 and n+3 leading to a ring-like topology (top: schematic). C. D. Same but with standard STDP rule [17, 26, 21]. Bidirectional connections are impossible.

Figure 5: Plasticity during rate coding. A network of 10 excitatory neurons is connected to 3 inhibitory neurons and receives feedforward inputs from 500 Poisson spike trains with a Gaussian profile of firing rates. The center of the Gaussian is shifted randomly every 100ms A. The schematic figure shows the network before and after the plasticity experiment. B-E. Learning with small amplitudes. Model parameters are taken from table 1B (visual cortex data) except for the amplitudes A_{LTP} and A_{LTD} which are reduced by a factor 100. B. Mean feedforward weights (left) and recurrent excitatory weights (right) averaged over 100s. The grey level graph for the feedforward weights (left) indicates that neurons develop receptive fields that are localized in the input space. The recurrent weights (right) are classified into: light blue - weak (less than 2/3 of the maximal weight), yellow - strong (more than 2/3 of the maximal weight) unidirectional, brown - strong reciprocal connections. The diagonal is white, since self-connections do not exist in the model. C. Same as (B) but for the sake of visual clarity the index of neurons is reordered so that neurons with similar receptive fields have adjacent numbers, highlighting that neurons with similar receptive fields (e.g., neurons 1 to 4) have strong bilateral connections. D. Three snap shots of the recurrent connections taken 5s apart indicating that recurrent connections are stable. E. Histogram of reciprocal, unidirectional and weak connections in the recurrent network averaged over 100s as in (B). The total number of weight fluctuations during 100s is 79 (noted on the figure). The histogram shows an average of 10 repetitions (errorbars are the standard deviation). F-I. Rate code during learning with normal amplitudes. Same network as before but standard set of parameters (table 1B, visual cortex). F. Receptive fields are localized; G. Reordering allows to visualize that the strong bidirectional give rise to clusters of neurons. These clusters are stable when averaged over 100 seconds, but H connections can change from one time step to the next. I. The percentage of reciprocal connections is high, but because of fluctuations (fluc) more than 1000 transitions between strong unidirectional to strong bidirectional or back occur during 100 seconds.

Figure 6: Temporal coding paradigm. The setting is the same as in Fig. 5 (parameters from table 1B, visual cortex) but the input patterns are moved successively every 20ms, corresponding to a step-wise motion of the Gaussian stimulus profile across the input neurons. A. The schematic figure shows the network before and after the plasticity experiment. B. Receptive fields are localized, but in the recurrent network no reciprocal connections appear. C. Reordering of neurons shows that the network develops a ring-like structure with strong unidirectional connections from neuron 8 (vertical axis) to neuron 7 and 6 (horizontal axis); from neuron 7 to neuron 6, 5, and 4; from neuron 4 to neuron 3, 2, and 1 etc. D. Some of the strong unilateral connections appear or disappear from one time step to the next, but the ring-like network structure persists, since the lines just below the diagonal are much more populated than the line above the diagonal. E. Reciprocal connections are completely absent, but unidirectional connections fluctuate several times between 'weak' and 'strong' during 100s.

Figure 7: A small patch of 16x16 pixels is chosen from the whitened natural images benchmark [40]. The patch is selected randomly and is presented as input to 512 neurons for 200ms. The positive part of the image is used as the firing rate to generate Poisson spike trains of the 256 "ON" inputs and the negative one for the 256 "OFF" inputs. B. The weights after convergence are shown for the "ON" inputs and the "OFF" inputs rearranged on a 16x16 image. The filter is calculated by subtracting the "OFF" weights from the "ON" weights. The filter is localized and bimodal, corresponding to an oriented receptive field.

Table 1: A. Parameters for the neuron model. B. Plasticity rule parameters for the various experiments. VC stands for Visual Cortex cells (for experimental details see [23], * standard set of parameters), SC for Somatosensory Cortex cells (see [35]) and HP for Hippocampal cells (see [15]). Bold numbers indicate the free parameters fitted to experimental data. Other parameters are set in advance to values based on the literature.

6 Methods

6.1 Neuron Model

In contrast to standard models of STDP, the plasticity model presented in this paper involves the postsynaptic membrane potential $u(t)$. Hence, predicting the weight change in a given experimental paradigm requires a neuron model that describes the temporal evolution of $u(t)$. For this purpose we chose the adaptive Exponential Integrate-and-Fire (AdEx) model [53] with an additional current describing the depolarizing spike after potential [54]. The neuron model is described by a voltage equation:

$$C\frac{d}{dt}u = -g_L(u - E_L) + g_L\Delta_T \exp\left(\frac{u - V_T}{\Delta_T}\right) - w_{ad} + z + I$$

where C is the membrane capacitance, g_L the leak conductance, E_L the resting potential and I the stimulating current. The exponential term describes the activation of a rapid sodium current. The parameter Δ_T is called the slope factor and V_T the threshold potential [53]. A hyperpolarizing adaptation current is described by the variable w_{ad} with dynamics

$$\tau_{w_{ad}}\frac{d}{dt}w_{ad} = a(u - E_L) - w_{ad},$$

where $\tau_{w_{ad}}$ is the time constant of the adaption of the neuron. Upon firing the variable u is reset to a fixed value V_{reset} whereas w_{ad} is increased by an amount b. The main difference to the Izhikevich model [55] is that

the voltage is exponential rather than quadratic allowing a better fit to data [54]. The spike afterpotential of the cells used in typical STDP experiments [23] have a long depolarizing spike after potential. We therefore add an additional current z which is set to a value I_{sp} immediately after a spike occurs and decays otherwise with a time constant τ_z.

$$\tau_z \frac{d}{dt} z = -z,$$

Finally, refractoriness is shown in pyramidal cells [54] and therefore is modeled with the adaptive threshold V_T. Therefore V_T is set to $V_{T_{max}}$ after a spike and decays to $V_{T_{rest}}$ with a time constant τ_{V_T} as measured in [54], i.e.

$$\tau_{V_T} \frac{d}{dt} V_T = -(V_T - V_{T_{rest}}).$$

Parameters for the neuron model are taken from [53] for the AdEx, τ_z is set to 40ms in agreement with [23, 54] and kept fixed throughout all simulations (see table 1A).

6.2 Plasticity Model

Since synaptic depression and potentiation take place through different pathways [33] our model exhibits separate additive contributions to the plasticity rule, one for LTD and another one for LTP.

For the LTD part, we assume that presynaptic spike arrival at synapse i induces depression of the synaptic weight w_i by an amount $-A_{\rm LTD}\,[\overline{u}_-(t) - \theta_-]_+$ that is proportional to the average postsynaptic depolarization $\overline{u}_-$. The brackets $[\,]_+$ indicate rectification, i.e. any value $\bar{u}_- < \theta_-$ does not lead to a change and implement experimental findings showing that postsynaptic depolarization should exceed a certain value θ^- to establish depression of the synapse [14] (see Fig. 1H). The quantity $\overline{u}_-(t)$ is an exponential low-pass filtered version of the postsynaptic membrane potential $u(t)$ with a time constant τ_-:

$$\tau_- \frac{d}{dt} \overline{u}_-(t) = -\overline{u}_-(t) + u(t).$$

The variable $\bar{u}_-$ is an abstract variable which could, for instance, reflect the level of calcium concentration [24] or the release of endocannabinoids [56], though such an interpretation is not necessary for our rule. Since the presynaptic spike train is described as a series of short pulses at time t_i^n where i is the index of the synapse and n an index that counts the spike, $X_i(t) = \sum_n \delta(t - t_i^n)$, depression is modeled as the following update rule, see also Fig. 1:

$$\frac{d}{dt} w_i^- = -A_{\rm LTD}(\bar{\bar{u}})\, X_i(t)\, [\overline{u}_-(t) - \theta_-]_+ \qquad \text{if } w_i > w_{\rm min}, \tag{1}$$

where $A_{\rm LTD}(\bar{\bar{u}})$ is an amplitude parameter that is under the control of homeostatic processes [34]. For slice experiment the parameter has a fixed value extracted from experiment. For network simulations, we make it depend on the mean depolarization $\bar{\bar{u}}$ of the postsynaptic neuron, averaged over a time scale of 1 second. Eq. 1

is a simple method to implement homeostasis; other methods such as weight rescaling would also be possible [34].

For the LTP part, we assume that each presynaptic spike at the synapse w_i increases the trace $\bar{x}_i(t)$ of some biophysical quantity, which decays exponentially with a time constant τ_x in the absence of presynaptic spikes, similar to previous work [17, 32]. The temporal evolution of $\bar{x}_i(t)$ is described by:

$$\tau_x \frac{d}{dt}\bar{x}_i(t) = -\bar{x}_i(t) + X_i(t),$$

where X_i is the spike train defined above. The quantity $\bar{x}_i(t)$ could for example represent the amount of glutamate bound to postsynaptic receptors [32] or the number of NMDA receptors in an activated state. The potentiation of w_i is modeled by the following expression, which is proportional to the trace $\bar{x}_i(t)$ (see also Fig. 1):

$$\frac{d}{dt}w_i^+ = +A_{\rm LTP}\,\bar{x}_i(t)\,[u(t)-\theta_+]_+\,[\overline{u}_+(t)-\theta_-]_+ \qquad \text{if } w_i < w_{\rm max}. \tag{2}$$

Here, $A_{\rm LTP}$ is a free amplitude parameter fitted to the data and $\overline{u}_+(t)$ is another low-pass filtered version of $u(t)$ similar to $\overline{u}_-(t)$ but with a shorter time constant τ_+ around 10ms. Thus positive weight changes can occur if the momentary voltage $u(t)$ surpasses a threshold θ_+ and, at the same time the average value $\overline{u}_+(t)$ is above θ_-.

The final rule used in the simulation is described by the equation

$$\frac{d}{dt}w_i = -A_{\rm LTD}(\bar{\bar{u}})\,X_i(t)\,[\overline{u}_-(t)-\theta_-]_+ + A_{\rm LTP}\,\bar{x}_i(t)\,[u(t)-\theta_+]_+\,[\overline{u}_+(t)-\theta_-]_+, \tag{3}$$

combined with hard bounds $0 \le w_i \le w_{max}$. For network simulation, $A_{\rm LTD}(\bar{\bar{u}}) = A_{\rm LTD}\,\frac{\bar{\bar{u}}^2}{u_{ref}^2}$ where u_{ref}^2 is a reference value.

6.3 Parameters and Data Fitting

For the plasticity experiments in slices, we take $\bar{\bar{u}} = u_{ref}$ as fixed and fit the parameters $A_{\rm LTD}$. The total number of parameters of the plasticity model is then 7. For all data sets, except the one taken from [15], the threshold θ^- is set to the resting potential and θ^+ to the firing threshold of the AdEx model, i. e. $\theta^- = -70.6$mV and $\theta^+ = -45.3$mV. The remaining five parameters τ_x, τ_-, τ_+, $A_{\rm LTD}$ and $A_{\rm LTP}$ are fitted to each data set individually by the following procedure. We calculate the theoretically predicted weight change $\Delta w_i^{{\rm th},j}$ by integrating (analytically or numerically) Eq. (**3**), for a given experimental protocol j, as a function of the free parameters. We then estimate the free parameters by minimizing the mean-square error E between the theoretical calculations and the experimental data $\Delta w_i^{{\rm exp},j}$:

$$E = \sum_j \left(\Delta w_i^{{\rm th},j} - \Delta w_i^{{\rm exp},j}\right)^2.$$

For the data set in hippocampus [15], we also fit the two parameters θ^- and θ^+ since completely different preparations and cell type were used. Moreover for this data set, the time constant τ_x is taken from physiological measurements given in [13] and fixed to the values of 16ms. The parameters for the various experiments are summarized in table 1B.

6.4 Protocols and mathematical methods

Voltage clamp experiment. (Fig. 1H) The postsynaptic membrane potential was switched in the simulations to a constant value u_{clamp} chosen from -80mV to 0mV while presynaptic fibers were stimulated with either 25 (blue line) or 100 pulses (red line) at 50Hz. Due to voltage clamping, the actual value of the voltage u itself and the low-pass filtered versions $\bar{u}$ are constant and equal to u_{clamp}. Hence, the synaptic plasticity rule becomes $\frac{d}{dt}w_i = -A_{\text{LTD}}\, X_i(t)\, [u_{\text{clamp}} - \theta_-]_+ + A_{\text{LTP}}\, \bar{x}_i(t)\, [(u_{\text{clamp}} - \theta_-)(u_{\text{clamp}} - \theta_+)]_+$.

Frequency dependence experiment. (Fig. 2B) Presynaptic spikes in the simulation were paired with postsynaptic spikes that were either advanced by +10ms or delayed by -10ms with respect to the presynaptic spike. This pairing was repeated 5 times with different frequencies ranging from 0.1 to 50Hz. These 5 pairings were repeated 15 times at 0.1Hz. However, the 5 pairing at 0.1Hz were repeated only 10 times to mimic the experimental protocol [23].

Burst-timing-dependent plasticity. (Fig. 3A) The presynaptic spike is paired Δt =+10ms before (or Δt =-10ms after) 1, 2 or 3 postsynaptic spikes. The frequency of the burst is 50Hz. The neuron receives 60 pairings at a frequency of 0.1Hz. Fig. 3B: The presynaptic spike is paired with a burst of 3 action potentials (Δt =+10ms and -10ms), while the burst frequency varies from 20 to 100Hz. Fig. 3C: A presynaptic spike is paired with a burst of 3 postsynaptic action potentials with burst frequency of 50Hz. The time Δt between the presynaptic spike and the first postsynaptic action potential varies from -80 to 40 ms. For a detailed description of the experiments see [35].

Poisson input for functional scenarios.(Fig. 4-7) Poisson inputs are used in all the following experiments. They are generated by a stochastic process where the spike is elicited with a stochastic intensity ν.

Relation between connectivity and coding: Toy model. (Fig. 4) Weights of ten all-to-all connected neurons are initialized at 1, bounded between 0 and 3. Weights evolve with the voltage-based rule with homeostasis (Eq. 3) for 100s. The model is compared to a canonical pair-based STDP model written as $\frac{d}{dt}w_i = -A^{\text{pair}}_{\text{LTD}}\, X_i\, \bar{y} + A^{\text{pair}}_{\text{LTP}}\, \bar{x}_i\, Y$, where Y is the postsynaptic spike train defined the same way as the presynaptic spike train X_i with a filter of the postsynaptic spikes $\bar{y}$ similar to $\bar{x}_i$. The parameters are chosen $A^{\text{pair}}_{\text{LTD}} = A^{\text{pair}}_{\text{LTP}} = 1e^{-5}$ for the amplitudes and τ_x for the time constant of $\bar{x}_i$ as well as for the time constant of the postsynaptic low-pass filter $\bar{y}$. Rate code: Neuron 1 fire at 2Hz, neuron 2 at 4Hz... neuron 10 at 20Hz

following a Poisson statistics, i.e. short current pulses are injected to make the neuron fire with Poisson statistics at this frequency. The neurons have different reference values from $u^2_{ref} = 60$ to 600mV^2. Temporal code: Neurons fire successively every 20ms, first neuron 1 fires then 20ms later neuron 2 then... 10 then 1 etc, in a loop. The neurons have a reference value set to $u^2_{ref} = 60\text{mV}^2$.

Rate coding in network simulation. (Fig. 5) Five hundred presynaptic Poisson neurons with firing rates ν_i^{pre} ($1 \leq i \leq 500$) are connected to 10 postsynaptic excitatory neurons. The inputs rates ν_i^{pre} follow a Gaussian profile, i. e. $\nu_i^{\text{pre}} = A \cdot \exp(-(i-\mu)^2/(2\sigma^2))$, with variance $\sigma = 10$ and amplitude $A = 30\text{Hz}$. The center μ of the Gaussian shifts randomly every 100ms between 10 different positions equally distributed. Circular boundary conditions are assumed, i.e. neuron $i = 500$ is considered as neighbor of $i = 1$. Synaptic weights of the feedforward connections are initialized randomly (uniformly in [0.5,2]) and hard bound are set to 0 and 3. The 10 excitatory neurons are all to all recurrently connected with a starting synaptic weight of 0.25 (hard bounds set to 0 and 0.75). In addition, 3 inhibitory neurons are randomly driven by 8 excitatory neurons and project on 6 excitatory neurons, also chosen randomly. Those random connections are fixed and have a weight equal to 1. The reference value is set to $u^2_{ref} = 60\text{mV}^2$ and the simulation time to 1000s. Parameters are normally chosen as in table 1B, visual cortex data, except for Fig. 5 B-E, where A_{LTP} and A_{LTD} where reduced by a factor 100.

Temporal coding in network simulation. (Fig. 6) Same setting than rate code but the patterns are presented for 20ms successively (from center position 500, to 450, to 400 etc in a circular manner). The reference value has been set to $u^2_{ref} = 80\text{mV}^2$.

ICA-like computation - Orientation selectivity with natural images. (Fig. 7) Ten natural images have been taken from the benchmark of Olshausen et al. [40]. A small patch of 16 by 16 pixels from any of the images is randomly chosen every 200ms. After prewhitening, the inputs for the "ON" ("OFF") image are Poisson spike trains generated by the positive (negative) part of the patch (with respect to a reference grey value reflecting the ensemble mean) with maximum frequency of 50Hz. The 2x16x16 inputs are connected to one postsynaptic neuron. The initial weights are set randomly between 0 and 2 and hard bounds are set between 0 and 3. The connections follow the synaptic rule (Eq. 3), where the reference value is set to $u^2_{ref} = 50\text{mV}^2$. Parameters are chosen as in table 1B (visual cortex data) but A_{LTP} and A_{LTD} where reduced by a factor 10. Every 20 s an extra normalization is applied to equalize the norm of the "ON" weights to the one of the "OFF" weights [31].

References

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

[2] Y. Fregnac and D. Shulz. Activity-dependent regulation of receptive field properties of cat area 17 by supervised hebbian learning. *J. Neurobiol*, 41:69–82, 1999.

[3] Robert C. Froemke, Michael M. Merzenich, and Christoph E. Schreiner. A synaptic memory trace for cortical receptive field plasticity. *Nature*, 450:425–429, 2007.

[4] G. H. Recanzone, C. E. Schreiner, and M. M. Merzenich. Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *The Journal of Neuroscience*, 13:87–103, 1993.

[5] S. Song, P.J. Sjöström, M. Reigl, S. Nelson, and D.B. Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, 3:507–519, 2005.

[6] S. Lefort, C. Tomm, J.C.F. Sarria, and C.C.H. Petersen. The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61:301–316, 2009.

[7] W Denk and H Horstmann. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol*, 2(11):e329. doi:10.1371/journal.pbio.0020329, 2004.

[8] R. Yuste and T. Bonhoeffer. Genesis of dendritic spines: insights from ultrastructural and imaging studies. *Nat Rev Neurosci*, 5(1):24–34, 2004.

[9] J. T. Trachtenberg, B. E. Chen, G. W. Knott, G. Feng, J. R. Sanes, E. Welker, and K. Svoboda. Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*, 420:788–794, 2002.

[10] D. O. Hebb. *The Organization of Behavior*. Wiley, New York, 1949.

[11] Robert C. Malenka and Mark F. Bear. LTP and LTD: An embarassment of riches. *Neuron*, 44:5–21, 2004.

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

[13] G.-q. Bi and M.-m. Poo. Synaptic modification of correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci.*, 24:139–166, 2001.

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

[15] A. Ngezahayo, M. Schachner, and A. Artola. Synaptic activation modulates the induction of bidirectional synaptic changes in adult mouse hippocamus. *J. Neuroscience*, 20:2451–2458, 2000.

[16] S. M. Dudek and M. F. Bear. Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus. *J. Neuroscience*, 13:2910–2918, 1993.

[17] W. Gerstner, R. Kempter, J.L. van Hemmen, and H. Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–78, 1996.

[18] P.D. Roberts and C.C. Bell. Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. *Computational Neuroscience*, 9:67–83, 2000.

[19] R. Legenstein, C. Naeger, and W. Maass. What can a neuron learn with spike-timing dependent plasticity. *Neural Computation*, 17:2337–2382, 2005.

[20] R. Guyonneau, R. VanRullen, and S.J. Thorpe. Neurons tune to the earliest spikes through stdp. *Neural Computation*, 17(4):859–879, 2005.

[21] W. Gerstner and W. K. Kistler. *Spiking Neuron Models*. Cambridge University Press, Cambridge UK, 2002.

[22] John Lisman and Nelson Spruston. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nature Neuroscience*, 8(7):839–841, 2005.

[23] P.J. Sjöström, G.G. Turrigiano, and S.B. Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32:1149–1164, 2001.

[24] H. Z. Shouval, M. F. Bear, and L. N. Cooper. A unified model of nmda receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA*, 99:10831–10836, 2002.

[25] J.E. Lisman and A.M. Zhabotinsky. A model of synaptic memory: A CaMKII/PP1 switch that potentiates transmission by organizing an AMPA receptor anchoring assembly. *Neuron*, 31:191–201, 2001.

[26] S. Song and L. F. Abbott. Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, 32:339–350, 2001.

[27] E. V. Lubenov and A. G. Siapas. Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron*, 58:118–131, 2008.

[28] A. Morrison, A. Aertsen, and M. Diesmann. Spike-timing dependent plasticity in balanced random networks. *Neural Computation*, 19:1437–1467, 2007.

[29] Eugene M. Izhikevich and Gerald M. Edelman. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, 105:3593–3598, 2008.

[30] L.N. Cooper, N. Intrator, B.S. Blais, and H. Z. Shouval. *Theory of cortical plasticity.* World Scientific, Singapore, 2004.

[31] K. D. Miller. A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity dependent competition between ON- and OFF-center inputs. *J. Neurosci.*, 14:409–441, 1994.

[32] J. P. Pfister and W. Gerstner. Triplets of spikes in a model of spike timing-dependent plasticity. *J. Neursci.*, 26:9673–9682, 2007.

[33] D.H. O'Connor, G.M. Wittenberg, and S.S.H. Wang. Dissection of Bidirectional Synaptic Plasticity Into Saturable Unidirectional Processes. *Journal of Neurophysiology*, 94:1565–1573, 2005.

[34] G.G. Turrigiano and S.B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5:97–107, 2004.

[35] Thomas Nevian and Bert Sakmann. Spine ca2+ signaling in spike-timing-dependent plasticity. *J. Neurosci.*, 26(43):11001–11013, 2006.

[36] J. Kozloski and G. A. Cecchi. Topological effects of spike timing-dependent plasticity. *arxiv.org*, abs:0810.0029, 2008.

[37] S. P. Jadhav, J. Wolfe, and D. E. Feldman. Sparse temporal coding of elementary tactile features during active whisker sensation. *Nature Neuroscience*, page doi:10.1038/nn.2328, 2009.

[38] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, 79:2554–2558, 1982.

[39] B. Blais, H. Shouval, and L. Cooper. Receptive field formation in natural scene environments: comparison of single-cell learning rules. *Neural Computation*, 10:1797–1813, 1998.

[40] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609, 1996.

[41] S. Martinez-Conde, S. Macknik, and D. Hubel. The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience*, 5:229–240, 2004.

[42] E. Oja. A simplified neuron model as a principal component analyzer. *J. Mathematical Biology*, 15:267–273, 1982.

[43] A. Hyvaerinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. Wiley-Interscience, 2001.

[44] H.-X. Wang, R.C. Gerkin, D.W. Nauen, and G.-Q. Wang. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nature Neuroscience*, 8:187–193, 2005.

[45] A. Saudargiene, B. Porr, and F. Wörgötter. How the shape of pre- and postsynaptic signals can influence stdp: A biophysical model. *Neural Computation*, 16:595–626, 2003.

[46] J.M. Brader, W. Senn, and S. Fusi. Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Computation*, 19:2881–2912, 2007.

[47] Claudia Clopath, Lorric Ziegler, Eleni Vasilaki, Lars Bsing, and Wulfram Gerstner. Tag-trigger-consolidation: A model of early and late long-term-potentiation and depression. *PLoS Comput Biol*, 4(12), Dec 2008.

[48] P.J. Sjöström and M. Häusser. A Cooperative Switch Determines the Sign of Synaptic Plasticity in Distal Dendrites of Neocortical Pyramidal Neurons. *Neuron*, 51(2):227–238, 2006.

[49] M. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Academy of Sci., USA*, 94:719–723, 1997.

[50] U. Frey and R.G.M. Morris. Synaptic tagging and long-term potentiation. *Nature*, 385:533 – 536, 1997.

[51] A.B. Barrett, G.O. Billings, R.G.M. Morris, and M.C.W. van Rossum. State based model of long-term potentiation and synaptic tagging and capture. *Plos Comp Biol*, 5(1):e1000259. doi:10.1371/journal.pcbi.1000259, 2009.

[52] JM. Bekkers and M. Häusser. Targeted dendrotomy reveals active and passive contributions of the dendritic tree to synaptic integration and neuronal output. *Proc. Natl. Acad. Sci. USA*, 104(27):11447–11452, 2007.

[53] R. Brette and W. Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.*, 94:3637 – 3642, 2005.

[54] L Badel, S Lefort, R Brette, CC Petersen, W Gerstner, and MJ Richardson. Dynamic i-v curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *J Neurophysiol*, 99:656 – 666, 2008.

[55] E.M. Izhikevich. Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15:1063–1070, 2004.

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

Figure-1(Clopath)

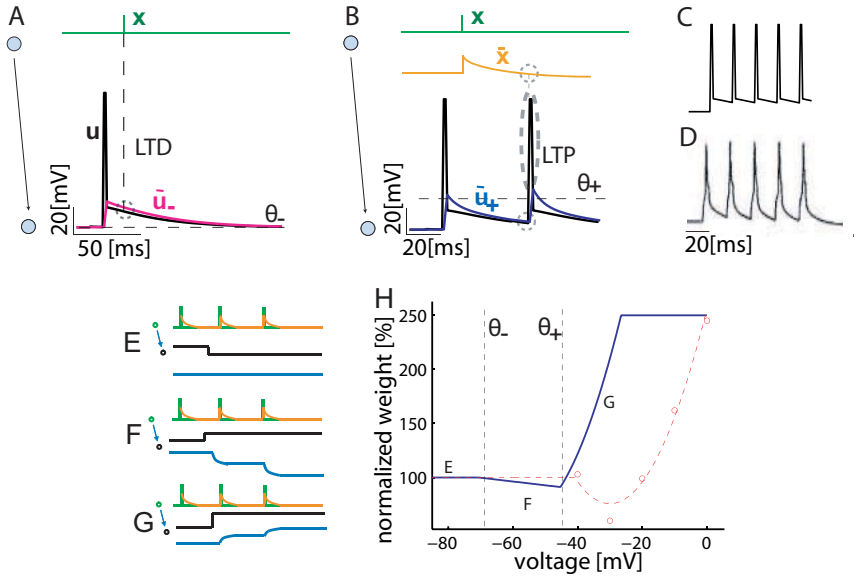

Figure-2(Clopath)

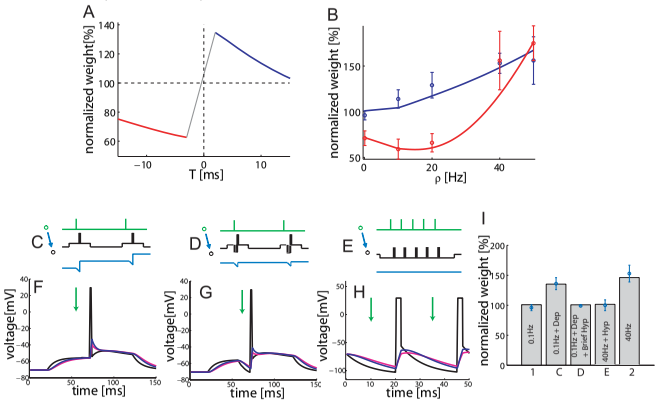

Figure-3(Clopath)

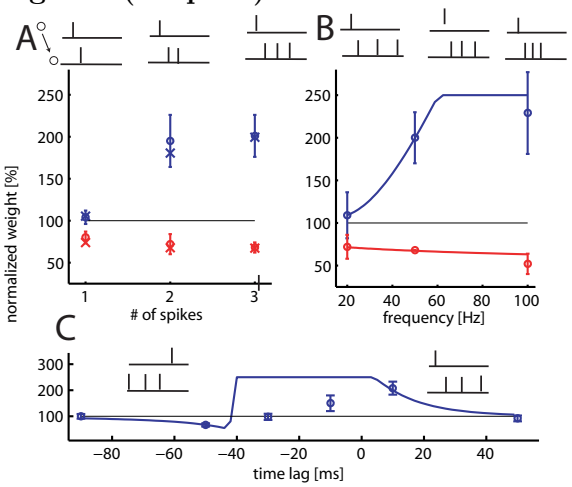

Figure-4(Clopath)

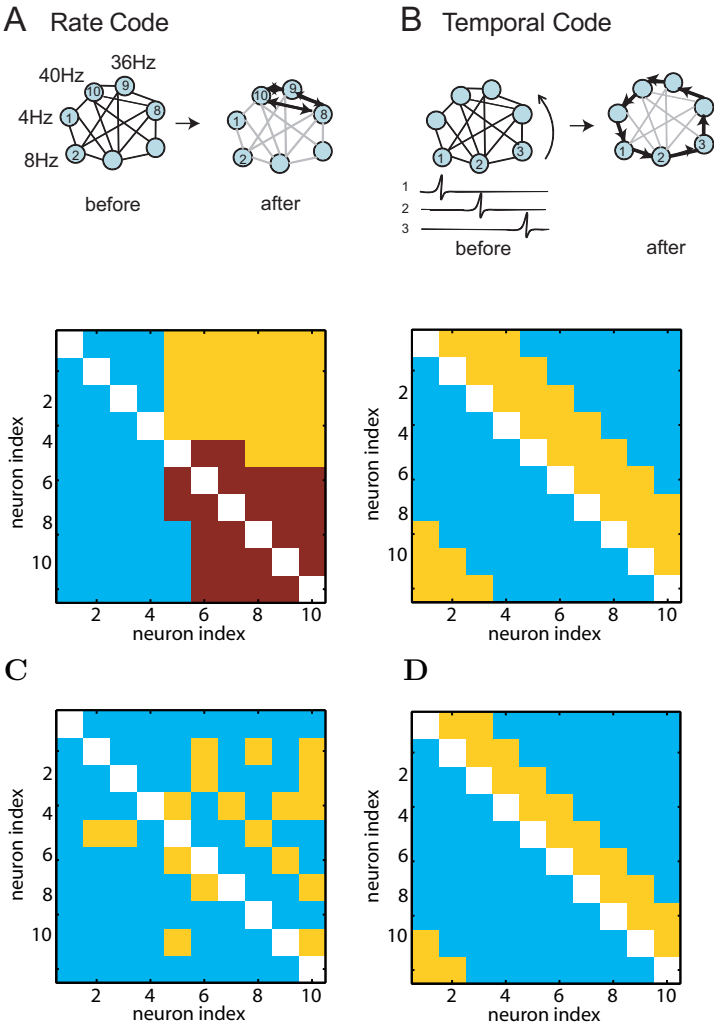

Figure-5(Clopath)

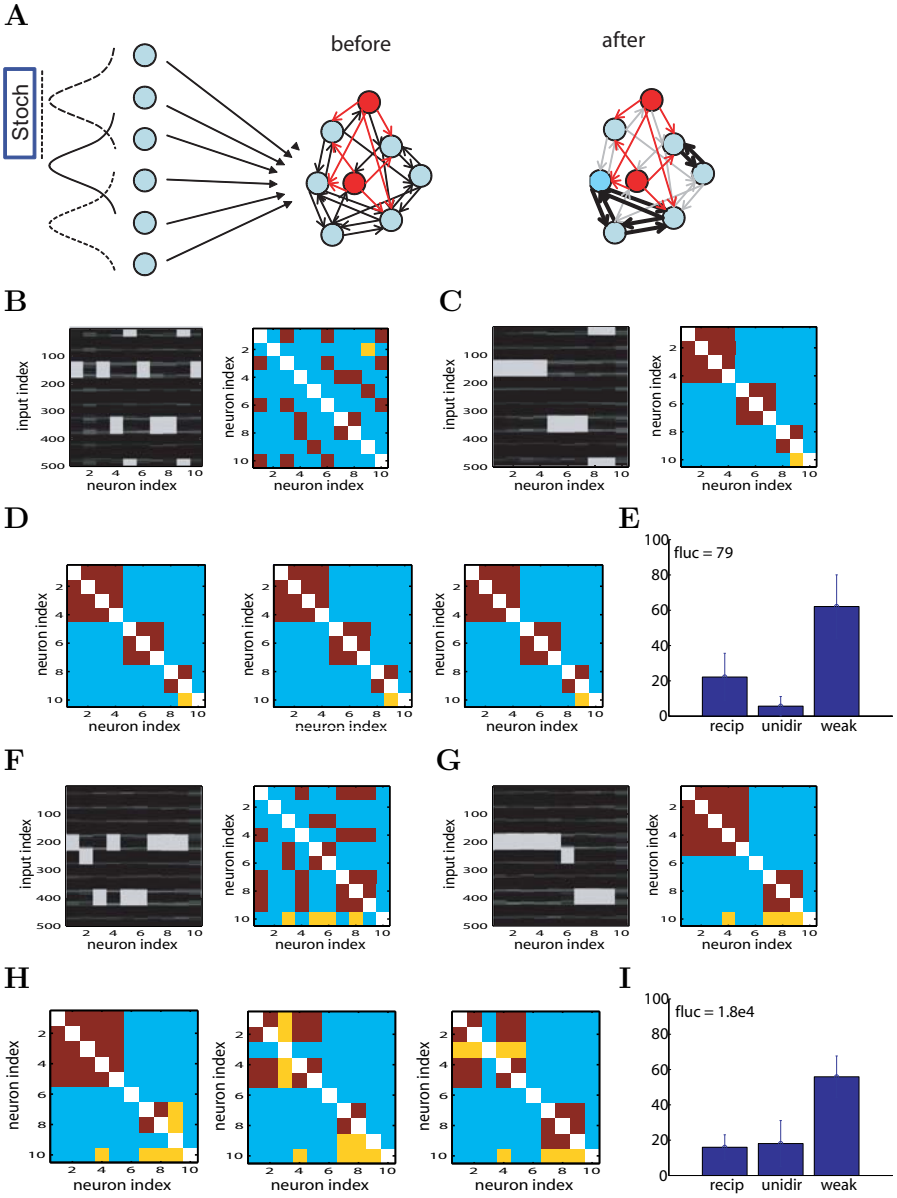

Figure-6(Clopath)

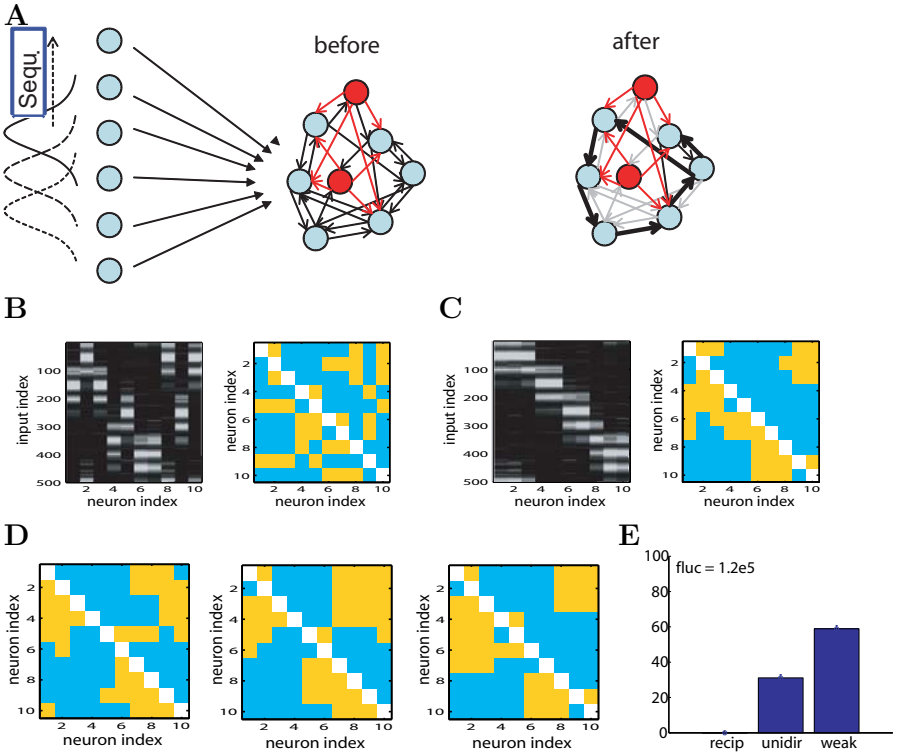

Figure-7(Clopath)

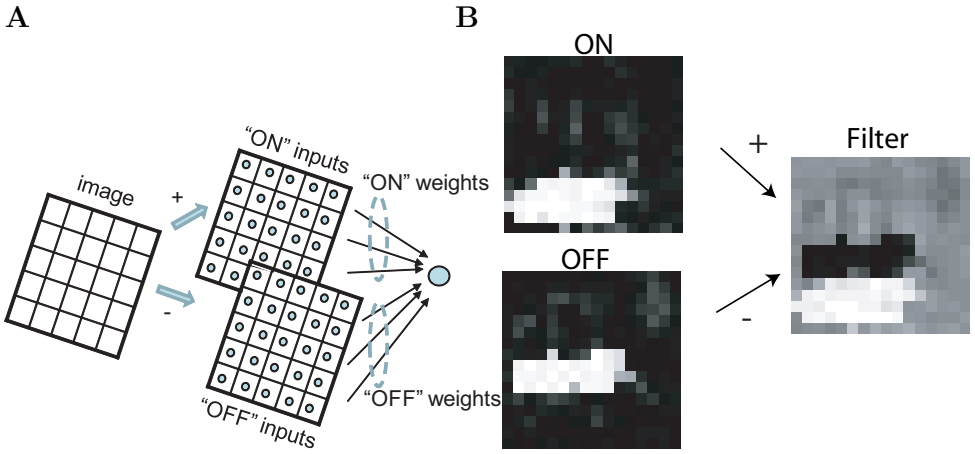

Table-1(Clopath)

A

Parameters	Value
C - membrane capacitance	281pF
g_L - leak conductance	30nS
E_L - resting potential	-70.6mV
Δ_T - slope factor	2mV
$V_{T_{rest}}$ - threshold potential at rest	-50.4mV
$\tau_{w_{ad}}$ - adaptation time constant	144ms
a - subthreshold adaptation	4nS
b - spike triggered adaptation	80.5pA
I_{sp} - spike current after a spike	400nA
τ_z - spike current time constant	40ms
τ_{V_T} - threshold potential time constant	50ms
$V_{T_{max}}$ - threshold potential after a spike	$-30.4mV$

B

Exper.	$\theta_-(mV)$	$\theta_+(mV)$	$A_{\mathrm{LTD}}(mV)^{-1}$	$A_{\mathrm{LTP}}(mV)^{-2}$	$\tau_x(ms)$	$\tau_-(ms)$	$\tau_+(ms)$
VC*	-70.6	-45.3	$\mathbf{14e^{-5}}$	$\mathbf{8e^{-5}}$	**15**	**10**	**7**
SC	-70.6	-45.3	$\mathbf{21e^{-5}}$	$\mathbf{67e^{-5}}$	**15**	**8**	**5**
HP	**−41**	**−38**	$\mathbf{38e^{-5}}$	$\mathbf{2e^{-5}}$	16		

2.2 Appendix 2

Mayr 2008 C. Mayr, J. Partzsch, and R. Schffny, BCM and Membrane Potential: Alternative Ways to Timing Dependent Plasticity, 15th International Conference on Neural Information Processing (ICONIP 2008), Springer LNCS, vol. 5506, pp. 137-144, 2009

BCM and Membrane Potential: Alternative Ways to Timing Dependent Plasticity

Johannes Partzsch*, Christian Mayr**, and Rene Schüffny

Chair for Parallel VLSI Systems and Neural Circuits,
University of Technology Dresden, Germany
{partzsch,mayr,schueffn}@iee.et.tu-dresden.de

Abstract. The Bienenstock-Cooper-Munroe (BCM) rule is one of the best-established learning formalisms for neural tissue. However, as it is based on pulse rates, it can not account for recent spike-based experimental protocols that have led to spike timing dependent plasticity (STDP) rules. At the same time, STDP is being challenged by experiments exhibiting more complex timing rules (e.g. triplets) as well as simultaneous rate- and timing dependent plasticity. We derive a formulation of the BCM rule which is based on the instantaneous postsynaptic membrane potential as well as the transmission profile of the presynaptic spike. While this rule is neither directly rate nor timing based, it can replicate BCM, conventional STDP and spike triplet experimental data, despite incorporating only two state variables. Moreover, these behaviors can be replicated with the same set of only four free parameters, avoiding the overfitting problem of more involved plasticity rules.

1 Introduction

One of the major research areas of neurobiology is long term learning (i.e. plasticity) of synapses in neural tissue [1,2,3,4]. Synapses are the contact points between neurons, where information from the sending neuron arrives at the so-called presynaptic side and is transmitted via the synapse as a postsynaptic current (PSC) pulse to the receiving neuron. The concept of long term plasticity is used to describe the phenomenon that certain types of pre- and/or postsynaptic stimuli can have long lasting effects on the efficacy of this transmission [2,5,6], i.e. the size of the PSC, ranging from days up to a year. These phenomena are called long term depression (LTD) for diminished synaptic responses respectively long term potentiation (LTP) for enhanced responses. Long term plasticity has defied easy modeling, with a host of in vivo and in vitro research showing diverse expressions of this plasticity [2,3]. Depending on the induction protocol, spike rates [6,7], different spike patterns [8,9,10], or membrane voltage control [3,7,10] have been found to elicit changes in plasticity.

* To whom correspondence should be addressed.

** The first two authors contributed equally to the research described in this manuscript.

M. Köppen et al. (Eds.): ICONIP 2008, Part I, LNCS 5506, pp. 137–144, 2009.

Various models have tried to incorporate these findings e.g. in implementations of the classical rate-based BCM rule [5,11,12] or the newer spike-based STDP rule [2,13]. Since both rules describe phenomena which have been well established experimentally as well as theoretically, several models try to achieve a synthesis of both rules. Deriving BCM from STDP formulations is relatively easy [14], but those rules cannot explain nonlinear STDP without introducing secondary state variables [4,9,15]. Trying to reach from BCM to STDP results in STDP with biologically unrealistic characteristics [11,12]. In contrast, in section 2.1 we derive a model of BCM which exhibits major BCM characteristics like LTD at low postsynaptic frequency and LTP at high ones, as well as a frequency threshold (section 3.2). At the same time, this model can be linked analytically to conventional STDP formulations, so the parameters of our model can be computed from those derived experimentally for STDP (section 2.2). Additionally, in section 3.3 we show that this model can replicate experimental data on higher order STDP effects, namely the triplet experiments by Froemke and Dan [9].

2 Methods

2.1 Model

Our plasticity rule is motivated by the original BCM rule formulation [5]. In this model, the weight m of a synapse changes dependent on presynaptic activity $d(t)$ and postsynaptic activity $c(t)$:

$$\frac{\mathrm{d}m}{\mathrm{d}t} = \phi(c(t) - \Theta_M) \cdot d(t) \tag{1}$$

The activities are instantaneous variables denoting changes of spiking frequency around a mean value, so that they can take on positive as well as negative values. $\phi(.)$ is an arbitrary function that changes sign at zero, so that the weight m moves towards d for $c > \Theta_M$, whereas it moves in opposite direction for $c < \Theta_M$.

Since some of the classical BCM studies [7] as well as newer spike-timing studies [3,10] have found a dependence of synaptic plasticity on the postsynaptic membrane potential, we re-interpret the meaning of the variables in the BCM rule. We make two observations: First, the presynaptic activity $d(t)$ at a certain location of a neuron's dendrite (down to a single synapse) can be read out via conductance changes in the postsynaptic membrane that are mediated by presynaptic neurotransmitter release. Often this is modelled by an exponentially decaying synaptic conductance variable (see [16] for an overview). Second, postsynaptic activity $c(t)$ is directly related to postsynaptic membrane potential: Below spiking threshold, the more excitatory input a neuron receives, the higher is its membrane potential. Additionally, the membrane potential involves a similar temporal asymmetry with respect to postsynaptic spikes as standard STDP rules: shortly before a postsynaptic spike, the membrane potential usually is above rest (depolarized), whereas it is below rest (hyperpolarized) shortly after a spike.

Motivated by the first observation, we introduce a presynaptic activity variable, which is a low-pass filtered version of the presynaptic spike train:

$$s(t) = \hat{S} \cdot \mathrm{e}^{-\frac{t-t_j^{\mathrm{pre}}}{\tau_s}} \ , \ t_j^{\mathrm{pre}} \leq t < t_{j+1}^{\mathrm{pre}} \ , \tag{2}$$

where τ_s is the decay time constant and $\hat{S}$ denotes the amplitude of the response to a presynaptic pulse.

Following the second observation, we use a simple spike-response model [16] of the membrane potential as postsynaptic activity, denoted as $u(t)$. This model consists of a Dirac pulse and an exponential decay to account for hyperpolarization after a spike:

$$u(t) = U_{\mathrm{p}} \cdot \delta(t - t_n^{\mathrm{post}}) + U_{\mathrm{refr}} \cdot \mathrm{e}^{-\frac{t-t_n^{\mathrm{post}}}{\tau_{\mathrm{refr}}}} \ , \ t_n^{\mathrm{post}} - 0 < t < t_{n+1}^{\mathrm{post}} - 0 \tag{3}$$

In this equation, U_{p} represents the area under the pulse curve, determining the pulse amplitude. $U_{\mathrm{refr}} < 0$ is the amplitude of post-spike hyperpolarization and τ_{refr} is the membrane time constant. Sub-threshold variations of membrane potential before a spike are neglected for simplicity.

For most experimental protocols, we included an attenuation of the postsynaptic spike. This attenuation was implemented by weighting the action potential amplitude U_{p} with the negative membrane potential:

$$U_{\mathrm{p}} \ \rightarrow \ U_{\mathrm{p}}(t^{\mathrm{post}}) = U_{\mathrm{p}} \cdot (1 - \alpha_{\mathrm{att}} \frac{u(t^{\mathrm{post}} - 0)}{U_{\mathrm{refr}}}) \tag{4}$$

With this formulation, a postsynaptic spike occuring shortly after a previous one will have amplitude $U_{\mathrm{p}} \cdot (1 - \alpha_{\mathrm{att}})$. Note that, due to setting the membrane voltage to U_{refr} after each pulse, the pulse amplitude will only depend on the time course of the membrane voltage since the last action potential.

The resulting voltage-based rule reads as follows:

$$\frac{\mathrm{d}m}{\mathrm{d}t} = (u(t) - \Theta_u) \cdot s(t) \ , \tag{5}$$

where Θ_u represents the threshold between LTP and LTD. This equation is similar to that of an ion channel with reversal potential Θ_u, activated by presynaptic pulses. Interestingly, similar rules have been synthesized for solving image processing tasks in VLSI realizations of neural networks [17].

Following our motivation, s should be a conductance, i.e. $[s]$ =A/V=S, and u should be a potential, i.e. $[u]$ =V, so that $[U_{\mathrm{p}}]$ =Vs. Thus, the weight m has units of a charge: $[m]$ =As. This may be interpreted as the amount of neurotransmitters released at the synapse. LTP would thus raise the available neurotransmitter amount, whereas LTD would lower it.

2.2 Explicit Expression and Parameter Fitting

Equation (5) can be integrated to arrive at an explicit expression for the weight at time t. For determination of the model's parameters, we will derive such

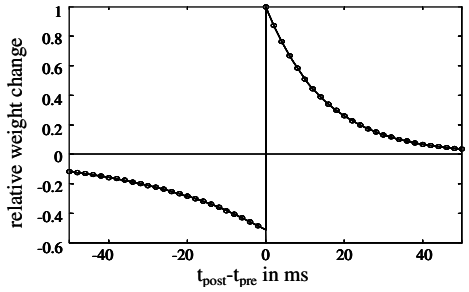

Fig. 1. Normalized STDP window of our rule for the parameters of Froemke and Dan [9] (see Fig. 2), derived analytically (solid line) and via simulations (circles, 60 pairings at 1Hz, protocol of [9,8])

expressions for a standard pairing protocol (see e.g. [8]). For clarity, we set $\Theta_u = 0$ and calculate the weight change triggered by a single pair of spikes with timing difference $\Delta t_n = t_n^{\text{post}} - t_n^{\text{pre}}$. Since the pairings have a low repetition frequency, the attenuation of U_p can be neglected. For pre-post pairings ($\Delta t_n > 0$) the following equation for times $t > t_n^{\text{post}}$ holds:

$$\Delta m_n(t) = \hat{S}\left(U_\text{p} + \frac{U_\text{refr}}{\frac{1}{\tau_s}+\frac{1}{\tau_\text{refr}}}\right)\mathrm{e}^{-\frac{|\Delta t_n|}{\tau_s}} - \frac{\hat{S}U_\text{refr}}{\frac{1}{\tau_s}+\frac{1}{\tau_\text{refr}}}\cdot\mathrm{e}^{-\frac{t-t_n^\text{pre}}{\tau_s}}\cdot\mathrm{e}^{-\frac{t-t_n^\text{post}}{\tau_\text{refr}}} \tag{6}$$

Post-pre pairings ($\Delta t_n < 0$) lead to a similar relationship for times $t > t_n^{\text{pre}}$:

$$\Delta m_n(t) = \hat{S}\cdot\frac{U_\text{refr}}{\frac{1}{\tau_s}+\frac{1}{\tau_\text{refr}}}\mathrm{e}^{-\frac{|\Delta t_n|}{\tau_\text{refr}}} - \frac{\hat{S}U_\text{refr}}{\frac{1}{\tau_s}+\frac{1}{\tau_\text{refr}}}\cdot\mathrm{e}^{-\frac{t-t_n^\text{pre}}{\tau_s}}\cdot\mathrm{e}^{-\frac{t-t_n^\text{post}}{\tau_\text{refr}}} \tag{7}$$

For small pairing frequency, i.e. $t \to \infty$, the terms dependent on t diminish in equations (6) and (7), so that the final weight change Δm_n merely depends on the timing difference Δt. Due to the exponential relationship, our model can be directly fitted to the exponential time window of standard STDP rules, parameterized by the amplitudes $A_{+/-}$ and time constants $\tau_{+/-}$ for LTP and LTD, respectively:

$$U_\text{refr} = A_- \cdot \left(\frac{1}{\tau_s}+\frac{1}{\tau_\text{refr}}\right)\frac{1}{\hat{S}} \qquad \tau_\text{refr} = \tau_- \tag{8}$$

$$U_\text{p} = (A_+ - A_-)\,\frac{1}{\hat{S}} \qquad \tau_s = \tau_+ \tag{9}$$

We have set $\hat{S} = 1\text{nS}$ and added units fAs to the amplitudes $A_{+/-}$ to arrive at membrane potential variations in the order of biological values.

Our analytical findings are confirmed by simulations of our BCM model with a low-frequency STDP pairing protocol, see Figure 1. As expected, the agreement with the exponential time window as derived above is excellent.

3 Results

3.1 Model Behaviour

Our model employs different mechanisms for LTP and LTD: While LTP results from sampling of presynaptic activity at postsynaptic spikes, LTD is continuously activated by coincidence of presynaptic activity and membrane hyperpolarization directly after a postsynaptic spike, see Figure 2A. Amplitudes of postsynaptic spikes are attenuated if they occur shortly after each other, so that for high postsynaptic firing rates the amount of LTP per spike is reduced. This counteracts the reduced amount of LTD per postsynaptic spike that results from shortened hyperpolarization periods in this regime.

It is important to note that our model parameters, being strongly linked to biophysical parameters, take on biologically realistic values when they are derived from standard STDP parameters (see section 2.2). The PSC time constant τ_s equals the STDP time constant for LTP, which is in the order of 20ms, a value compatible with NMDA synapse conductance changes [16,13]. The STDP time constant for LTD equals the membrane time constant τ_{refr}, which was found to be in the order of 10-40ms [1]. If we scale our membrane model, so that $U_{\mathrm{refr}} = -5\mathrm{mV}$, we arrive at a value $U_{\mathrm{p}} = 151\mu\mathrm{Vs}$ for realistic STDP parameters (see Fig. 2B, [9]), corresponding to a (rectangular) pulse of length 2ms and height 75mV, which are reasonable values for action potentials [1].

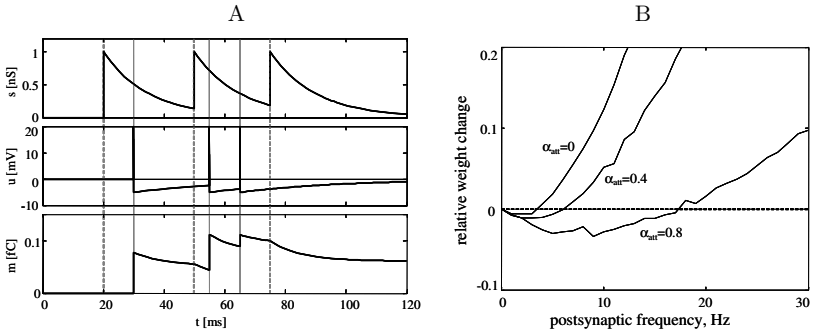

Fig. 2. A: progression of presynaptic activity s, postsynaptic activity u and weight m for a sample spike train. B: Weight change for different attenuation factors α_{att} when stimulating with Poisson spike trains at varying postsynaptic rates. Parameters for equation (5) as derived from Froemke and Dan [9] ($\tau_+ = 14.8\mathrm{ms}$, $\tau_- = 33.8\mathrm{ms}$, $A_+ = 1.01$, $A_- = -0.52$) via equations (8) and (9).

3.2 Frequency Dependence

In a first experiment, we applied our rule to a protocol recently used by Izhikevich and Desai [14] to test STDP models on their BCM characteristics. Thereby,

random Poisson spike trains of length 200s where generated, with the presynaptic firing rate held constant at 10Hz, but varying postsynaptic firing rate.

Figure 2B shows resulting frequency curves for different degrees of action potential attenuation. They all show typical BCM behaviour, depending on the attenuation factor α_{att}: The amount of LTD as well as the threshold between LTD and potentation could be increased by raising α_{att}. This is because the postsynaptic action potential is the only source of LTP, so that its attenuation will directly result in a higher amount of LTD.

On the other hand, a positive threshold Θ_u would add a constant amount of LTD to the curves in Figure 2B, which shifts the entire curve down instead of changing its slope as for α_{att}. This is consistent with experimental findings on the plasticity-depolarization relationship [10]. Also, due to this shifting property, the frequency threshold Θ_M between LTD and LTP of the original BCM rule is directly related to the voltage threshold Θ_u in our model.

3.3 Triplet Protocols

We tested the performance of our model on triplet protocols using the data of Froemke and Dan [9]. Thereby, either two presynaptic and one postsynaptic or one presynaptic and two postsynaptic pulses were generated, and weight changes monitored dependent on the two independent time differences in the protocol. Such experiments are only poorly replicated by standard STDP models [4]. We did not fit our model especially to the triplet data, but used the STDP parameters measured in their control experiment, as before (see Figure 2B).

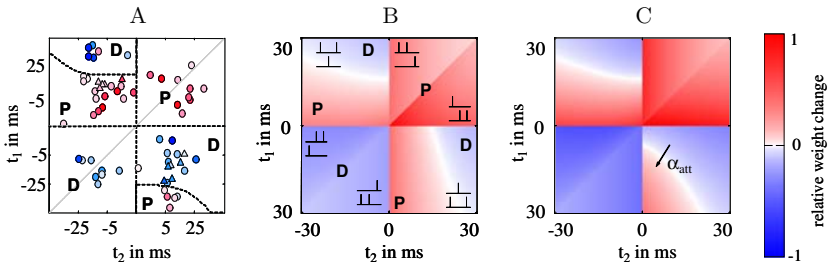

Fig. 3. Triplet experiments. Upper left part: 2 pre-, 1 post-synaptic spike; lower right part: 1 pre-, 2 post-synaptic spikes; A: experimental data of Froemke and Dan [9]; B: our basic model, using parameters from Froemke and Dan [9] (see Fig. 2); C: same as B, but with additional attenuation of postsynaptic action potentials, $\alpha_{\text{att}} = 0.8$.

Figure 3 shows relative weight changes after application of the spike triplets. The time differences are measured post-minus-pre, as in [9]. Our basic model matches the experimental results for most triplets, comparable to the biophysical model in [13]. Particularly, it reproduces the responses to pre-post-pre pairings (upper left quarter), but fails to account for the post-pre-post responses.

We can make our model consistent with the post-pre-post triplets if we add attenuation of postsynaptic spikes, see Figure 3C: The amount of LTD for the post-pre-post triplets is adjusted by the attenuation α_{att} (see arrow), so that our model can replicate the experimental data of [9] very well.

4 Discussion

In this paper, we have introduced a simple learning rule that is motivated by the original BCM formulation and membrane potential variations of biological neurons. Using a simple spike response neuron, we can directly map our model to standard pair-based STDP rules. Like iterative implementations of STDP, our model samples presynaptic activity at postsynaptic spikes for LTP. In contrast to that, postsynaptic hyperpolarization after action potentials is continuously integrated with presynaptic activity to arrive at a mechanism for LTD.

When replicating experimental data with such models, major challenges have been the integration of triplets [9] and higher order pairings [10,4] as well as frequency effects [14,15,6]. Several higher-order models with secondary state variables [15,9,4] have been proposed to reproduce some of these effects. However, incorporating no secondary state variables, our model can account for BCM effects, conventional STDP and triplet experiments with only four free parameters. The generalization capacity of our model can be seen from the fact that all these effects can be generated with one biologically realistic parameter set, while even higher order models have to fit several parameter sets to account for different experimental data [15,13]. In addition, our model is able to link the parameters of an STDP model to the ones of the neuron model and to BCM parameters.

Our model is also suited for implementation in neuromorphic hardware. Implementing a membrane-based Hebb rule similar to ours requires only low circuit complexity [17]. Compared to more involved realizations of neurobiological behavior, which replicate both STDP and complex membrane dynamics [18], our model is easier to implement. This is because in those models, the dynamic variables required by our model are already contained as membrane potential and synaptic conductance, so that the weight change may be directly derived from these quantities, requiring no separate circuitry for STDP state variables.

Because in our model, plasticity is directly linked to depolarization, experimental protocols employing this variable [10,3] could be qualitatively replicated, an advantage that purely spike-based models are lacking. Further analysis will also try to create a stronger link between the frequency threshold of conventional BCM and our voltage threshold. The metaplasticity of this threshold [5,11] might be accounted for by variations in resting potential [1,10] caused by metabolic changes in the dendritic membrane.

Acknowledgements. This work is supported by funding under the Sixth Framework Programme of the European Union under the grant no. 15879 (FACETS). We would thank W. Gerstner for fruitful discussions on plasticity.

References

1. Koch, C.: Biophysics of computation. Information processing in single neurons. In: Computational Neuroscience. Oxford University Press, Oxford (1999)
2. Morrison, A., Diesmann, M., Gerstner, W.: Phenomenological models of synaptic plasticity based on spike timing. Biological Cybernetics 98, 459–478 (2008)
3. Lisman, J., Spruston, N.: Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. Nature Neuroscience 8(7), 839–841 (2005)
4. Pfister, J.P., Gerstner, W.: Triplets of spikes in a model of spike timing-dependent plasticity. Journal of Neuroscience 26(38), 9673–9682 (2006)
5. Bienenstock, E., Cooper, L., Munro, P.: Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. Journal of Neuroscience 2(1), 32–48 (1982)
6. Dudek, S., Bear, M.: Homosynaptic long-term depression in area CAl of hippocampus and effects of N-methyl-D-aspartate receptor blockade. PNAS 89, 4363–4367 (1992)
7. Artola, A., Bröcher, S., Singer, W.: Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. Nature 347, 69–72 (1990)
8. Bi, G.Q., Poo, M.M.: Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. Journal of Neuroscience 18(24), 10464–10472 (1998)
9. Froemke, R., Dan, Y.: Spike-timing-dependent synaptic modification induced by natural spike trains. Nature 416, 433–438 (2002)
10. Sjöström, P., Turrigiano, G., Nelson, S.: Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 32, 1149–1164 (2001)
11. Shouval, H., Bear, M., Cooper, L.: A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. PNAS 99(16), 10831–10836 (2002)
12. Kurashige, H., Sakai, Y.: BCM-type synaptic plasticity model using a linear summation of calcium elevations as a sliding threshold. In: King, et al. (eds.) ICONIP 2006. LNCS, vol. 4232, pp. 19–29. Springer, Heidelberg (2006)
13. Badoual, M., Zou, Q., Davison, A., Rudolph, M., Bal, T., Fregnac, Y., Destexhe, A.: Biophysical and phenomenological models of multiple spike interactions in spike-timing dependent plasticity. International Journal of Neural Systems 16(2), 79–97 (2006)
14. Izhikevich, E., Desai, N.: Relating STDP to BCM. Neural Computation 15, 1511–1523 (2003)
15. Lu, B., Yamada, W., Berger, T.: Asymmetric synaptic plasticity based on arbitrary pre- and postsynaptic timing spikes using finite state model. In: Proceedings of International Joint Conference on Neural Networks (2007)
16. Gerstner, W., Kistler, W.: Spiking neuron models: single neurons, populations, plasticity. Cambridge University Press, Cambridge (2002)
17. Schreiter, J., Ramacher, U., Heittmann, A., Matolin, D., Schüffny, R.: Cellular pulse coupled neural network with adaptive weights for image segmentation and its VLSI implementation. In: Proceedings 16th International Symposium on Electronic Imaging: Science and Technology, vol. 5298, pp. 290–296 (2004)
18. Schemmel, J., Brüderle, D., Meier, K., Ostendorf, B.: Modeling synaptic plasticity within networks of highly accelerated I&F neurons. In: ISCAS 2007 (2007)

2.3 Appendix 3

C. Clopath, L. Ziegler, E. Vasilaki, L. Bsing and W. Gerstner, Tag-Trigger-Consolidation: A Model of Early and Late Long-Term-Potentiation and Depression, PLoS Comput Biol, Vol. 4, Nr. 12, 2008.

OPEN ACCESS Freely available online

Tag-Trigger-Consolidation: A Model of Early and Late Long-Term-Potentiation and Depression

Claudia Clopath[◔], **Lorric Ziegler**[◔], **Eleni Vasilaki, Lars Büsing**[¤], **Wulfram Gerstner***

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

Abstract

Changes in synaptic efficacies need to be long-lasting in order to serve as a substrate for memory. Experimentally, synaptic plasticity exhibits phases covering the induction of long-term potentiation and depression (LTP/LTD) during the early phase of synaptic plasticity, the setting of synaptic tags, a trigger process for protein synthesis, and a slow transition leading to synaptic consolidation during the late phase of synaptic plasticity. We present a mathematical model that describes these different phases of synaptic plasticity. The model explains a large body of experimental data on synaptic tagging and capture, cross-tagging, and the late phases of LTP and LTD. Moreover, the model accounts for the dependence of LTP and LTD induction on voltage and presynaptic stimulation frequency. The stabilization of potentiated synapses during the transition from early to late LTP occurs by protein synthesis dynamics that are shared by groups of synapses. The functional consequence of this shared process is that previously stabilized patterns of strong or weak synapses onto the same postsynaptic neuron are well protected against later changes induced by LTP/LTD protocols at individual synapses.

Citation: Clopath C, Ziegler L, Vasilaki E, Büsing L, Gerstner W (2008) Tag-Trigger-Consolidation: A Model of Early and Late Long-Term-Potentiation and Depression. PLoS Comput Biol 4(12): e1000248. doi:10.1371/journal.pcbi.1000248

Editor: Lyle J. Graham, UFR Biomédicale de l'Université René Descartes, France

Received August 18, 2008; **Accepted** November 10, 2008; **Published** December 26, 2008

Funding: This work was partially supported by the European community via th FACETS project. CC was supported by the Swiss National Science Foundation. The sponsors did not influence the design or analysis of the study.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: wulfram.gerstner@epfl.ch

¤ Current address: Institut für Grundlagen der Informationsverarbeitung, TU Graz, Graz, Austria

◔ These authors contributed equally to this work.

Introduction

Changes in the connection strength between neurons in response to appropriate stimulation are thought to be the physiological basis for learning and memory formation [1,2]. A minimal requirement for proper memory function is that these changes, once they are induced, persist for a long time. For several decades, experimentalists have therefore focused on Long-Term Potentiation (LTP) and Long-Term Depression (LTD) of synapses in hippocampus [3,4] and cortical areas [5,6]. LTP can be induced at groups of synapses by strong 'tetanic' high-frequency stimulation of the presynaptic pathway [3] while stimulation at lower frequency leads to LTD Dudek92. Both LTP and LTD can also be induced at a single synapse or a small number of synaptic contacts if presynaptic activity is paired with either a depolarization of the postsynaptic membrane [5,7] or tightly timed postsynaptic spikes [8,9].

While the induction protocol for LTP and LTD is often as short as a few seconds, the changes in synaptic efficacy persist for much longer [9]. In typical slice experiments on LTP [and similarly for LTD or Spike-Timing Dependent Plasticity (STDP)] the persistence of the change is monitored for 30 minutes to 1 hour. Accumulating evidence suggests, however, that after this early phase of LTP (E-LTP) different biochemical processes set in that are necessary for the further maintenance of potentiated synapses during the late phase of LTP (L-LTP) [10,11]. For an understanding of the transition from early to late LTP, the concept of 'synaptic tagging and capture' has become influential [12,13]. During induction of the early phase of LTP, each potentiated synapse sets a tag that marks that it has received a specific afferent signal. A candidate molecule, involved in the tag signaling LTP induction in apical dendrites of hippocampal neurons, is the calcium-calmodulin dependent kinase II (CaMKII) [13]. Newly synthesized plasticity-related proteins are 'captured' by the tagged synapse and transform E-LTP into L-LTP that can be maintained over hours or days. A candidate protein involved in the maintenance of potentiated hippocampal synapses is the protein kinase Mζ (PKMζ) [11,14].

The stabilization and maintenance of potentiated synapses poses a number of theoretical challenges. First, on the level of single synapses we must require synaptic strength to remain stable, despite the fact that AMPA channels in the postsynaptic membrane are continuously exchanged and recycled [15–17]. Thus the synapse is not 'frozen' but part of a dynamic loop. Second, on the level of neuronal representation in cortical areas, one finds representations of input features that are stable but at the same time sufficiently plastic to adjust to new situations [18]. In the theoretical community, this paradox has been termed the stability-plasticity dilemma in unsupervised learning [19]. Third, humans keep the ability to memorize events during adulthood, but can also remember earlier episodes years back. However, continued learning of new patterns in theoretical models of associative memory networks forces the erasure or 'overwriting' of old ones, the so-called palimpsest property [20,21]. In the context of

Author Summary

Humans and animals learn by changing the strength of connections between neurons, a phenomenon called synaptic plasticity. These changes can be induced by rather short stimuli (lasting sometimes only a few seconds) but should then be stable for months or years in order to be useful for long-term memory. Experimentalists have shown that synapses undergo a sequence of steps that transforms the rapid change during the early phase of synaptic plasticity into a stable memory trace in the late phase. In this paper we introduce a model with a small number of equations that can describe the phenomena of induction of synaptic changes during the early phase of synaptic plasticity, the trigger process for protein synthesis, and the final stabilization. The model covers a broad range of experimental phenomena known as tagging experiments and makes testable predictions. The ability to model the stabilization of synapses is crucial to understand learning and memory processes in animals and humans and a necessary ingredient for any large-scale model of the brain.

continued learning, theoretical arguments show that synaptic plasticity on multiple time scales cannot prevent, but at most delay the erasure of memories in the presence of ongoing synaptic activity [22]. This suggests that additional mechanisms are necessary to further protect existing memories and 'gate' the learning of new ones.

Despite these challenges for the long-term stability of synapses, most classical models of synaptic plasticity focus on the induction and early phase of LTP or LTD and completely ignore the question of maintenance. Traditional models of associative memories separate the learning phase from the retrieval phase [23] and the same holds for standard models of STDP [24–26]. Detailed biophysical models of LTP and LTD describe calcium dynamics and Calcium/Calmodulin-Dependent Protein Kinase II (CaMKII) phosphorylation during the induction and early phase of LTP [27–29]. While these models show that switches built of CaMKII proteins can be stable for years, they do not address aspects of tagging leading to heterosynaptic interaction during L-LTP and L-LTD. Moreover, while CaMKII phosphorylation is necessary for induction of LTP and mediate tags in the apical dendrites of hippocampal CA1 neurons [30], it is less clear whether it is necessary for its maintenance [31]. On the other hand protein kinase Mζ is essential for maintenance of some synapse types [11,13,14] but the same molecule is potentially relevant for *induction* in others [30].

We wondered whether a simple model that connects the process of LTP induction with that of maintenance would account for experimental results on tagging and 'cross-tagging' [11–13,32] without specific assumptions about the (partially unknown) molecular pathways involved in the maintenance process. If so, the model should allow us to discuss functional consequences that are generic to the tagging hypothesis independent of the details of a biophysical implementation in the cell. Even though we believe that the model principles are more general, we focus on synapses from the Schaffer-Collaterals onto the CA1 neurons in hippocampus as an experimentally well-studied reference system for synaptic plasticity. Since typical tagging experiments involve the extracellular stimulation of one or several *groups* of synapses (rather than single synapses), our model of early and late LTP/LTD is developed in the context of a neuron model with hundreds of synapses. The application of the principles of synaptic consolidation to experiments inducing E-LTP/E-LTD at *single* synapses is considered in the discussion section.

Results

We study a model with a large number of synapses i onto a single postsynaptic neuron. To be specific, we think of a pyramidal neuron in the CA1 area of hippocampus. Our model combines features of traditional models for the *induction* of potentiation [24–26,33–36] with a simple description of tagging and synthesis of plasticity related proteins that finally lead to the *maintenance* of the induced changes. The section is organized as follows: We first introduce the essential components of the model step by step ('Constructing the Model'). We then test the performance of the model with a set of stimuli typically used to induce long-term changes of synapses ('Testing the Model').

Constructing the Model

Our model contains three elements, Figure 1. The first one sets the *tag* during the induction of E-LTP or E-LTD. A tag is indicated by a value $h=1$ for LTP or $l=1$ for LTD. In the absence of tags we have $h=l=0$. The second one describes the process that *triggers* the synthesis of plasticity related proteins. The final component describes the up-regulation of a maintenance-related process from a low value ($z=0$) to a high value ($z\approx 1$). The dynamics of this component is intrinsically bistable and leads to a *consolidation* of the previously induced change at the labeled synapses upon interaction with the protein p ('protein capture'). The total change Δw of the synaptic strength reported in experiments contains contributions [13] of the early components l and h as well as the late component z. Since the model describes a sequence of three steps 'Tag-Trigger-Consolidation' we call it in the following the TagTriC-Model (Figure 1).

Tag and Induction of LTP/LTD

Results from minimal stimulation protocols which putatively activate only a single synapse suggest that the induction of LTP is a switch-like process [7,37]. We therefore model individual synapses as discrete quantities that can switch, during the induction of LTP, from an initial 'non-tagged state' (N) to a 'high state' (H) with a transition rate ρ_H that depends on the induction protocol. Similarly, induction of LTD moves the synapse from the initial non-tagged state (N) to a 'low state' (L) at a rate ρ_L. If synapse i is in the high state, the synaptic variable h_i is equal to one. If it is in the low state, another local variable l_i is set to one. These local variables h_i and l_i do not only control the weight of the synapse during E-LTP and E-LTD, but also serve as 'tags' for up- or down-regulation of the synapse. Tags reset to zero stochastically with a rate k_h and k_l, respectively. If both tags are zero, the synapse is in the non-tagged state N. Since the synapse is either up-regulated OR down-regulated, at most one of the tags can be non-zero (Figure 1A).

The stochastic transitions from the initial state N with $h_i=0$ and $l_i=0$ to the down-regulated state $l_i=1$ or an upregulated state $h_i=1$ depend in a Hebbian manner on presynaptic activity and the state of the postsynaptic neuron. In the absence of presynaptic activity, the LTD rate ρ_L vanishes. Presynaptic activity combined with a time-averaged membrane potential $\bar{u}$ above a critical value $\vartheta_{\rm LTD}$ leads in the TagTriC model to a LTD transition rate ρ_L proportional to $[\bar{u}(t)-\vartheta_{\rm LTD}]$. For a transition from the initial state to the high state, we require in addition that the *momentary* membrane potential is above a second threshold $\vartheta_{\rm LTP}$. Hence the transition rate ρ_H is proportional to $[\bar{u}(t)-\vartheta_{\rm LTD}][u-\vartheta_{\rm LTP}]$

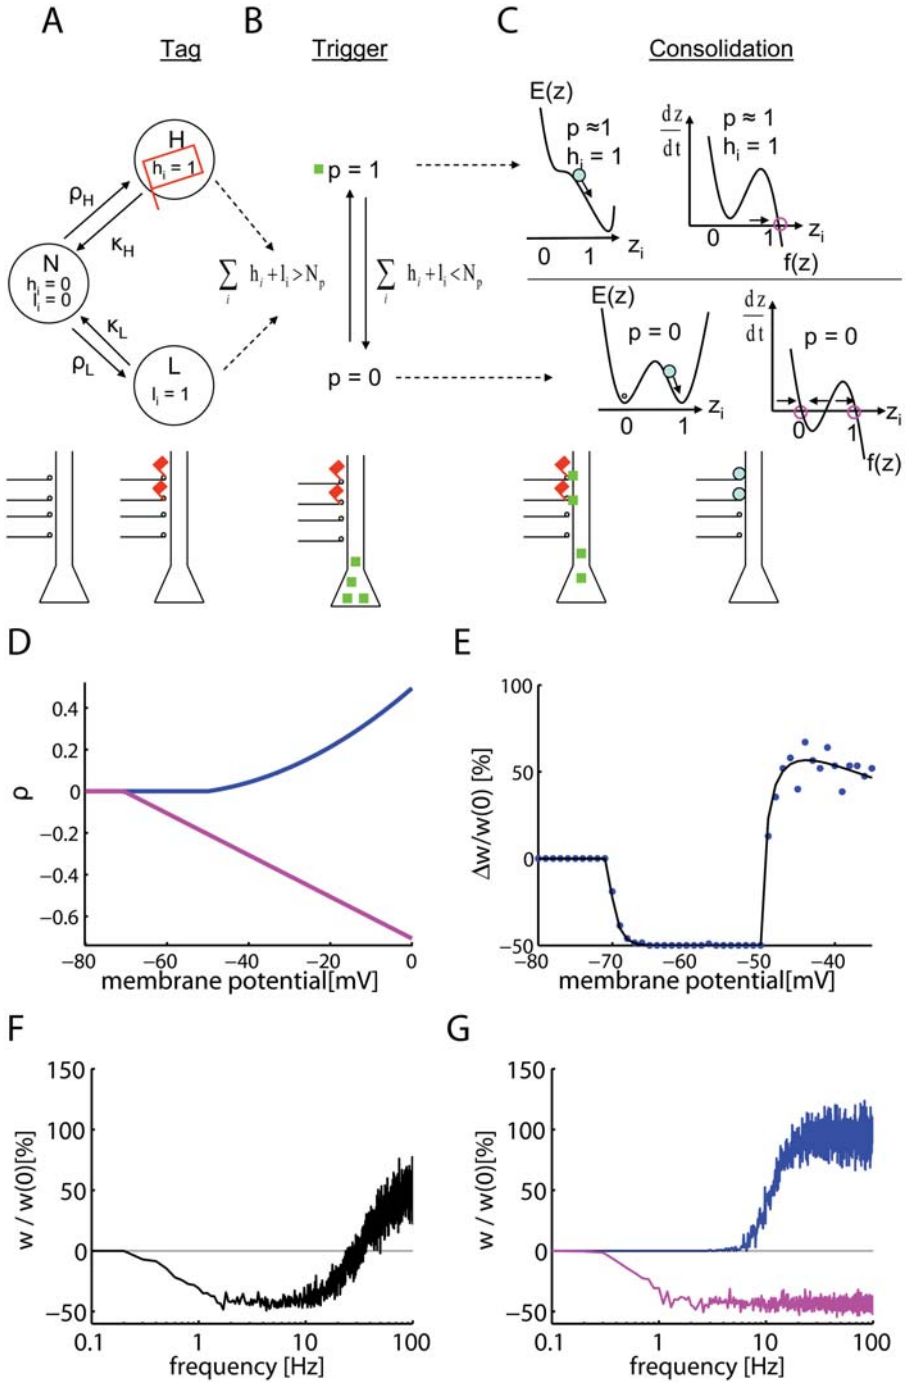

Figure 1. The three components of the Tag-Trigger-Consolidation (TagTriC) model. (A) A synapse can be in the non-tagged state N, the high state H or the low state L. A synapse i in H (or L) has a tag $h_i = 1$ (or $l_i = 1$, respectively). Transitions to a tagged state occur with rates ρ_H for potentiation and ρ_L for depression. The tag $h_i = 1$ is indicated by a red flag in both the flow graph and the schematic drawing below. (B) Synthesis of plasticity related proteins p (green squares) is triggered if the total number of set tags is larger than a critical number N_p. If the trigger threshold N_p is not reached, the protein concentration decays back to zero. (C) The consolidation dynamics can be visualized as downward motion in a potential surface $E(z)$. The function $f(z)$ (shown to the right) is the derivative of E and characterizes the dynamics $dz/dt = f(z)$. If a tag is set at the synapse ($h_i = 1$) and protein synthesis has been triggered ($p \approx 1$), the dynamics can be imagined as downward motion into the right well of the potential $E(z)$. In this case, $z \approx 1$ is the only fixed point of the dynamics (magenta circle). In the absence of tags ($h_i = l_i = 0$, below) the consolidation variable z_i of synapse i is bistable and approaches (direction of flow indicated by arrows) stable fixed points at $z_i = 0$ or $z_i = 1$ (magenta circles). The steps of synaptic tagging and capture are indicated immediately below the flow diagram. (D) The tagging rates for depression ($-\rho_L$,(magenta)) and for potentiation ρ_H (blue) are shown as a function of the clamped voltage under the assumption that a presynaptic spike has arrived less than 1 millisecond before. Note that for depression we plot the negative rate $-\rho_L$ rather than ρ_L to emphasize the fact that depression leads to a down-scaling of the synapse. (E) Voltage dependence of early LTP and LTD. The weight change $\Delta w/w(0)$ induced by a stimulation of 100 synapses at 2 Hz during 50 s while the postsynaptic voltage is clamped is shown as a function of voltage. The percent change $\Delta w/\hat{w}$ in simulations (circles) of LTP/LTD induction experiments can be predicted from a theory (solid line) based on the difference in transition rates $\rho_H - \rho_L$. The simulation reflects the voltage dependence seen in experiments [5,39]. (F,G) Frequency dependence of early LTP and LTD. Simultaneous stimulation of 100 synapses by 3 trains (separated by 5 min) of 100 pulses at rates ranging 0.03 to 100 Hz shows LTD at low frequencies and LTP at frequencies above 30 Hz. (G) If LTP is blocked in the model, LTD (pink line) occurs up to high frequencies as in experiments [7]. Blue line: LTP with blocked of LTD.
doi:10.1371/journal.pcbi.1000248.g001

whenever these threshold conditions are satisfied; see Methods for details.

Our assumptions regarding the transition rates essentially summarize the qualitative voltage dependence seen in the Artola-Bröcher-Singer experiments [5]. Indeed, when 100 synapses in the TagTriC model are stimulated at low frequency during 50 seconds while the membrane voltage is kept fixed at different values (Figure 1D), the total weight change summed across all synapses exhibits LTD at low voltage and LTP at high voltage [38,39]. As expected, the resulting weight changes in the simulations of Figure 1E reflect the voltage dependence of the transition rates in Figure 1D.

Trigger for Protein Synthesis

Previously induced LTP or LTD needs to be consolidated in order to last for more than one hour. Consolidation requires that protein synthesis is triggered. Experimental evidence indicates that triggering of protein synthesis needs the presence of neuromodulators such as dopamine (in the apical CA1 region) or other modulators (in other regions). In typical tagging experiments, extracellular stimulation co-stimulates dopaminergic input leading to a phasic dopamine signal [13,40]. In our model, induction of E-LTP or E-LTD through appropriate stimulation protocols changes the synaptic efficacy and sets tags at the modified synapses, both described by the variables $h_i = 1$ or $l_i = 1$. Protein synthesis in the model is triggered (see methods for details) if the total number of tags $\Sigma_i(h_i+l_i)$ (which indirectly reflects the phasic dopamine signal) reaches a threshold N_p which depends on the level of background dopamine (and other neuromodulators). More specifically, N_p decreases with the concentration of background dopamine so that the presence of dopamine facilitates the trigger process [32].

If the trigger criterion is satisfied, the concentration p of synthesized plasticity related proteins approaches with rate k_p a value close to one. If the number of tags falls below the threshold N_p, the protein concentration p decays with a time constant τ_p back to zero. Further details on the role of the trigger threshold and its relation to neuromodulators can be found in the discussion section.

Consolidation and Late LTP

The total weight w_i of a synapse i depends on the present value of the tags h_i or l_i as well as on its long-term value z_i. The slow variable z_i is a continuous variable with one or two stable states described by a generic model of bistable switches, that could be implemented by suitable auto-catalytic processes [16]. While the concentration p of plasticity related proteins is zero, the variable z_i has two stable states at $z_i = 0$ and $z_i = 1$, respectively. If the protein concentration takes a value of $p \approx 1$, one of the stable states disappears and, depending on the tag that was set, the long term-value of the synapse can be up- or down-regulated; see methods and Figure 1C for details.

In order to illustrate the mechanism of induction of L-LTP, let us suppose that the synapse has been initially close to the state $z_i = 0$. The dynamics of the synapse can be imagined as downward motion in a 'potential' E. The current stable state of the synapse is at the bottom of the left well in the potential pictured in Figure 1C. We assume that during a subsequent LTP induction protocol the synapse has been tagged with $h_i = 1$ and that the total number of tags set during the LTP induction protocol surpasses the trigger threshold N_p. If the protein concentration p approaches one, the potential surface is tilted so that the synapse now moves towards the remaining minimum at $z \approx 1$. After decay of the tags, p returns to zero, and we are back to the original potential, but now with the synapse trapped in the state $z = 1$. It can be maintained in this state for a long time, until another strong tagging event occurs during which the synapse is tagged with $l_i = 1$ as a result of LTD induction. In this case the potential surface can be tilted towards the left so that the only equilibrium point is at $z = 0$. Since consolidation is typically studied in animals that are more than 20 days old [13], we assume that before the beginning of the experiment 30 percent of the synapses are already in the upregulated state $z = 1$ and the remaining 70 percent in the state $z = 0$; see also [7]. Because of the bistable dynamics of consolidation, only synapses that are initially in the upregulated state $z = 1$ can undergo L-LTD and only synapses that start from $z = 0$ can undergo L-LTP; compare [7]. Note, however, that tags for potentiation and depression can be set independently of the value of z. We may speculate that the variable z is related to the activity of PKMζ [11,14], or to the self-sustained clustering of AMPA receptors [41], but the exact biochemical signaling chain is irrelevant for the functional consequences of the model discussed in the results section. In our model, the bistable dynamics of the z-variable captures the essence of synaptic persistence despite molecular turnover [15,16,28] and mobility of AMPA receptors [41].

Tests of the Model

The TagTriC model has been tested on a series of stimulation protocols that reflect induction of LTP and LTD as well as the consolidation of plasticity events.

Induction of Synaptic Changes

A typical LTP induction experiment starts with extracellular stimulation of a bundle of presynaptic fibers (i.e., the Schaffer collaterals leading from CA3 to CA1) that activate a large number (typically hundreds [13]) of presynaptic terminals. With an extracellular probe electrode placed close to one of the postsynaptic neurons, a change in synaptic efficacy is measured via the amplitude (or initial slope) of the evoked postsynaptic potential, representing the total response summed across all the stimulated synapses. In our simulations, we mimic these experiments by simultaneous stimulation of 100 synapses. The state of the postsynaptic neuron is described by the adaptive exponential integrate-and-fire model [42] and can be manipulated by current injection.

In a preliminary set of simulation experiments done with presynaptic stimulation alone (no manipulation of the postsynaptic neuron), the TagTriC model exhibits LTD or LTP depending on the frequency of the presynaptic stimulation (Figure 1F) in agreement with experimental results [4,43]. Moreover, under the assumption that LTP has been blocked pharmacologically ($\rho_H = 0$ in the model), our model shows LTD even for high stimulation frequencies (Figure 1G). This stems from the fact that LTD and LTP are represented in the TagTriC model by two independent pathways (Figure 1A) which are under control condition in competition with each other, but show up individually if one of the paths is blocked [43]. Together with the voltage dependence of Figure 1E, the above simulation results indicate that our model of LTP and LTD induction can account for a range of experiments on excitatory synapses in the hippocampal CA1 region, in particular, voltage and frequency dependence.

Consolidation of Synaptic Changes

In order to study whether consolidation of synaptic changes in our model follows the time course seen in experiments, we simulate standard experimental stimulation protocols [12,13]. A weak tetanus consisting of a stimulation of 100 synapses at 100 Hz for 0.2 seconds (21 pulses) leads in our model to the induction of LTP (change by +15 percent) which decays back to baseline over

the time course of two hours (Figure 2A). Thus, after the early phase of LTP the synapses are not consolidated. A stronger stimulus consisting of stimulating the same group of hundred synapses by 100 pulses at 100 Hz (repeated 3 times every 10 minutes) yields stronger LTP that consolidates and remains elevated (weight change by 22±5 percent) for as long as the simulations are continued (more than 10 hours, only the first 5 hours are shown in Figure 2B). Thus our model exhibits a transition from early to late LTP if E-LTP is induced by the strong tetanic stimulation protocol, but not the weak one, consistent with results in experiments [12,13]. If, however, the weak tetanus at a first group of 100 synapses is given 30 minutes before or after a strong tetanus at a second group of 100 synapses, the synapses in both the weakly and strongly stimulated groups are consolidated (Figure 2C and 2D). If the weak tetanus in group one is given 120 minutes after the strong tetanus in group two, then consolidation of the synapses in the weakly stimulated group does not occur (Figure 2E). Thus our model exhibits a time course of heterosynaptic interaction between the two groups of synapses as reported in classical tagging experiments [12,13].

An advantage of a modeling approach is that we can study the dependence of the heterosynaptic interaction between the two groups of synapses upon model parameters. A critical parameter in the model is the trigger threshold N_p that needs to be reached in order to start protein synthesis (Figure 1B). With our standard choice of parameters, where $N_p = 40$, we can plot the consolidated weight change $\Delta w/w(0)$ in the weakly stimulated group (measured 10 hours after the induction) as a function of the time difference between the stimulation of the group receiving the strong tetanus and that receiving the weak tetanus. The curve in Figure 2F shows that for a time difference up to 1 hour there is significant interaction between the two groups of synapses leading to synaptic consolidation, whereas for time differences beyond 2 hours this is no longer the case. If the trigger threshold is increased to $N_p = 60$ (corresponding to less available neuromodulator), then the maximal time difference that still yields L-LTP in the weakly stimulated group of synapses is reduced to about 20 minutes (Figure 2F) whereas a reduction of N_p yields an increased time window of interaction (data not shown). If N_p is reduced much further, the weak tetanus alone will be sufficient to allow a transition from the early to the late phase of LTP. We speculate that N_p could depend on the age of the animal as well as on the background level of dopamine or other neuromodulators so as to enable a tuning of the degree of plasticity (see discussion for details).

LTD and Cross-Tagging

We consider two experimental protocols known to induce LTD—a weak low-frequency protocol consisting of 900 pulses at 1 Hz and a strong low-frequency protocol consisting of 900 repetitions at 1 Hz of a short burst of three pulses at 20 Hz. This strong low-frequency protocol applied to 100 model synapses leads to a significant level of LTD (reduction of weights to 70±4 percent of initial value) which is consolidated 5 hours later at a level of 83±3 percent of initial value. If a group of 100 synapses is stimulated with the weak low-frequency protocol, an early phase of LTD is induced that is not consolidated but decays over the time course of 3 hours (Figure 3A and 3B). However, if the weak low-frequency stimulation occurs after another group of 100 synapses had been stimulated by the strong low-frequency protocol, then the group that has received the weak stimulation shows consolidated synapses (at 90±2 percent 5 hours after stimulus induction, Figure 3C). Moreover, consolidation of LTD (at 92±3 percent 5 hours after stimulus induction) in the group of synapses receiving the weak low-frequency protocol also occurs if it was stimulated thirty minutes after the stimulation of a second group of synapses by a strong tetanus, leading to LTP (Figure 3D). Thus, the TagTriC model exhibits cross-tagging consistent with experiments [11,32]. In our model, cross-tagging occurs because the tags for LTP and LTD (h_i and l_i, respectively) enter in a symmetric fashion into the trigger criterion for the synthesis of plasticity-related proteins (see Figure 1 and Methods).

Model Mechanism for Tagging, Cross-Tagging, and Consolidation

In order to elucidate how the model gives rise to the series of results discussed in the preceding paragraphs, we have analyzed the evolution of the model variables during and after induction of LTP (Figure 4). Critical for consolidation is the synthesis of plasticity related proteins, characterized by the variable p in the model. Synthesis is only possible while the total number of tags $\sum_i^N h_i + l_i$ is above the protein triggering threshold N_p. For the strong tetanic stimulus this criterion is met for about 90 minutes (shaded region in Figure 4A) leading to high levels of plasticity related proteins. After 90 minutes the concentration of proteins starts to decay back to baseline. While the level of proteins is sufficiently elevated the consolidation variable z_i of each tagged synapse moves towards $z_i \approx 1$ since this is the only stable fixed point of the dynamics (Figure 1C). This leads to a consolidation time of about 2 hours, enough to switch a large fraction of synapses into the up-regulated state $z \approx 1$ (green line, Figure 4A). Hence the average weight of the stimulated synapses stabilizes at a value above baseline, indicating L-LTP (Figure 4A, solid line).

If, in a different experiment, 100 synapses are stimulated by the weak tetanus, the synthesis of plasticity related proteins is only possible during a few minutes (Figure 4B, red line), which is not sufficient to switch tagged synapses from $z = 0$ into the upregulated state $z \approx 1$. Hence the weights (Figure 4B, black line) decay together with the tags (Figure 4B, magenta line) back to baseline and the transition from early to late LTP does not occur. The decay of the weights is controlled by the rate k_H at which tags stochastically return to zero. The evolution of the protein concentration p and the consolidation variable z after a strong tetanus that leads to 90 minutes of protein synthesis and a weaker tetanus that only leads to 40 minutes of protein synthesis has been illustrated in (Figure 5A).

The total amount of available protein that is synthesized depends in our model on the time that the total number of tags stays above the protein triggering threshold N_p. Even though always 100 synapses are stimulated in our model, not all receive tags in each experiment; moreover because of the competition for potentiation tags ($h_i = 1$) and depression tags ($l_i = 1$) during induction of plasticity, different synapses can receive different tags in the same experiment. With our strong tetanus protocol, on average 70 (out of 100) synapses receive a potentiation tag and 30 a depression tag while with the weak tetanus the numbers are 30 and 10, respectively. For the depression protocols, on average 10 synapses receive a potentiation tag and 90 a depression tag under strong low-frequency stimulation, and typically zero a potentiation tag and 40 a depression tag under the weak low-frequency protocol. These numbers vary from one trial to the next so that sometimes the protein trigger threshold $N_p = 40$ is reached with the weak protocols and sometimes not. The important aspect is that even if the threshold is reached for a short time, the duration of protein synthesis is not long enough to provide a sufficient protein concentration p for consolidation of the tagged synapses; see Figure 4B and Figure 5A.

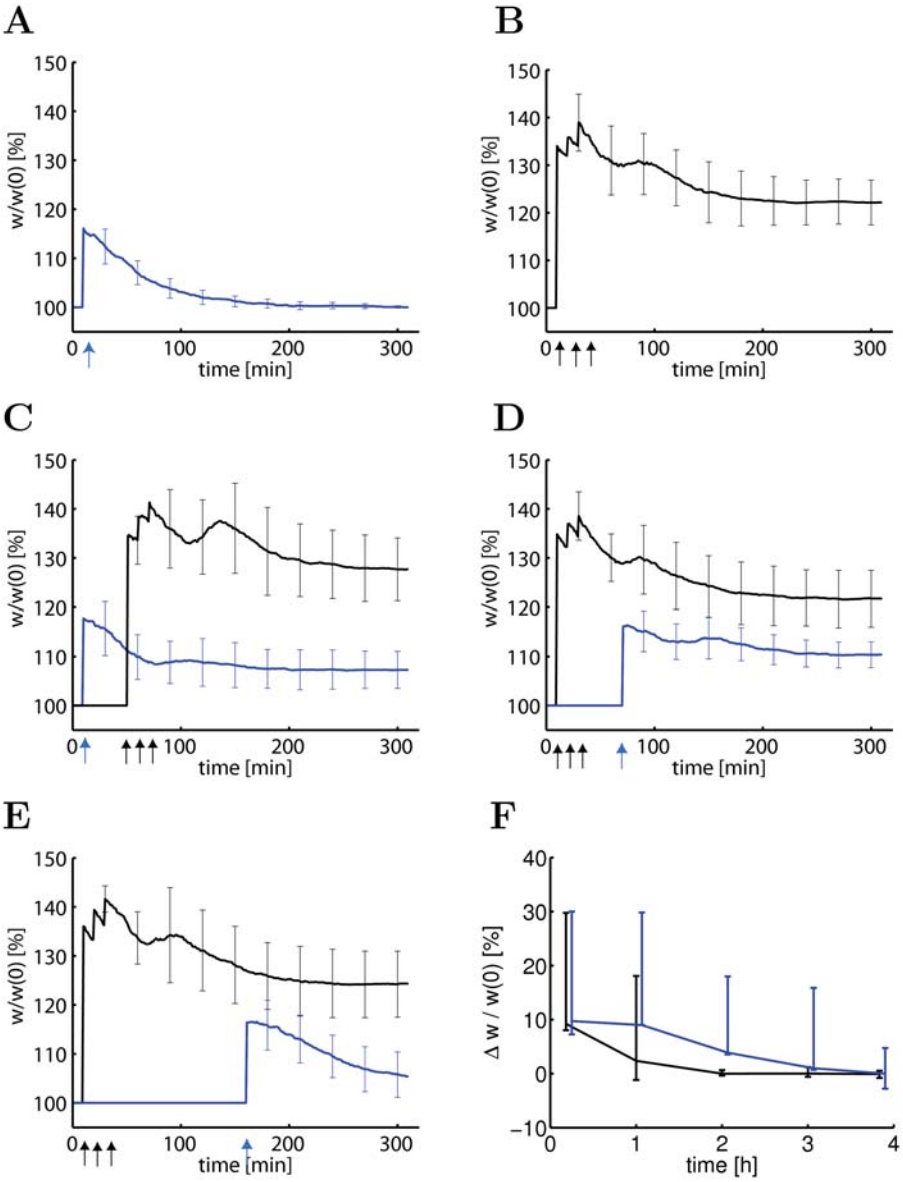

Figure 2. The model accounts for tagging paradigms. (A) A weak tetanus (21 pulses at 100 Hz) applied at a group of 100 synapses at $t = 10$ min (arrow) leads to an increased connection weight ($w/w(0)$, blue line) that decays back to baseline. (B) A strong tetanus (100 pulses at 100 Hz repeated three times, arrows) leads to late LTP that is sustained for 5 hours (black line). (C) If the weak tetanus (blue arrow) in a first group of synapses is followed thirty minutes later by a strong tetanus (black arrows) in a second group of synapses, the weights in the first group (blue line) and the second group (black line) are stabilized above baseline. (D) Stimulating a group of synapses by a weak tetanus (blue arrow) 30 minutes *after* the end of the strong tetanic stimulation of a second group also leads to stabilization of the weights in both groups above baseline. (E) If the weak tetanic stimulation occurs 2 hours after the strong tetanic stimulation of the other group, only synapses in the strongly stimulated group will be stabilized (black line), but not those in the weakly stimulated group (blue line). (F) Fraction of stabilized weights $\Delta w/w(0)$ in the weakly stimulated group measured 10 hours after induction of LTP as a function of the time difference between the weak stimulation and the end of the strong tetanic stimulation in the second group. Blue line: normal set of parameters ($N_p = 40$). Black line: protein trigger threshold increased to $N_p = 60$. In panels A–E, lines indicate the result averaged over 10 repetitions of the simulation experiments and bars standard deviation. In panel F, line indicates the result averaged over 100 repetitions. 90 of the 100 individual trials stayed within the bounds indicated by the error bars.
doi:10.1371/journal.pcbi.1000248.g002

Since the concentration p of plasticity related proteins is crucial for the transition from early to late LTP we wondered how a block of protein synthesis would interfere with the consolidation of weights in the TagTriC model. Application of a protein synthesis inhibitor (modeled by setting the rate k_p of protein synthesis to zero) during 1 hour starting thirty minutes before a strong tetanus is given to a group of 100 synapses that would normally lead to L-LTP, induced E-LTP but prevented consolidation into L-LTP (data not shown). However, if the same simulation experiment was repeated after a second group of synapses had received a strong tetanic stimulation 35 minutes prior to the application of protein synthesis blocker, then both groups of synapses showed consolidation of weights (Figure 4D), consistent with experiments [12]. Closer inspection of the lower panel in Figure 4D shows that two components contribute to consolidation: Firstly, the concentration of plasticity related proteins (red line) that has increased because of

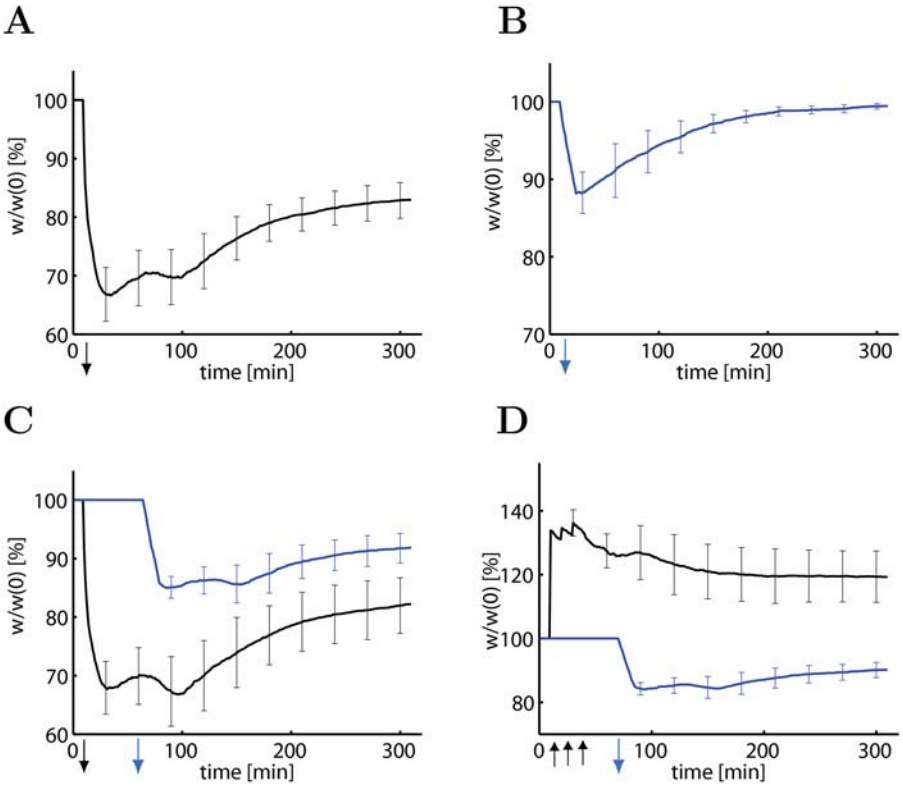

Figure 3. The model accounts for cross-tagging between LTP and LTD. (A) A strong low-frequency stimulus (3 pulses at 20 Hz, repeated 900 times every second) applied to a group of $N = 100$ synapses induces LTD with mean weights ($w/w(0)$) stabilized at $83 \pm 3\%$ of initial value after 5 hours (black line). (B) A weak low-frequency stimulus (1 pulse repeated 900 times at 1 Hz) induces early LTD, which is not consolidated. (C) If the weak low-frequency stimulus is applied 30 minutes after a second group of synapses has received the strong low-frequency protocol, the weights in both groups (blue, weak stimulus; black, strong stimulus) are consolidated at values below baseline. (D) Consolidation of LTD in the group receiving weak low-frequency stimulation (blue line) also happens if induction occurs 30 minutes after stimulating a second group of synapses with a strong tetanic protocol (see Figure 2) inducing LTP (black line). Downward arrows indicated the period of weak (blue arrow) or strong (black arrow) low-frequency protocols. The black upward arrows indicate strong tetanic stimulation. Lines show mean results, averaged over 10 repetitions of the simulation experiment. Error bars are standard deviation.
doi:10.1371/journal.pcbi.1000248.g003

the first strong tetanic stimulus decreases only slowly back to baseline enabling the switching of the slow components (variable z, green line) even in the presence of protein synthesis blocker. Secondly, even after the end of the application of the blocker, the total number of tags that has been set by LTP induction is still above the critical value N_p (shaded region in Figure 4D) so that protein synthesis can be resumed after the end of the blocking period. In summary, the detailed analysis of the TagTriC model allows to account for many aspects of tagging experiment in terms of a limited number of variables.

Discussion

Relation of Models to Experiments

Synaptic plasticity is based on intricate signal transduction chains involving numerous processing steps and a large number of different molecules [2,13,17]. Despite the complexity of the molecular processes, synaptic plasticity has experimentally been characterized by a small set of distinct phenomena such as short-term plasticity [44] as well as early and late phases of LTP and LTD [13].

Existing models of synaptic plasticity have focused on the description of short-term plasticity [44] and on the induction of LTP and LTD [24–26,33–36]. The question of maintenance has received much less attention and was mainly addressed in the context of bistability of the CaMKII auto-phosphorylation process [27–29], AMPA receptor aggregation [41], or four identified kinase pathways [45]. While CaMKII is necessary for induction of long-term potentiation [46], it is probably too narrow to focus modeling studies only on a single or a few kinases such as CaMKII and neglect other proteins and signaling cascades that are involved in synaptic maintenance [13]. For example, there is strong evidence that PKMζ is involved in synaptic maintenance and necessary for the late phase of LTP in vitro [11] and in vivo [14]. However, the actual processes are complex and the molecules involved in setting tags may differ between different parts of the dendrite. For example PKMζ is involved in setting tags during E-LTP in the basal dendrite, whereas CaMKII (or MAPK for E-LTD) plays a similar role in apical dendrites [30].

Instead of focusing on specific signaling cascades, the TagTriC model presented in this papers aims at describing the essential ingredients of any possible functional model of L-LTP and tagging. These ingredients include (i) a bistable switch (described by the dynamics of the z_i-variable) for each synapse that guarantees long-term stability in the presence of molecular turnover [16]; (ii) a global triggering signal for protein synthesis (described by the dynamics of the p variable); a formalism to (iii) induce early forms of LTP and LTD and (iv) set synaptic tags. Since we aimed for the simplest possible model, we have identified the synaptic tags h_i and l_i for potentiation and depression with the

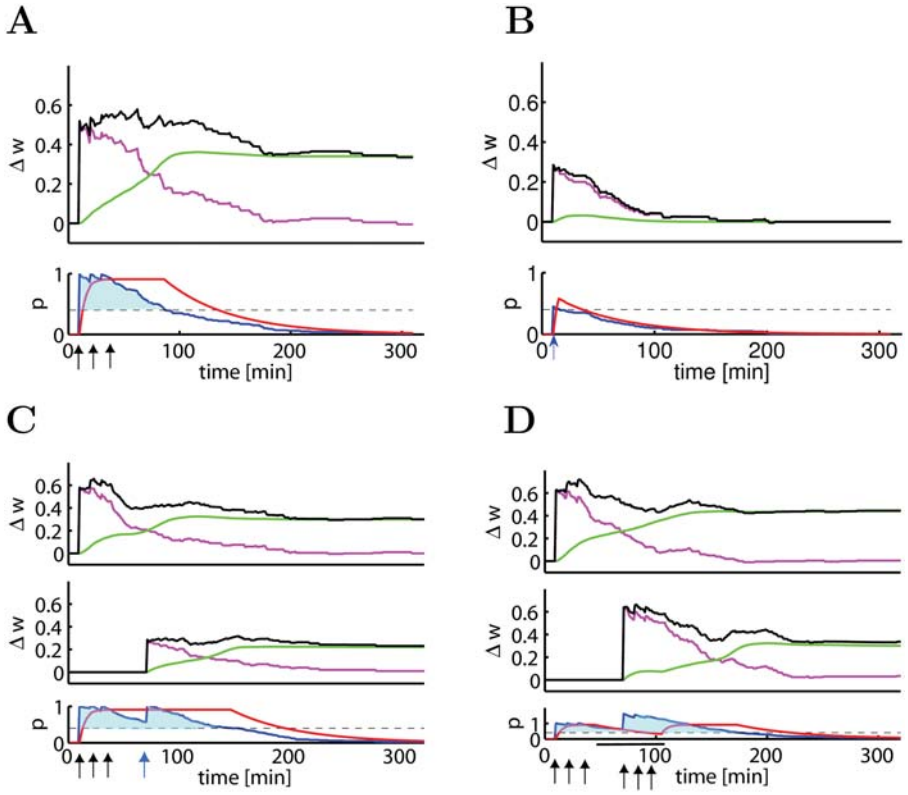

Figure 4. Dynamics of the TagTriC Model during different tagging protocols and protein synthesis blocking. The change of the total synaptic weight (top panels, black line $\Delta w = \sum_{i=1}^{N} [w_i(t) - w_i(0)/N]$) has contribution from early LTP (top panels, magenta line represents $\sum_{i=1}^{N} (h_i - \alpha l_i/N)$) and from late LTP (top panels, green line represents $\sum_{i=1}^{N} \beta(z_i - z_i(0))/N$). The protein variable p (red line, bottom panels) grows as long as the average number of tags ($\sum_{i=1}^{N} (h_i + l_i)/N$, blue line) is above the protein synthesis trigger threshold (N_p/N, dashed horizontal line). For better visibility, the regions where the blue line is above the trigger threshold is shaded. (A) A strong tetanus ($N = 100$ synapses, stimulated by 100 pulses at 100 Hz, repeated three times every ten minutes) leads to a sustained period of about 90 minutes where the number of tagged synapses is above the protein synthesis triggering threshold (lower panel, blue shaded). During this time the protein synthesis variable p is close to one (red line, lower panel), causing an increase in the fraction of consolidated weights (green line, top panel). (B) During a weak tetanus ($N = 100$ synapses, stimulated by 21 pulses at 100 Hz) the number of tags surpasses the protein triggering threshold only for a short time which does not enable switching of the z variable (top panel, green line) to the up-regulated state. (C) If the weak tetanus is given 30 minutes after the strong one, the number of tags set by the strong tetanus is still above the threshold, which allows protein synthesis stabilizing both the group of 100 synapses receiving the strong tetanus (top panel) and the group of 100 synapses receiving the weak tetanus (middle panel). (D) Protein synthesis is blocked for 1 hour (indicated by black bar at bottom of panel) starting 35 minutes after a first group of 100 synapses has been stimulated by a strong tetanus. Despite protein synthesis blocking, both the first group of synapses (top panel) and a second group of 100 synapses that received a strong tetanus during the blocking period (middle panel) develop late LTP because proteins synthesized during the induction of early LTP in the first group decay only slowly (bottom panel).
doi:10.1371/journal.pcbi.1000248.g004

synaptic weights during the early phase of LTP and LTD, respectively, so that points (iii) and (iv) are described by the same transition of the synapse from an initial non-tagged state to the high or low state, respectively. Variants of the model where the weight during the early phase of LTP and LTD is not directly proportional to the value of the tags are conceivable.

Even though we do not want to identify the synaptic variables h_i, l_i, z_i with specific biochemical signals, a couple of candidate molecules and signaling chains should be mentioned. The setting of the tag for LTP under normal physiological conditions involves NMDA receptor activation and elevated levels of calcium which in turn trigger a signaling chain involving Calmodulin and CaMKII. We therefore think that the h_i variable (representing both the tag for LTP induction and the weight increase during the early phase of LTP) should be related to the activation of CaMKII [13,46]. The molecular interpretation of the tag l_i for LTD is less clear [13]. In our model we have taken the tags as discrete quantities that decay stochastically, but a model with continuous tags that decrease exponentially gives qualitatively the same results (data not shown). The reason is that triggering protein synthesis in our model requires a large number of tags to be set, so that even in the stochastic model only the *mean* number of tags is relevant–and the mean (more precisely, its expectation value) is a continuous variable. Nevertheless, we prefer the model with discrete values over the continuous one in view of the switch-like transitions of synapses after induction of LTP and LTD [7,37]. Maintenance of enhanced synaptic weights is probably implemented by an increased number of AMPA receptors in the postsynaptic membrane. Whether the stability arises from a self-organization process of receptors [41] or from interaction with persistently activated CaMKII molecules [46] or from additional kinases such as PKMζ [11,14], is an open problem of experimental investigation. Similarly, the exact identity of many plasticity related proteins is still unknown [13]. In our model we assume that recently synthesized plasticity related proteins are accessible to all synapses onto the same postsynaptic neuron. However, a distinction between proteins synthesized in, say, basal dendrites and that synthesized in apical dendrites would be possible by

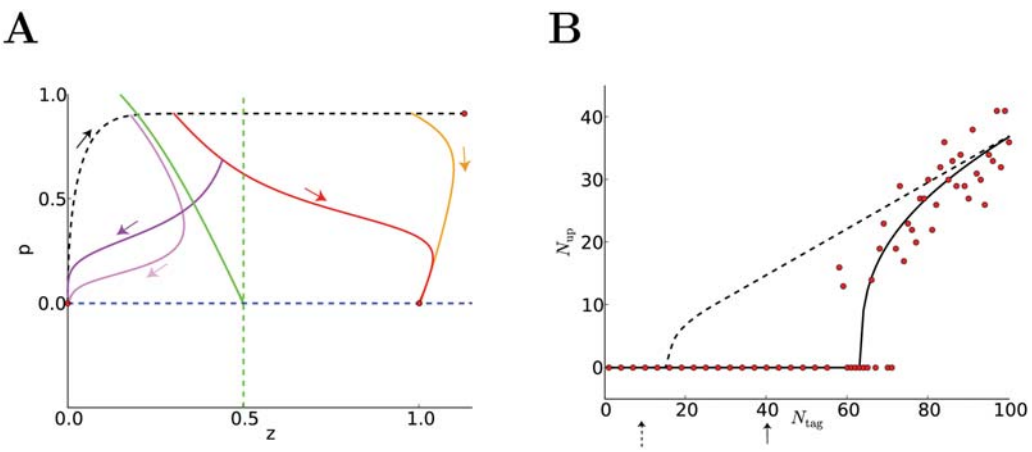

Figure 5. Theory and predictions. (A) Evolution of the variables p and z during tagging. If protein synthesis is 'ON' and the synapse tagged, p and z move along the black dashed line towards the stable fixed point on the upper right ($p \approx 1$, $z \approx 1$) (red filled circle). If protein synthesis stops after some time (yellow line, after 90 min; orange line, after 40 minutes) but the synapse remains tagged, the dynamics converges towards the fixed point $p = 0$, $z = 1$ (red filled circle) indicating that the synapse is consolidated (yellow and orange trajectories). However, if protein synthesis stops too early (after 25 min, pink line), or if the synaptic tag is lost too early (after 60 min, magenta line), the synapse is not consolidated and the trajectories converge towards the non-tagged initial state $p = 0$, $z = 0$ (red filled circle). The green dashed vertical line at $z = 0.5$ indicates the threshold beyond which a loss of the tag does not affect consolidation; the green solid line indicates the separatrix between the stable fixed points at $z = 0$ and $z = 1$. The minimal duration of protein synthesis to allow any consolidation is given by the intersection of the black dashed line with the separatrix. (B) Number of consolidated synapses (N_{up}, vertical axis) as a function of the number of initially tagged synapses (N_{tag}, horizontal axis) in simulations (red filled circles) and theory (solid line). Some of the initially tagged synapses fail to be consolidated because either they lose their tag or protein synthesis stops too early (see A). With a protein synthesis threshold $N_p = 40$ (arrow) we need about 60 initially tagged synapses to achieve any consolidation (solid line). If the protein synthesis threshold is reduced to $N_p = 10$ (dashed arrow), we need at least 15 tagged synapses to see any consolidation (dashed line).
doi:10.1371/journal.pcbi.1000248.g005

replacing the variable p by two or more distinct variables p_k with similar dynamics (but potentially different trigger thresholds N_p), allowing for a compartmentalization of tagging [13].

Experimental cross-tagging results clearly indicate that there are two different types of synaptic tags, one for LTP and one for LTD [13,32], which we called h_i for LTP and l_i for LTD, leading to three different states during tagging (Figure 1A). Since we have identified the tagging with the early phase of LTP and LTD, our model of E-LTP and E-LTD also has three different states (whereas our model of late LTP/LTD has only two states characterized by $z_i = 0$ and $z_- = 1$). The three-state model of early LTP/LTD presented in this paper would predict that all non-tagged synapses can undergo a transition to E-LTP or E-LTD depending on the induction protocol–whereas experiments suggest that about 70 percent of synapses show LTP but not LTD and the remaining 30 percent LTD but not LTP [7]. Moreover, only those synapses that are initially weak can be potentiated and only those that are initially strong can be depressed [7]. This aspect can be included in our model if we replace the induction rates ρ_H for LTP by $\rho_H(1-z_i)$ and ρ_L for LTD by $\rho_l z_i$ so LTP is only possible from a state with $z_i = 0$ and LTD only from an initial state $z_i = 1$ — in agreement with a two-state model of early LTP/LTD [7]. For the tagging and induction experiments presented in this paper, the results do not change significantly when we implement this extension of the induction model.

Functional Consequences and Predictions

One of the advantages of a simple phenomenological model is that it should be capable of illustrating the functional consequences of tagging and L-LTP or L-LTD in a transparent manner. What are these functional consequences?

A characteristic feature that is made transparent in our model (and which we expect to be present in any model of tagging) is that, under typical experimental conditions, the transition from early to late LTP is only possible if a sizable group of synapses have undergone E-LTP or E-LTD. Hence, while induction of E-LTP is a local Hebbian process that is likely to take place at the postsynaptic site of the synapse (e.g., the dendritic spine), the transition from the early to the late phase of LTP requires a minimum number of synapses to be activated by appropriate stimulation including co-activation of neuromodulatory input so as to trigger synthesis of plasticity related proteins. A direct consequence of this is that synapses cannot be considered as independent. In order to predict whether a synapse memorizes an item for a long time or forgets it and re-learns some other item, it is not sufficient to consider a 'Hebbian' induction model, where synaptic changes depend only on the activity of pre- and postsynaptic neurons. For maintenance, it is not the synapse which decides individually, but it is the neuron as a whole (or a large functional compartment sharing the same site of synthesis of plasticity-related proteins [13,30,47]) which 'decides' whether it is going to store the present information, or not. Hence, classical [20,21,34] and recent [22] theoretical models which studied memory maintenance in the presence of ongoing neuronal activity on the level of *single* synapses need to be reconsidered, since the assumption of independent synapses does not hold (Figure 5A and 5B). In particular, our model predicts that, after an ensemble of identical neurons have received the same stimulus, some neurons learn (adapt a *large* fraction of their synapses to the stimulus) and others don't (keep all their synapses unchanged). With our choice of parameters, this happens in the TagTriC model if the number of synapses that have been tagged during the induction protocol is between 55 and 70 (Figure 5B). This neuronal, rather than synaptic, decision about memorizing an input (see also [48]) is potentially attractive for prototype learning–a standard paradigm in neuronal clustering and categorization algorithms, e.g., [19]. In contrast to traditional neuronal clustering models where learned

memories need to be protected against overwriting by completely different memory items [19], a model based on tagging would have an intrinsic vigilance threshold via the trigger threshold N_p. Hence it is resistant to changes at a single synapse.

In our view, the protein synthesis trigger threshold N_P is an important control parameter in the model. The results of Figure 2F show that an increase of the trigger threshold reduces the maximal delay after which a weak tetanus leads to L-LTP after a strong tetanic stimulation in a different group of synapses. With our normal value of $N_p = 40$ we need around 60 synapses to be initially tagged in order to retain any memory. If we decrease the trigger threshold to $N_p = 10$ and keep all other parameters of the model unchanged, then we need at least a group of 15 synapses tagged during the induction protocol to get any consolidation since some of the initially tagged synapses loose their tag too early to get consolidated (Figure 5B). Only for a very small trigger threshold, say $N_p = 1$, (which could occur at high concentration of neuromodulators) synapses become (nearly) independent, since a tag at a single synapse would be sufficient to trigger the synthesis of proteins which would then become available at that synapse. Repeated stimulation of the synapse alone would then be sufficient to transform E-LTP into L-LTP.

In our opinion, the trigger threshold N_p is significantly lower in the presence of neuromodulators such as, for example, dopamine (for synapses from Schaffer collaterals onto CA1 pyramidal neurons) or noradrenaline (for synapses in the dentate gyrus). A simple model for the dependence of N_p on dopamine would be $N_p = n_0/(\mathrm{DA}_{\mathrm{bg}}+c_0)$ where n_0 is some arbitrary number (say $n_0 = 1$), c_0 a small number (say 0.001) and DA denotes the stationary 'background' concentration of dopamine (that is, before the start of the experiment), normalized to $0<\mathrm{DA}_{\mathrm{bg}}<1$. The phasic dopamine signal caused by co-stimulation of dopaminergic input during tagging experiments is assumed to be proportional to the number of tags $\sum_i^N h_i + l_i$. The trigger condition $\sum_i^N h_i + l_i > N_p$ becomes then equivalent to the condition $\left(\sum_i^N h_i + l_i\right)\left(\mathrm{DA}_{\mathrm{bg}} + c_0\right) > n_0$ which shows a trade-off between the phasic dopamine signal and the stationary background level of dopamine. In particular in the presence of a large concentration of dopamine ($\mathrm{DA} \approx 1$), single synapses can be consolidated. With the assumption that standard tagging experiments in a large group of synapses are performed at a low dopamine concentration of $\mathrm{DA} = 0.024$ before stimulation, we retrieve the value of $N_p = 40$ used in the main part of the results section. The dependence of the trigger criterion on the number of tags $\sum_i^N h_i + l_i$ takes implicitly the co-activation of neuromodulatory input during the experimental stimulation protocol into account: the larger the number of stimulated neurons and the stronger the stimulus, the higher the probability of co-activation of dopaminergic fibers. Blocking dopamine receptors amounts in the model to setting both the background and the phasic dopamine signal to zero. In this case, protein synthesis is not possible.

Our model of LTP/LTD induction does not only account for voltage and frequency dependence of LTP/LTD induction, but also for spike timing dependence. In fact, for a stimulation paradigm where postsynaptic spikes are induced by short current pulses of large amplitude either a few milliseconds before or after presynaptic spike arrival, the model of LTP/LTD induction used in the TagTriC model becomes formally equivalent to a recent model of spike-timing dependent plasticity [35] which can be seen as an extension of classical models of STDP [24–26]. In the case of stochastic spiking of pre- and postsynaptic neurons our model shares important features with the Bienenstock-Cooper-Munro model [33], in particular the quadratic dependence upon the postsynaptic variables. In addition, our model also accounts for the voltage dependence of the Artola-Bröcher-Singer model [38]. Thus, the model of LTP/LTD induction shares features with numerous established theoretical models and covers a large range of experimental paradigms known to induce LTP or LTD [3–6,8].

Since the subsequent steps of protein synthesis trigger and stabilization are independent of the way early phase of LTP is induced, our model predicts that tagging experiments repeated with different stimulation paradigms, but otherwise identical experimental preparation and age of animal, should give similar results as standard tagging protocols. In particular we propose to stimulate a group of synapses in hippocampal slices by 40–60 extracellular current pulses at 10 Hz while the postsynaptic neuron is receiving intracellular current injection that triggers action potential firing either a few milliseconds before or after presynaptic spike arrival and keeps the membrane potential at a depolarized level between postsynaptic action potential firing. Our model predicts that this will induce early LTD or LTP depending on spike timing and depolarization level that is not maintained beyond 1 or 2 hours. However, if the same stimulation occurs after a second group of synapses has received a strong tetanus, then stabilization of synapses at potentiated or depressed levels should occur, similar to standard tagging and cross-tagging experiments. In our opinion, these predictions should not depend on model details, but hold for a broad class of models that combine a mathematical description of induction of synaptic plasticity with a mechanism of consolidation.

Another finding—which is somewhat unexpected and in contrast to other conceptual models of synaptic tagging and capture [12,13,47]—is that during a strong tetanic stimulation a fraction of synapses receives tags for depression (while most, but not all, receive tags for potentiation). This is due to the fact that during induction of plasticity, transition to E-LTP and E-LTD act in parallel [7]. The prediction is that after consolidation (say 2 hours after the strong tetanic stimulation) a small fraction of synapses would show L-LTD, rather than L-LTP.

An essential ingredient of our model that allows long-term stability of consolidated synapses is the bistable dynamics of the variable z. In our opinion, such bistability (or possibly multi-stability [49] with three or four stable states) is necessary for synaptic maintenance in the presence of molecular turn-over, as recognized in earlier theoretical work [15,16,34]. Our model therefore predicts that L-LTP and L-LTD should have bistable, switch-like properties. While there is evidence for switch like transitions during the induction of E-LTP and E-LTD [7,37], the bistability of the late phase of synaptic plasticity has so far not been shown. A possible experiment would be to combine a minimal stimulation protocol (e.g., a weak tetanus) at a single synapse [7,37] with a medium to strong stimulus at a group of other synapses (e.g., tetanic stimulus varying between 30 and 100 pulses). The prediction is that the weight of the single synapse shows an all-or-none phenomenon with transition probabilities that depend on the stimulation of the group of other synapses. In particular, as the number of pulses of the tetanic stimulation is reduced (covering a continuum from strong to weak tetanic stimulation), the maintenance in the potentiated state should become less likely (averages across many experiments decrease) whereas the results of individual experiments show either full potentiation or none, which should give rise to a bimodal distribution of normalized synaptic weights.

Open Questions and Perspectives

A lot of questions remain open and need to be addressed in future studies. First, can a synapse that has been potentiated in the past and is maintained after a transition to late LTP undergo a

further potentiation step [13]? In our current model this is not possible since the consolidation variable z has only two stable fixed points. If we replace the function $f(z)$ depicted in Figure 1 by another one with more than two stable fixed points, then the answer to the above question would be positive. Indeed, there have been suggestions that self-organization of receptors into stable sub-groups could lead to multiple stable states [49].

Second, induction of LTP or LTD is not only possible by strong extracellular stimulation of groups of synapses, but also at single synapses if presynaptic activity is paired with either a depolarization of the postsynaptic membrane [5,7] or tightly timed postsynaptic spikes as in STDP experiments [6,8]. How can it be that the change induced by STDP seems to be maintained over one hour without visible degradation? [6,7]. Are synapses in these experiments consolidated, and if so what is the concentration of neuromodulators? In the TagTriC model with the choice of parameters used in the present paper, consolidation would not be possible, since the minimum number of synapses that have undergone E-LTP or LTD is $N_p = 40$ in order to trigger protein synthesis, but, as explained above, an increased neuromodulator concentration would make consolidation possible.

Third, what is the role of NMDA receptor activation during synaptic consolidation? In our present model, protein synthesis is triggered by appropriate induction protocols, but is independent of synaptic activity during the consolidation process. However, recent experimental results suggest that protein synthesis blocker needs synaptic stimulation during the consolidation period to become effective [50], suggesting a subtle interplay between protein synthesis and synaptic activation that cannot be captured by our model.

Fourth, has each neuron a single protein synthesis unit or is protein synthesis a local process confined to each dendritic branch? In the first case, there is a single neuron-wide protein synthesis trigger threshold [12] and the neuron as a whole 'decides' whether early forms of synaptic potentiation and depression will be consolidated or not. This is the paradigm posited in the TagTriC model. In the alternative model of local protein synthesis [13,47], the critical unit for consolidation are local groups of synapses on the same dendritic branch. Thus, for the same number of tagged synapses, a local group of synapses on the same dendritic branch is more likely to undergo consolidation than a distributed set of tagged synapses, leading to a form of clustered plasticity [47]. The TagTriC model can be easily adapted to the case of clustered plasticity by (i) replacing the point-neuron model by a neuron model with spatially distributed synapses and (ii) replacing the neuron-wide trigger equation (see 4 and Figure 1B) by a finite number of analogous, but dendrite-specific equations.

Fifth, how can tags be reset? Experiments show that a depotentiating stimulus given 5 minutes after a weak tetanus erases the trace of E-LTP (resets the tag) whereas depotentiation 10 or 15 minutes after the strong tetanus only transiently suppresses the E-LTP, making the consolidation of the synapse by protein capture possible [51]. We have checked in additional simulations that our present model cannot account for these experiments. In our opinion, the above tag-reset experiments show that the synapse has additional hidden states currently not included in the TagTriC model. Additional states would allow to (i) separate the measured early LTP during the first 5 minutes from setting the tag; and (ii) distinguish between depotentiation and depression of synapses. One interpretation of the tag-reset experiments [51] is that during the first five minutes the tag is not yet set whereas early LTP is already visible. The tag would be set only with a delay of 5–10 minutes. Application of a depotentiating stimulus more than 10 minutes later would then leave the potentiation tag intact, but move the synapse to a transiently depotentiated state.

The final and potentially most interesting question is that of functional relevance: Can the TagTriC model be used to simulate reward-based learning in experiments in vivo [13]? The formal theory of reinforcement learning makes use of an eligibility trace [52] which can be interpreted as a synapse specific tag. In the future we want to check whether the TagTriC model can be linked to reinforcement learning models [53–56] under the assumption that reward prediction errors are represented by a dopamine signal [57] which influences the protein synthesis dynamics in our model. This open link to reward-based learning is of fundamental functional importance.

Methods

Model of Early LTP/LTD and Tagging

In our model we assume that presynaptic spike arrival needs to be combined with a depolarization of the postsynaptic membrane (e.g., [5]) in order to induce a change of the synapse. In voltage clamp experiments (e.g., [39]) the postsynaptic voltage would be constant. However, in general the voltage is time-dependent and described by a variable $u(t)$. In the TagTriC model, we assume that the low-pass-filtered voltage

$$\bar{u}(t) = \frac{1}{\tau_{\mathrm{lowP}}} \int_0^\infty \exp\left(-\frac{s}{\tau_{\mathrm{lowP}}}\right) u(t-s-\varepsilon)ds.$$

needs to be above a critical value ϑ_{LTD} to make a change of the synapse possible. τ_{lowP} is the time constant of the low-pass filter and $\varepsilon = 1$ ms is a short delay twice the width of a spike (see Table 1). This short delay ensures that $\bar{u}$ includes effects of previous presynaptic inputs and postsynaptic spikes, but not of an ongoing postsynaptic action potential.

Table 1. Parameter values used throughout all simulations, except Figure 1E–G where $N_p = 10$ and initial percentage of $z_i = 1$ was 10%, because these simulations refer to experiments with younger animals.

Tag	Trigger	Consolidation
$N = 100$	$k_p = 1/(6\ \mathrm{min})$	$N = 100$
$A_{\mathrm{LTD}} = 0.01$	$\tau_p = 60$ min	$\gamma = 0.1$
$A_{\mathrm{LTP}} = 0.014$	$N_p = 40$	$\tau_z = 6$ min
$\tau_x = 100$ ms		$\beta = 2$
$\tau_{\mathrm{lowP}}^{\mathrm{LTP}} = 100$ ms		Initialisation: $N(z_i = 1) = 30$
$\tau_{\mathrm{lowP}}^{\mathrm{LTD}} = 1$ s		
$\varepsilon = 1$ ms		
$k_h = 1/\mathrm{h}$		
$k_l = 1/(1.5\ \mathrm{h})$		
$\Theta_{\mathrm{LTD}} = -70.6$ mV		
$\Theta_{\mathrm{LTP}} = -50$ mV		
$\alpha = 0.5$		
Initialisation: $l_i = h_i = 0$		

doi:10.1371/journal.pcbi.1000248.t001

Combining presynaptic spike arrival at synapse i (represented by x_i) with a depolarization $\bar{u}$ of the postsynaptic neuron above a threshold ϑ_{LTD} we get a rate of LTD

$$\rho_L = A_{\mathrm{LTD}} x_i(t)[\bar{u}(t) - \vartheta_{\mathrm{LTD}}]^+ \quad (1)$$

where $A_{\mathrm{LTD}}>0$ is a parameter and $[.]^+$ denotes rectification, i.e., $[y]^+ = y$ if $y>0$ and zero otherwise. Here $x_i(t) = \sum_f \delta\left(t - t_i^f\right)$ denotes the presynaptic spike train with pulses at time t_i^f and δ the Dirac-delta function. Formally, ρ_L describes the rate of stochastic transitions from the non-tagged state $h=0$, $l=0$ to the low state $l=1$, Figure 1. In simulations we work with discrete time steps of $\Delta = 1$ ms. Eq. 1 indicates that the probability $P_{l=0\to l=1}$ of a transition to the low-state during the time step Δ vanishes in the absence of presynaptic spike arrival and takes a value of $P_{l=0\to l=1} = 1-\exp(-A_{\mathrm{LTD}}[\bar{u}(t)-\vartheta_{\mathrm{LTD}}]^+\Delta) \approx A_{\mathrm{LTD}}[\bar{u}(t)-\vartheta_{\mathrm{LTD}}]^+\Delta$ if a presynaptic spike arrives at the synapse i during the time step Δ. Note that the transition from $l=0$ to $l=1$ is only possible if $h=0$ and h remains zero during the transition.

Similarly, a switch from the non-tagged state $h=0$, $l=0$ to the high state $h=1$ occurs at a rate ρ_H which also depends on postsynaptic voltage and presynaptic spike arrival. We assume that each presynaptic spike at synapse i leaves a trace $\bar{x}_i$ that decays exponentially with time constant τ_x. The exact biophysical nature of the trace is irrelevant, but could, for example, represent the amount of glutamate bound to the postsynaptic receptor. The value of the trace at time t caused by earlier spike arrivals at time t_i^f is then $\bar{x}_i(t) = (1/\tau_x)\sum_f \exp\left[-\left(t-t_i^f\right)\Big/\tau_x\right]$ where the sum runs over all firing times $t_i^f < t$. With the trace $\bar{x}_i$ we write

$$\rho_H = A_{\mathrm{LTP}}\bar{x}_i(t)[\bar{u}(t) - \vartheta_{\mathrm{LTD}}]^+ [u(t) - \vartheta_{\mathrm{LTP}}]^+ \quad (2)$$

which indicates that, in addition to the conditions for LTD induction we also require the *momentary* membrane potential $u(t)$ to be above a second threshold ϑ_{LTP}. This threshold could change on the time scale of minutes or hours as a function of homeostatic processes. To summarize, the rate of LTP transition ρ_H is different from ρ_L in five aspects. First, the constant A_{LTP} is not the same as A_{LTD}. Second, LTP is caused by the *trace* $\bar{x}_i$ left by presynaptic spikes, rather than the spikes themselves. This trace-formulation ensures that presynaptic spikes can interact with later postsynaptic spikes as in classical models of STDP [24–26]. Third, the time constant of the low-pass filter in $\bar{u}$ is different; fourth, the momentary voltage needs to be above a threshold ϑ_{LTP}; and fifth, the total dependence upon the postsynaptic voltage is quadratic, rather than linear. The quadratic dependence ensures that for large depolarization LTP dominates over LTD [39]. Tagged synapses with $h_i=1$ decay with probability $P_{h=1\to h=0} = k_H\Delta$ back to the non-tagged state (and analogously, but with rate k_L for the transition $l_i=1 \to l_i=0$).

In the TagTriC model, the local synaptic values $h=1$ for potentiation or $l=1$ for depression act as tags indicating potential sites for further consolidation, but are also directly proportional to the weight of the synapse after induction of LTP or LTD. Since in minimal stimulation experiments LTD leads to a reduction of about 50 percent of the synaptic efficacy whereas LTP leads to an increase by up to 100 percent [7], we model the weight change during the early phase of LTP as $\Delta w_i = (h_i - \alpha l_i)\hat{w}$ where $\hat{w}$ is the weight of the non-tagged synapse and $\alpha = 0.5$. The total weight change $\Delta w/\hat{w}$ measured shortly after induction of LTP or LTD with extracellular protocols corresponds to the fraction of synapses in the high or low states, respectively, hence, if all synapses start from the non-tagged state the measured weight change is $\Delta w\Big/\hat{w} = \sum_{i=1}^N (h_i - \alpha l_i)/N = \langle h\rangle - \alpha\langle l\rangle$ where N is the number of synapses stimulated by the protocol. The set of parameters of LTP/LTD induction and tagging is given in table 1.

Trigger

The triggering process is controlled by the dynamics of a variable p which describes the amount of plasticity related proteins synthesized in the postsynaptic neuron. Protein synthesis is triggered and the variable p increases while the concentration of dopamine exceeds a critical level ϑ_p [58]. If the dopamine concentration DA falls below ϑ_p, the protein concentration decays with a time constant τ_p. Assuming standard first-order kinetics we have

$$\frac{dp}{dt} = k_p(1-p)\Theta\left[\mathrm{DA} - \vartheta_p\right] - \frac{p}{\tau_p} \quad (3)$$

Protein synthesis has a maximum rate dp/dt of k_p and saturates if the amount of protein approaches a value one. $\Theta[y]$ denotes the unit step function with $\Theta[y] = 1$ for $y>0$ and zero otherwise.

Dopamine is present at a low stationary background value. In addition a phasic dopamine component is induced in standard tagging experiments in hippocampal slices, because of co-stimulation of dopaminergic inputs during extracellular stimulation of presynaptic fibers [40]. To describe the time course of the phasic dopamine component in our model, we assume that the dopamine is proportional to the total number of tags $\Sigma_i(h_i+l_i)$ induced by the stimulation protocol. The stationary background level of dopamine $\mathrm{DA_{bg}}$ is included in the threshold $\vartheta_p = N_p(\mathrm{DA_{bg}})$ for protein synthesis. Hence Eq. 3 can be rewritten in the form

$$\frac{dp}{dt} = k_p(1-p)\Theta\left[\sum_i (h_i + l_i) - N_p\left(\mathrm{DA_{bg}}\right)\right] - \frac{p}{\tau_p} \quad (4)$$

Note that we have chosen units so that the threshold for protein synthesis N_p can be interpreted as the minimal number of tags necessary to stimulate protein synthesis. This interpretation is important for the discussion of the model results, in particular Figures 4 and 5.

A suitable model for dependence of the protein synthesis threshold on the background level of dopamine is $N_p(\mathrm{DA_{bg}}) = n_0/(\mathrm{DA_{bg}}+c_0)$ where $n_0 = 1$ is a scaling factor, $c_0 = 0.001$ a constant and $0 \le \mathrm{DA_{bg}} \le 1$ is the normalized dopamine concentration. We note that the trigger condition $[\Sigma_i(h_i+l_i) - N_p(\mathrm{DA_{bg}})] > 0$ is then equivalent to the condition $(\mathrm{DA_{bg}}+0.001)[\Sigma_i(h_i+l_i)] > 1$. This formulation shows that there is a trade-off between background levels and phasic dopamine. Unless stated otherwise we always use in the simulation a fixed dopamine level $\mathrm{DA_{bg}} = 0.024$ so that $N_p = 40$. The specific model $N_p(\mathrm{DA_{bg}})$ of the dependence upon background dopamine levels is therefore irrelevant.

We assume that the plasticity related protein p synthesized in the postsynaptic neuron is diffused in the dendrite of the postsynaptic neuron and hence available to all the synapses under consideration. Hence, the tags h_i and l_i have indices, since they are synapse-specific, whereas p in Eq. 4 does not.

Consolidation and Late LTP

The consolidation variable z describes the late phase of LTP and follows the dynamics

$$\tau_z \frac{dz_i}{dt} = f(z_i) + \gamma(\mathrm{DA})(h_i - l_i)p. \quad (5)$$

The scaling factor γ is a function of the dopamine level *DA*. In the simulations we always assumed a fixed dopamine level and set $\gamma(\mathrm{DA})=0.1$.

In the absence of plasticity related proteins ($p=0$), or if no tags are set ($h_i=l_i=0$), the function $f(z)=z(1-z)(z-0.5)$ generates a bistable dynamics with stable fixed points at $z=0$ and $z=1$ and an unstable fixed point at $z=0.5$ marked by the zero crossings of the function f, Figure 1C. In the presence of a finite amount of proteins $p>0$ and a non-zero tag, the location of the fixed points changes and for $p>0.47$, only one of the stable fixed points remains. The potential shown in Figure 1C is a function E with $dE/dz=-f(z)$ so that $dz/dt=-dE/dz$. We note that a synapse i can change its consolidated value only if both a tag ($h_i=1$ or $l_i=1$) and protein $p>0.47$ is present–summarizing the essence of 'synaptic tagging and capture' [12,13].

Synaptic Weight

The synaptic weights have contributions from early and late LTP and LTD. The total synaptic weight of a synapse i is $w_i=\hat{w}(1+h_i-\alpha l_i+\beta z_i)$ where $\hat{w}$ is the value of a non-tagged synapse, $\alpha=0.5$ and $\beta=2$ are parameters, h_i and l_i are binary values indicating E-LTP and E-LTD, respectively, and z_i is the value of the L-LTP trace of synapse i. Since we model slice experiments in animals older than 20 days, we assume that 30 percent of the synapses have undergone previous potentiation and have $z=1$ while the remaining 70 percent of synapses are in the state $z=0$ [7]. In all simulation experiments we stimulate one or several groups of $N=100$ synapses each. Assuming that no tags have been set in the recent past ($h_i=l_i=0$), the initial value of the average weight in a group of N synapses is then $w(0)=\hat{w}\left[\sum_{i=1}^{N}1+\beta z_i\right]\Big/N=1.6\hat{w}$.

Neuron Model

For all simulations in this paper we use the adaptive exponential integrate-and-fire model [42] as a compact description of neuronal firing dynamics. Briefly, it consists of two equations. The voltage equation has an exponential and a linear term as measured in experiments [59]. The second equation describes adaptation. Although firing rate adaptation is not important for the present study, it would be relevant in the context of other stimulation paradigms. Parameters for the neuron model are as in [42] and are kept fixed for all simulations presented in this paper. The voltage threshold V_s of spike initiation by a short current pulse is 25 mV above the resting potential of -70.6 mV [42]. Synaptic input is simulated as a short current pulse. The initial connection weight $\hat{w}$ was adjusted so that simultaneous activation of 40 or more synapses triggers spike firing in the postsynaptic neuron. Hence the amplitude of a single EPSP is about 0.6 mV.

The adaptive exponential integrate-and-fire model is defined in continuous time. If a spike is triggered by a strong current pulse, the voltage rises within less than 0.5 millisecond to a value of 20 mV where integration is stopped. The voltage is then reset to resting level, and integration restarted after a refractory time of 1 ms. In order to enable us to perform simulations of plasticity experiments with a time step of $\Delta=1$ ms, the voltage equation during the rising slope of the action potential was integrated once at a much higher resolution (time step 0.02 ms), so as to determine the exact contribution of each postsynaptic spike to the probability of LTP induction. Every postsynaptic spike was then treated as an event in the plasticity simulations that contributed a probability $P_{h=0\to h=1}$ of flipping the tag from $h=0$ to $h=1$ in a time step $\Delta=1$ ms which we can write as $P_{h=0\to h=1}=a_\Delta\bar{x}(t)[\bar{u}(t)-\vartheta_{\mathrm{LTD}}]^+$ with a numerical conversion factor $a_\Delta=A_{\mathrm{LTP}}$ 5 ms mV derived by the above procedure; see Eq. 2.

Number of Consolidated Synapses

In Figure 5 we plot the number of synapses that have been consolidated as a function of the number N_{tag} of initially tagged ($h_i=1$) synapses. Since the number of tags decays exponentially with rate k_H, the expected duration T_P^{ON} of protein synthesis is $T_P^{\mathrm{ON}}=(1/k_H)ln\left(N_{\mathrm{tag}}/N_p\right)$ where N_p is the protein trigger threshold. While protein synthesis is 'ON' the variables p and z move along the black dashed line in Figure 5A which crosses after a time t_1 the separatrix (green line in Figure 5A) and at a time t_2 the line $z=0.5$ (vertical dashed green line). Different cases have to be distinguished. (i) $T_P^{\mathrm{ON}}<t_1$, no consolidation takes place (see pink trajectory), hence $N_{\mathrm{up}}=0$. (ii) $T_P^{\mathrm{ON}}>t_2$, consolidation is guaranteed for all synapses that are still tagged at time t_2, hence $N_{\mathrm{up}}=N_{\mathrm{tag}}\exp(-kt_2)$. (iii) In the case of $t_1<T_P^{\mathrm{ON}}\le t_2$, the time t_{cross} needed to cross the vertical line $z=0.5$ is numerically calculated by integrating the equations $dp/dt=-p/(\tau_p)$ and $dz/dt=f(z)+\gamma\ p$ starting at $t=T_P^{\mathrm{ON}}$ at the point $p\left(T_P^{\mathrm{ON}}\right),z\left(T_P^{\mathrm{ON}}\right)$ on the black-dashed line (see orange line in Figure 5A for a sample trajectory). The number of consolidated synapses is then $N_{\mathrm{up}}=N_{\mathrm{tag}}\exp(-kt_{\mathrm{cross}})$. The solid line in Figure 5B represents N_{up} as a function of N_{tag} calculated for the cases (i)–(iii). With our standard set of parameters, we have $t_1\approx28$ min and $t_2\approx60$ min.

Acknowledgments

We thank Julietta Frey and Mark van Rossum for helpful discussions.

Author Contributions

Conceived and designed the experiments: CC WG. Performed the experiments: CC. Analyzed the data: CC. Wrote the paper: WG. Designed the model of late LTP and performed research: LZ. Participated in research on a precursor model of early LTP/LTD: EV LB.

References

1. Bliss TVP, Collingridge GL (1993) A synaptic model of memory: long-term potentiation in the hippocampus. Nature 361: 31–39.
2. Malenka RC, Bear MF (2004) LTP and LTD: an embarassment of riches. Neuron 44: 5–21.
3. Bliss T, Gardner-Medwin A (1973) Long-lasting potentation of synaptic transmission in the dendate area of unanaesthetized rabbit following stimulation of the perforant path. J Physiol 232: 357–374.
4. Dudek SM, Bear MF (1992) Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. Proc Natl Acad Sci U S A 89: 4363–4367.
5. Artola A, Bröcher S, Singer W (1990) Different voltage dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. Nature 347: 69–72.
6. Markram H, Lübke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postysnaptic AP and EPSP. Science 275: 213–215.
7. O'Connor D, Wittenberg G, Wang SH (2005) Graded bidirectional synaptic plasticity is composed of switch-like unitary events. Proc Natl Acad Sci USA 102: 9679–9684.
8. Bi G, Poo M (2001) Synaptic modification of correlated activity: Hebb's postulate revisited. Annu Rev Neurosci 24: 139–166.
9. Abraham W (2003) How long will long-term potentiation last? Philos Trans R Soc Lond B Biol Sci 358: 735–744.
10. Krug M, Lössner B, Ott T (1984) Anisomycin blocks the late phase of long-term potentiation in the dentate gyrus of freely moving rats. Brain Res Bull 13: 39–42.
11. Sajikumar S, Navakkode S, Sacktor T, Frey J (2005) Synaptic tagging and cross-tagging: the role of protein kinase Mζ in maintaining long-term potentiation but not long-term depression. J Neurosci 25: 5750–5756.
12. Frey U, Morris R (1997) Synaptic tagging and long-term potentiation. Nature 385: 533–536.

13. Reymann K, Frey J (2007) The late maintenance of hippocampal ltp: requirements, phases,synaptic tagging, late associativity and implications. Neuropharmacology 52: 24–40.
14. Pastalkova E, Serrano P, Pinkhasova D, Wallace E, Fenton A, et al. (2006) Storage of spatial information by the maintenance mechanism of LTP. Science 313: 1141–1144.
15. Crick F (1984) Memory and molecular turnover. Nature 312: 101.
16. Lisman J (1985) A mechanism for memory storage insensitive to molecular turnover: a bistable autophosphorylating kinase. Proc Natl Acad Sci U S A 82: 3055–3057.
17. Newpher TM, Ehlers MD (2008) Glutamate receptor dynamics in dendritic microdomains. Neuron 58: 472–497.
18. Buonomano D, Merzenich M (1998) Cortical plasticity: from synapses to maps. Annu Rev Neurosci 21: 149–186.
19. Carpenter G, Grossberg S (1987) ART 2: self-organization of stable category recognition codes for analog input patterns. Appl Opt 26: 4919–4930.
20. Nadal JP, Toulouse G, Changeux JP, Dehaene S (1986) Networks of formal neurons and memory palimpsests. Europhys Lett 1: 349–381.
21. Amit D, Fusi S (1994) Learning in neural networks with material synapses. Neural Comput 6: 957–982.
22. Fusi S, Drew P, Abbott L (2005) Cascade models of synaptically stored memories. Neuron 45: 599–611.
23. Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci U S A 79: 2554–2558.
24. Gerstner W, Kempter R, van Hemmen J, Wagner H (1996) A neuronal learning rule for sub-millisecond temporal coding. Nature 383: 76–78.
25. Kempter R, Gerstner W, van Hemmen JL (1999) Hebbian learning and spiking neurons. Phys Rev E 59: 4498–4514.
26. Song S, Miller K, Abbott L (2000) Competitive Hebbian learning through spike-time-dependent synaptic plasticity. Nat Neurosci 3: 919–926.
27. Lisman J (1989) A mechanism for Hebb and anti-Hebb processes underlying learning and memory. Proc Natl Acad Sci U S A 86: 9574–9578.
28. Miller P, Zhabotinsky A, Lisman J, Wang X (2005) The stability of s stochastic CaMKII switch: dependenc on the number of enzyme molecules and protein turnover. PLoS Biol 3: e107. doi:10.1371/journal.pbio.0030107.
29. Graupner M, Brunel N (2007) Stdp in a bistable synapse model based on CaMKII and associate signaling pathways. PLoS Comput Biol 3: e 221. doi:10.1371/journal.pcbi.0030221.
30. Sajikumar S, Navakkode S, Frey J (2007) Identification of compartment- and process-specific molecules required for 'synaptic tagging' during long-term potentiation and long-term depression in hippocampal CA1. J Neurosci 27: 5068–5080.
31. Othmakhov N, Griffith L, Lisman J (1997) Postsynaptic inhibitors of calcium/calmodulin-dependent protein kinase type II block induction but not maintenance of pairing induced long-term potentiation. J Neurosci 17: 5357–5365.
32. Sajikumar S, Frey J (2004) Late-associativity, synaptic tagging, and the role of dopamine during LTP and LTD. Neurobiol Learn Mem 82: 12–25.
33. Bienenstock E, Cooper L, Munroe P (1982) Theory of the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. J Neurosci 2: 32–48.
34. Fusi S (2002) Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. Biol Cybern 87: 459–470.
35. Pfister JP, Gerstner W (2006) Triplets of spikes in a model of spike timing-dependent plasticity. J Neurosci 26: 9673–9682.
36. Gerstner W, Kistler WK (2002) Spiking Neuron Models. Cambridge, UK: Cambridge University Press.
37. Petersen C, Malenka R, Nicoll R, Hopfield J (1998) All-or-none potentiation of ca3-ca1 synapses. Proc Natl Acad Sci U S A 95: 4732–4737.
38. Artola A, Singer W (1993) Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. Trends Neurosci 16: 480–487.
39. Ngezahayo A, Schachner M, Artola A (2000) Synaptic activation modulates the induction of bidirectional synaptic changes in adult mouse hippocamus. J Neurosci 20: 2451–2458.
40. Frey U, Schroeder H, Matthies H (1990) Dopaminergic antagonists prevent long-term maintenance of posttetanic LTP in the CA1 region of rat hippocampal slices. Brain Res 522: 69–75.
41. Hayer A, Bhalla US (2005) Molecular switches at the synapse emerge from receptor and kinase traffic. PLoS Comput Biol 1: e20. doi:10.1371/journal.pcbi.0010020.
42. Brette R, Gerstner W (2005) Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. J Neurophysiol 94: 3637–3642.
43. O'Connor D, Wittenberg G, Wang S (2005) Dissection of bidirectional synaptic plasticity into saturable unidirectional processes. J Neurophysiol 94: 1565–1573.
44. Markram H, Wu Y, Tosdyks M (1998) Differential signaling via the same axon of neocortical pyramidal neurons. Proc Natl Acad Sci U S A 95: 5323–5328.
45. Smolen P, Baxter D, Byrne J (2006) A model of the roles of essential kinases in the induction and expression of late long-term potentiation. Biophys J 90: 2760–2775.
46. Lisman J, Schulman H, Cline H (2002) The molecular basis of CaMKII function in synaptic and behavioural memory. Nat Rev Neurosci 3: 175–190.
47. Govindarajan A, Kelleher R, Tonegawa S (2006) A clustered plasticity model of long-term memory engrams. Nat Rev Neurosci 7: 575–583.
48. Toyoizumi T, Pfister JP, Aihara K, Gerstner W (2007) Optimality model of unsupervised spike-timing dependent plasticity: synaptic memory and weight distribution. Neural Comput 19: 639–671.
49. Lisman J (2003) Long-term potentiation: outstanding questions and attempted synthesis. Phil Trans R Soc Lond B: Biol Sci 358: 829–842.
50. Fonseca R, Nägerl U, Morris R, Bonhoeffer T (2003) Neuronal activity determines the protein synthesis dependence of long-term potentiation. Nat Neurosci 9: 478–480.
51. Sajikumar S, Frey J (2004) Resetting of 'synaptic tags' is time- and activity-dependent in rat hippocampal CA1 in vitro. Neuroscience 129: 503–507.
52. Sutton R, Barto A (1998) Reinforcement Learning. Cambridge, MA: MIT Press.
53. Arleo A, Gerstner W (2000) Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. Biol Cybern 83: 287–299.
54. Pfister JP, Toyoizumi T, Barber D, Gerstner W (2006) Optimal spike-timing dependent plasticity for precise action potential firing in supervised learning. Neural Comput 18: 1309–1339.
55. Izhikevich E (2007) Solving the distal reward problem through linkage of STDP and dopamine signaling. Cereb Cortex 17: 2443–2452.
56. Legenstein R, Pecevski D, Maass W (2008) A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. PLoS Comput Biol 4: e1000180. doi:10.1371/journal.pcbi.1000180.
57. Schultz W, Dayan P, Montague R (1997) A neural substrate for prediction and reward. Science 275: 1593–1599.
58. Navakkode S, Sajikumar S, Frey J (1007) Synergistic requirements for the induction of dopaminergic D1/D5-receptor-mediated LTP in hippocampal slices of rat CA1 in vitro. Neuropharmacology 52: 1547–1554.
59. Badel L, Lefort S, Brette R, Petersen C, Gerstner W, et al. (2008) Dynamic I-V curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. J Neurophysiol 99: 656–666.

2.4 Appendix 4

A Learning Theory for Reward-Modulated Spike-Timing-Dependent Plasticity with Application to Biofeedback Robert Legenstein, Dejan Pecevski, and Wolfgang Maass PLoS Comput Biol. 2008 October; 4(10): e1000180.

A Learning Theory for Reward-Modulated Spike-Timing-Dependent Plasticity with Application to Biofeedback

Robert Legenstein[☯], **Dejan Pecevski**[☯], **Wolfgang Maass***

Institute for Theoretical Computer Science, Graz University of Technology, Graz, Austria

Abstract

Reward-modulated spike-timing-dependent plasticity (STDP) has recently emerged as a candidate for a learning rule that could explain how behaviorally relevant adaptive changes in complex networks of spiking neurons could be achieved in a self-organizing manner through local synaptic plasticity. However, the capabilities and limitations of this learning rule could so far only be tested through computer simulations. This article provides tools for an analytic treatment of reward-modulated STDP, which allows us to predict under which conditions reward-modulated STDP will achieve a desired learning effect. These analytical results imply that neurons can learn through reward-modulated STDP to classify not only spatial but also temporal firing patterns of presynaptic neurons. They also can learn to respond to specific presynaptic firing patterns with particular spike patterns. Finally, the resulting learning theory predicts that even difficult credit-assignment problems, where it is very hard to tell which synaptic weights should be modified in order to increase the global reward for the system, can be solved in a self-organizing manner through reward-modulated STDP. This yields an explanation for a fundamental experimental result on biofeedback in monkeys by Fetz and Baker. In this experiment monkeys were rewarded for increasing the firing rate of a particular neuron in the cortex and were able to solve this extremely difficult credit assignment problem. Our model for this experiment relies on a combination of reward-modulated STDP with variable spontaneous firing activity. Hence it also provides a possible functional explanation for trial-to-trial variability, which is characteristic for cortical networks of neurons but has no analogue in currently existing artificial computing systems. In addition our model demonstrates that reward-modulated STDP can be applied to all synapses in a large recurrent neural network without endangering the stability of the network dynamics.

Citation: Legenstein R, Pecevski D, Maass W (2008) A Learning Theory for Reward-Modulated Spike-Timing-Dependent Plasticity with Application to Biofeedback. PLoS Comput Biol 4(10): e1000180. doi:10.1371/journal.pcbi.1000180

Editor: Lyle J. Graham, UFR Biomédicale de l'Université René Descartes, France

Received March 3, 2008; **Accepted** August 7, 2008; **Published** October 10, 2008

Funding: Written under partial support by the Austrian Science Fund FWF, project # P17229-N04, project # S9102-N04, as well as project # FP6-015879 (FACETS) and project # FP7-216886 (PASCAL2) of the European Union.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: maass@igi.tugraz.at

☯ These authors contributed equally to this work.

Introduction

Numerous experimental studies (see [1] for a review; [2] discusses more recent in-vivo results) have shown that the efficacy of synapses changes in dependence of the time difference $\Delta t = t_{post} - t_{pre}$ between the firing times t_{pre} and t_{post} of the pre- and postsynaptic neurons. This effect is called spike-timing-dependent plasticity (STDP). But a major puzzle for understanding learning in biological organisms is the relationship between experimentally well-established rules for STDP on the microscopic level, and adaptive changes of the behavior of biological organisms on the macroscopic level. Neuromodulatory systems, which send diffuse signals related to reinforcements (rewards) and behavioral state to several large networks of neurons in the brain, have been identified as likely intermediaries that relate these two levels of plasticity. It is well-known that the consolidation of changes of synaptic weights in response to pre- and postsynaptic neuronal activity requires the presence of such third signals [3,4]. In particular, it has been demonstrated that dopamine (which is behaviorally related to novelty and reward prediction [5]) gates plasticity at corticostriatal synapses [6,7] and within the cortex [8]. It has also been shown that acetylcholine gates synaptic plasticity in the cortex (see for example [9] and [10,11] contains a nice review of the literature).

Corresponding spike-based rules for synaptic plasticity of the form

$$\frac{d}{dt} w_{ji}(t) = c_{ji}(t) d(t) \tag{1}$$

have been proposed in [12] and [13] (see Figure 1 for an illustration of this learning rule), where w_{ji} is the weight of a synapse from neuron i to neuron j, $c_{ji}(t)$ is an eligibility trace of this synapse which collects weight changes proposed by STDP, and $d(t) = h(t) - \bar{h}$ results from a neuromodulatory signal $h(t)$ with mean value $\bar{h}$. It was shown in [12] that a number of interesting learning tasks in large networks of neurons can be accomplished with this simple rule in Equation 1. It has recently been shown that quite similar learning rules for spiking neurons arise when one applies the general framework of distributed reinforcement learning from [14] to networks of spiking neurons [13,15], or if one maximizes the likelihood of postsynaptic firing at desired firing times [16]. However no analytical tools have been available, which make it possible to predict for what learning tasks, and under which parameter settings, reward-modulated STDP will be successful. This article provides such analytical tools, and demonstrates their

Author Summary

A major open problem in computational neuroscience is to explain how learning, i.e., behaviorally relevant modifications in the central nervous system, can be explained on the basis of experimental data on synaptic plasticity. Spike-timing-dependent plasticity (STDP) is a rule for changes in the strength of an individual synapse that is supported by experimental data from a variety of species. However, it is not clear how this synaptic plasticity rule can produce meaningful modifications in networks of neurons. Only if one takes into account that consolidation of synaptic plasticity requires a third signal, such as changes in the concentration of a neuromodulator (that might, for example, be related to rewards or expected rewards), then meaningful changes in the structure of networks of neurons may occur. We provide in this article an analytical foundation for such reward-modulated versions of STDP that predicts when this type of synaptic plasticity can produce functionally relevant changes in networks of neurons. In particular we show that seemingly inexplicable experimental data on biofeedback, where a monkey learnt to increase the firing rate of an arbitrarily chosen neuron in the motor cortex, can be explained on the basis of this new learning theory.

applicability and significance through a variety of computer simulations. In particular, we identify conditions under which neurons can learn through reward-modulated STDP to classify temporal presynaptic firing patterns, and to respond with particular spike patterns.

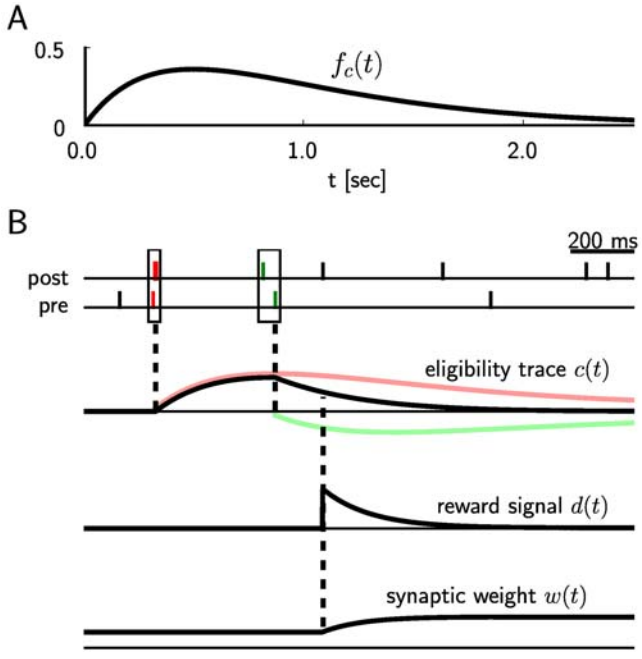

Figure 1. Scheme of reward-modulated STDP according to Equations 1–4. (A) Eligibility function $f_c(t)$, which scales the contribution of a pre/post spike pair (with the second spike at time 0) to the eligibility trace $c(t)$ at time t. (B) Contribution of a pre-before-post spike pair (in red) and a post-before-pre spike pair (in green) to the eligibility trace $c(t)$ (in black), which is the sum of the red and green curves. According to Equation 1 the change of the synaptic weight w is proportional to the product of $c(t)$ with a reward signal $d(t)$.
doi:10.1371/journal.pcbi.1000180.g001

We also provide a model for the remarkable operant conditioning experiments of [17] (see also [18,19]). In the simpler ones of these experiments the spiking activity of single neurons (in area 4 of the precentral gyrus of monkey cortex) was recorded, the deviation of the current firing rate of an arbitrarily selected neuron from its average firing rate was made visible to the monkey through the displacement of an illuminated meter arm, whose rightward position corresponded to the threshold for the feeder discharge. The monkey received food rewards for increasing (or in alternating trials for decreasing) the firing rate of this neuron. The monkeys learnt quite reliably (within a few minutes) to change the firing rate of this neuron in the currently rewarded direction. Adjacent neurons tended to change their firing rate in the same direction, but also differential changes of directions of firing rates of pairs of neurons are reported in [17] (when these differential changes were rewarded). For example, it was shown in Figure 9 of [17] (see also Figure 1 in [19]) that pairs of neurons that were separated by no more than a few hundred microns could be independently trained to increase or decrease their firing rates. Obviously the existence of learning mechanisms in the brain which are able to solve this extremely difficult credit assignment problem provides an important clue for understanding the organization of learning in the brain. We examine in this article analytically under what conditions reward-modulated STDP is able to solve such learning problem. We test the correctness of analytically derived predictions through computer simulations of biologically quite realistic recurrently connected networks of neurons, where an increase of the firing rate of one arbitrarily selected neuron within a network of 4000 neurons is reinforced through rewards (which are sent to all 142813 synapses between excitatory neurons in this recurrent network). We also provide a model for the more complex operant conditioning experiments of [17] by showing that pairs of neurons can be differentially trained through reward-modulated STDP, where one neuron is rewarded for increasing its firing rate, and simultaneously another neuron is rewarded for decreasing its firing rate. More precisely, we increased the reward signal $d(t)$ which is transmitted to all synapses between excitatory neurons in the network whenever the first neuron fired, and decreased this reward signal whenever the second neuron fired (the resulting composed reward corresponds to the displacement of the meter arm that was shown to the monkey in these more complex operant conditioning experiments).

Our theory and computer simulations also show that reward-modulated STDP can be applied to all synapses within a large network of neurons for long time periods, without endangering the stability of the network. In particular this synaptic plasticity rule keeps the network within the asynchronous irregular firing regime, which had been described in [20] as a dynamic regime that resembles spontaneous activity in the cortex. Another interesting aspect of learning with reward-modulated STDP is that it requires spontaneous firing and trial-to-trial variability within the networks of neurons where learning takes place. Hence our learning theory for this synaptic plasticity rule provides a foundation for a functional explanation of these characteristic features of cortical network of neurons that are undesirable from the perspective of most computational theories.

Results

We first give a precise definition of the learning rule in Equation 1 for reward-modulated STDP. The standard rule for STDP, which specifies the change $W(\Delta t)$ of the synaptic weight of an excitatory synapse in dependence on the time difference $\Delta t = t_{post} - t_{pre}$ between the firing times t_{pre} and t_{post} of the pre- and postsynaptic neuron, is

based on numerous experimental data (see [1]). It is commonly modeled by a so-called learning curve of the form

$$W(\Delta t)=\begin{cases} A_+ e^{-\Delta t/\tau_+}, & \text{if } \Delta t \geq 0 \\ -A_- e^{\Delta t/\tau_-}, & \text{if } \Delta t < 0 \end{cases}, \tag{2}$$

where the positive constants A_+ and A_- scale the strength of potentiation and depression respectively, and τ_+ and τ_- are positive time constants defining the width of the positive and negative learning window. The resulting weight change at time t of synapse ji for a presynaptic spike train S_i^{pre} and a postsynaptic spike train S_j^{post} is usually modeled [21] by the instantaneous application of this learning rule to all spike pairings with the second spike at time t

$$\begin{aligned}\left[\frac{d}{dt}w_{ji}(t)\right]_{STDP} &= \int_0^\infty dr\, W(r) S_i^{post}(t) S_i^{pre}(t-r) \\ &+ \int_0^\infty dr\, W(-r) S_j^{post}(t-r) S_i^{pre}(t).\end{aligned} \tag{3}$$

The spike train of a neuron i which fires action potentials at times $t_i^{(1)}, t_i^{(2)}, t_i^{(3)},\ldots$ is formalized here by a sum of Dirac delta functions $S_i(t)=\sum_n \delta\left(t-t_i^{(n)}\right)$.

The model analyzed in this article is based on the assumption that positive and negative weight changes suggested by STDP for all pairs of pre- and postsynaptic spikes at synapse ji (according to the two integrals in Equation 3) are collected in an eligibility trace $c_{ji}(t)$ at the site of the synapse. The contribution to $c_{ij}(t)$ of all spike pairings with the second spike at time $t-s$ is modeled for $s>0$ by a function $f_c(s)$ (see Figure 1A); the time scale of the eligibility trace is assumed in this article to be on the order of seconds. Hence the value of the eligibility trace of synapse ji at time t is given by

$$c_{ji}(t)=\int_0^\infty ds f_c(s)\left[\frac{d}{dt}w_{ji}(t-s)\right]_{STDP}, \tag{4}$$

see Figure 1B. The actual weight change $\frac{d}{dt}w_{ji}(t)$ at time t for reward-modulated STDP is the product $c_{ij}(t)\cdot d(t)$ of the eligibility trace with the reward signal $d(t)$ as defined by Equation 1. Since this simple model can in principle lead to unbounded growth of weights, we assume that weights are clipped at the lower boundary value 0 and an upper boundary w_{max}.

The network dynamics of a simulated recurrent network of spiking neurons where all connections between excitatory neurons are subject to STDP is quite sensitive to the particular STDP-rule that is used. Therefore we have carried out our network simulations not only with the additive STDP-rule in Equation 3, whose effect can be analyzed theoretically, but also with the more complex rule proposed in [22] (which was fitted to experimental data from hippocampal neurons in culture [23]), where the magnitude of the weight change depends on the current value of the weight. An implementation of this STDP-rule (with the parameters proposed in [22]) produced in our network simulations of the biofeedback experiment (computer simulation 1) as well as for learning pattern classification (computer simulation 4) qualitatively the same result as the rule in Equation 3.

Theoretical Analysis of the Resulting Weight Changes

In this section, we derive a learning equation for reward-modulated STDP. This learning equation relates the change of a synaptic weight w_{ji} over some sufficiently long time interval T to statistical properties of the joint distribution of the reward signal $d(t)$ and pre- and postsynaptic firing times, under the assumption that the weight and correlations between pre- and postsynaptic spike times are slowly varying in time. We treat spike times as well as the reward signal $d(t)$ as stochastic variables. This mathematical framework allows us to derive the expected weight change over some time interval T (see [21]), with the expectation taken over realizations of the stochastic input- and output spike trains as well as stochastic realizations of the reward signal, denoted by the ensemble average $\langle\cdot\rangle_E$

$$\begin{aligned}\frac{\langle w_{ji}(t+T)-w_{ji}(t)\rangle_E}{T} &= \frac{1}{T}\Big\langle\int_t^{t+T}\frac{d}{dt}w_{ji}(t')dt'\Big\rangle_E \\ &= \Big\langle\Big\langle\frac{d}{dt}w_{ji}(t)\Big\rangle_T\Big\rangle_E,\end{aligned} \tag{5}$$

where we used the abbreviation $\langle f(t)\rangle_T = T^{-1}\int_t^{t+T} f(t')\,dt'$. If synaptic plasticity is sufficiently slow, synaptic weights integrate a large number of small changes. In this case, the weight w_{ji} can be approximated by its average $\langle w_{ji}\rangle_E$ (it is "self-averaging", see [21]). We can thus drop the expectation on the left hand side of Equation 5 and write it as $\frac{d}{dt}\langle w_{ji}(t)\rangle_T$. Using Equation 1, this yields (see Methods)

$$\begin{aligned}\frac{d}{dt}\langle w_{ji}(t)\rangle_T &= \int_0^\infty dr\, W(r)\int_0^\infty ds\, f_c(s)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T \\ &+ \int_{-\infty}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T.\end{aligned} \tag{6}$$

This formula contains the *reward correlation* for synapse ji

$$\begin{aligned}&D_{ji}(t,s,r)= \\ &\langle d(t)|\text{Neuron } j \text{ spikes at } t-s, \text{ and neuron } i \text{ spikes at } t-s-r\rangle_E,\end{aligned} \tag{7}$$

which is the average reward at time t given a presynaptic spike at time $t-s-r$ and a postsynaptic spike at time $t-s$. The joint firing rate $v_{ji}(t,r)=\langle S_j(t)S_i(t-r)\rangle_E$ describes correlations between spike timings of neurons j and i, i.e., it is the probability density for the event that neuron i fires an action potential at time $t-r$ and neuron j fires an action potential at time t. For synapses subject to reward-modulated STDP, changes in efficacy are obviously driven by co-occurrences of spike pairings and rewards within the time scale of the eligibility trace. Equation 6 clarifies how the expected weight change depends on how the correlations between the pre- and postsynaptic neurons correlate with the reward signal.

If one assumes for simplicity that the impact of a spike pair on the eligibility trace is always triggered by the postsynaptic spike, one gets a simpler equation (see Methods)

$$\frac{d}{dt}\langle w_{ji}(t)\rangle_T = \int_0^\infty ds\, f_c(s)\int_{-\infty}^\infty dr\, W(r)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T. \tag{8}$$

The assumption introduces a small error for post-before-pre spike pairs, because for a reward signal that arrives at some time d_r after the pairing, the weight update will be proportional to $f_c(d_r)$ instead of $f_c(d_r+r)$. The approximation is justified if the temporal average is performed on a much longer time scale than the time scale of the learning window, the effect of each pre-post spike pair on the reward signal is delayed by an amount greater than the time scale of the learning window, and f_c changes slowly compared to the time scale of the learning window (see Methods for details). For the

analyzes presented in this article, the simplified Equation 8 is a good approximation for the learning dynamics. Equation 8 is a generalized version of the STDP learning equation $\frac{d}{dt}w_{ji}(t)=\int_{-\infty}^{\infty} dr\, W(r)\langle v_{ji}(t-s,r)\rangle_T$ in [21] that includes the impact of the reward correlation weighted by the eligibility function. To see the relation between standard STDP and reward-modulated STDP, consider a constant reward signal $d(t)=d_0$. Then also the reward correlation is constant and given by $D(t,s,r)=d_0$. We recover the standard STDP learning equation scaled by d_0 if the eligibility function is an instantaneous delta-pulse $f_c(s)=\delta(s)$. Furthermore, if the statistics of the reward signal $d(t)$ is time-independent and independent from the pre- and postsynaptic spike statistics of some synapse ji, then the reward correlation is given by $D_{ji}(t,s,r)=\langle d(t)\rangle_E=d_0$ for some constant d_0. Then, the weight change for synapse ji is $\frac{d}{dt}\langle w_{ji}(t)\rangle_T=d_0\int_{-\infty}^{\infty} dr\, W(r)\int_0^{\infty} ds f_c(s)\langle v_{ji}(t-s,r)\rangle_T$. The temporal average of the joint firing rate $\langle v_{ji}(t-s,r)\rangle_T$ is thus filtered by the eligibility trace. We assumed in the preceding analysis that the temporal average is taken over some long time interval T. If the time scale of the eligibility trace is much smaller than this time interval T, then the weight change is approximately $\frac{d}{dt}\langle w_{ji}(t)\rangle_T\approx d_0\left(\int_0^{\infty} ds f_c(s)\right)\int_{-\infty}^{\infty} dr\, W(r)\langle v_{ji}(t,r)\rangle_T$, and the weight w_{ji} will change according to standard STDP scaled by a constant proportional to the mean reward and the integral over the eligibility function. In the remainder of this article, we will always use the smooth time-averaged weight change $\frac{d}{dt}\langle w_{ji}(t)\rangle_T$, but for brevity, we will drop the angular brackets and simply write $\frac{d}{dt}w_{ji}(t)$.

The learning Equation 8 provides the mathematical basis for our following analyses. It allows us to determine synaptic weight changes if we can describe a learning situation in terms of reward correlations and correlations between pre- and postsynaptic spikes.

Application to Models for Biofeedback Experiments

We now apply the preceding analysis to the biofeedback experiment of [17] that were described in the introduction. These experiments pose the challenge to explain how learning mechanisms in the brain can detect and exploit correlations between rewards and the firing activity of one or a few neurons within a large recurrent network of neurons (the credit assignment problem), without changing the overall function or dynamics of the circuit.

We show that this phenomenon can in principle be explained by reward-modulated STDP. In order to do that, we define a model for the experiment which allows us to formulate an equation for the reward signal $d(t)$. This enables us to calculate synaptic weight changes for this particular scenario. We consider as model a recurrent neural circuit where the spiking activity of one neuron k is recorded by the experimenter (Experiments where two neurons are recorded and reinforced were also reported in [17]. We tested this case in computer simulations (see Figure 2) but did not treat it explicitly in our theoretical analysis). We assume that in the monkey brain a reward signal $d(t)$ is produced which depends on the visual feedback (through an illuminated meter, whose pointer deflection was dependent on the current firing rate of the randomly selected neuron k) as well as previously received liquid rewards, and that this signal $d(t)$ is delivered to *all* synapses in large areas of the brain. We can formalize this scenario by defining a reward signal which depends on the spike rate of the arbitrarily selected neuron k (see Figure 3A and 3B). More precisely, a reward pulse of shape $\varepsilon_r(r)$ (the reward kernel) is produced with some delay d_r every time the neuron k produces an action potential

$$d(t)=\int_0^{\infty} dr\, S_k^{post}(t-d_r-r)\varepsilon_r(r). \tag{9}$$

Note that $d(t)=h(t)-\bar{h}$ is defined in Equation 1 as a signal with zero mean. In order to satisfy this constraint, we assume that the reward kernel ε_r has zero mass, i.e., $\bar{\varepsilon}_r=\int_0^{\infty} dr\, \varepsilon_r(r)=0$. For the analysis, we use the linear Poisson neuron model described in Methods. The mean weight change for synapses to the reinforced neuron k is then approximately (see Methods)

$$\frac{d}{dt}w_{ki}(t)\approx \int_0^{\infty} ds f_c(s+d_r)\varepsilon_r(s)\int_{-\infty}^{\infty} dr\, W(r)\langle v_{ki}(t-d_r-s,r)\rangle_T. \tag{10}$$

This equation describes STDP with a learning rate proportional to $\int_0^{\infty} ds f_c(s+d_r)\varepsilon_r(s)$. The outcome of the learning session will strongly depend on this integral and thus on the form of the reward kernel ε_r. In order to reinforce high firing rates of the reinforced neuron we have chosen a reward kernel with a positive bump in the first few hundred milliseconds, and a long negative tail afterwards. Figure 3C shows the functions f_c and ε_r that were used in our computer model, as well as the product of these two functions. One sees that the integral over the product is positive and according to Equation 10 the synapses to the reinforced neuron are subject to STDP. This does not guarantee an increase of the firing rate of the reinforced neuron. Instead, the changes of neuronal firing will depend on the statistics of the inputs. In particular, the weights of synapses to neuron k will not increase if that neuron does not fire spontaneously. For uncorrelated Poisson input spike trains of equal rate, the firing rate of a neuron trained by STDP stabilizes at some value which depends on the input rate (see [24,25]). However, in comparison to the low spontaneous firing rates observed in the biofeedback experiment [17], the stable firing rate under STDP can be much higher, allowing for a significant rate increase. It was shown in [17] that also low firing rates of a single neuron can be reinforced. In order to model this, we have chosen a reward kernel with a negative bump in the first few hundred milliseconds, and a long positive tail afterwards, i.e. we inverted the kernel used above to obtain a negative integral $\int_0^{\infty} ds f_c(s+d_r)\varepsilon_r(s)$. According to Equation 10 this leads to anti-STDP where not only inputs to the reinforced neuron which have low correlations with the output are depressed (because of the negative integral of the learning window), but also those which are causally correlated with the output. This leads to a quick firing rate decrease at the reinforced neuron.

The mean weight change of synapses to non-reinforced neurons $j\neq k$ is given by

$$\frac{d}{dt}w_{ji}(t)\approx \int_0^{\infty} ds f_c(s)\int_{-\infty}^{\infty} dr\, W(r)\int_0^{\infty} dr'\varepsilon_r(r') \left\langle \frac{v_{kj}(t-d_r-r',s-d_r-r')}{v_j(t-s)}v_{ji}(t-s,r)\right\rangle_T, \tag{11}$$

where $v_j(t)=\langle S_j(t)\rangle_E$ is the instantaneous firing rate of neuron j at time t. This equation indicates that a non-reinforced neuron is trained by STDP with a learning rate proportional to its correlation with the reinforced neuron given by $v_{kj}(t-d_r-r',s-d_r-r')/v_j(t-s)$. In fact, it was noted in [17] that neurons nearby the reinforced neuron tended to change their firing rate in the same direction. This observation might be explained by putative correlations of the recorded neuron

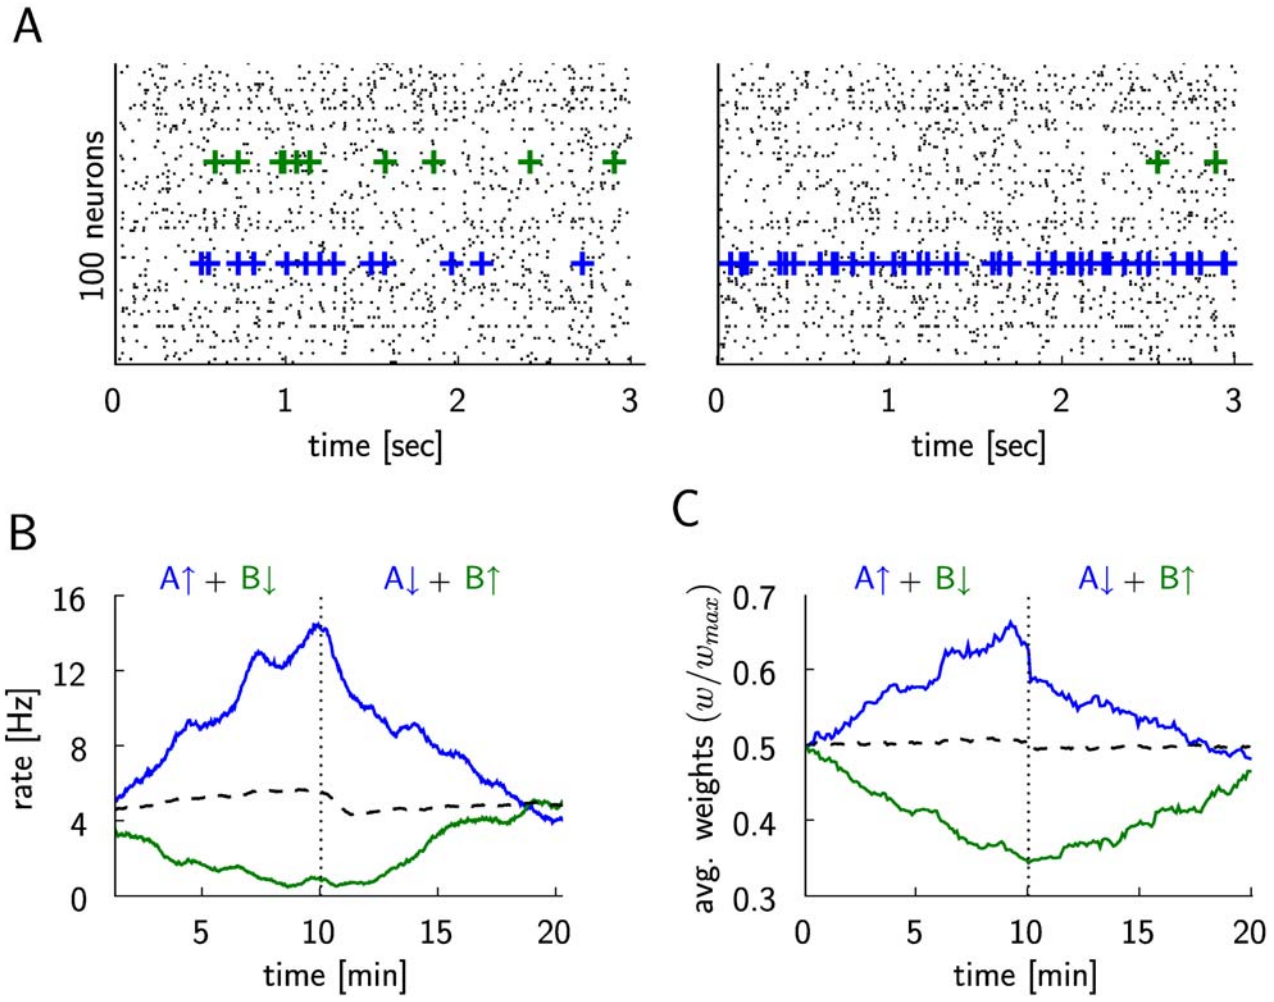

Figure 2. Differential reinforcement of two neurons (within a simulated network of 4000 neurons, the two rewarded neurons are denoted as A and B), corresponding to the experimental results shown in Figure 9 of [17] and Figure 1 of [19]. (A) The spike response of 100 randomly chosen neurons at the beginning of the simulation (20 sec–23 sec, left plot), and at the middle of simulation just before the switching of the reward policy (597 sec–600 sec, right plot). The firing times of the first reinforced neuron A are marked by blue crosses and those of the second reinforced neuron B are marked by green crosses. (B) The dashed vertical line marks the switch of the reinforcements at $t = 10$ min. The firing rate of neuron A (blue line) increases while it is positively reinforced in the first half of the simulation and decreases in the second half when its spiking is negatively reinforced. The firing rate of the neuron B (green line) decreases during the negative reinforcement in the first half and increases during the positive reinforcement in the second half of the simulation. The average firing rate of 20 other randomly chosen neurons (dashed line) remains unchanged. (C) Evolution of the average weight of excitatory synapses to the rewarded neurons A and B (blue and green lines, respectively), and of the average weight of 1744 randomly chosen excitatory synapses to other neurons in the circuit (dashed line).
doi:10.1371/journal.pcbi.1000180.g002

with nearby neurons. On the other hand, if a neuron j is uncorrelated with the reinforced neuron k, we can decompose the joint firing rate into $v_{kj}(t-d_r-r',s-d_r-r') = v_k(t-d_r-r')v_j(t-s)$. In this case, the learning rate for synapse ji is approximately zero (see Methods). This ensures that most neurons in the circuit keep a constant firing rate, in spite of continuous weight changes according to reward-modulated STDP.

Altogether we see that the weights of synapses to the reinforced neuron k can only change if there is spontaneous activity in the network, so that in particular also this neuron k fires spontaneously. On the other hand the spontaneous network activity should not consist of repeating large-scale spatio-temporal firing patterns, since that would entail correlations between the firing of neuron k and other neurons j, and would lead to similar changes of synapses to these other neurons j. Apart from these requirements on the spontaneous network activity, the preceding theoretical results predict that stability of the circuit is preserved, while the neuron which is causally related to the reward signal is trained by STDP, if $\int_0^\infty ds\, f_c(s+d_r)\varepsilon_r(s)$ is positive.

Computer Simulation 1: Model for Biofeedback Experiment

We tested these theoretical predictions through computer simulations of a generic cortical microcircuit receiving a reward signal which depends on the firing of one arbitrarily chosen neuron k from the circuit (reinforced neuron). The circuit was composed of 4000 LIF neurons, with 3200 being excitatory and 800 inhibitory, interconnected randomly by 228954 conductance based synapses with short term dynamics (All computer simulations were also carried out as a control with static current based synapses, see Methods and Suppl.). In addition to the explicitly modeled synaptic connections, conductance noise (generated by

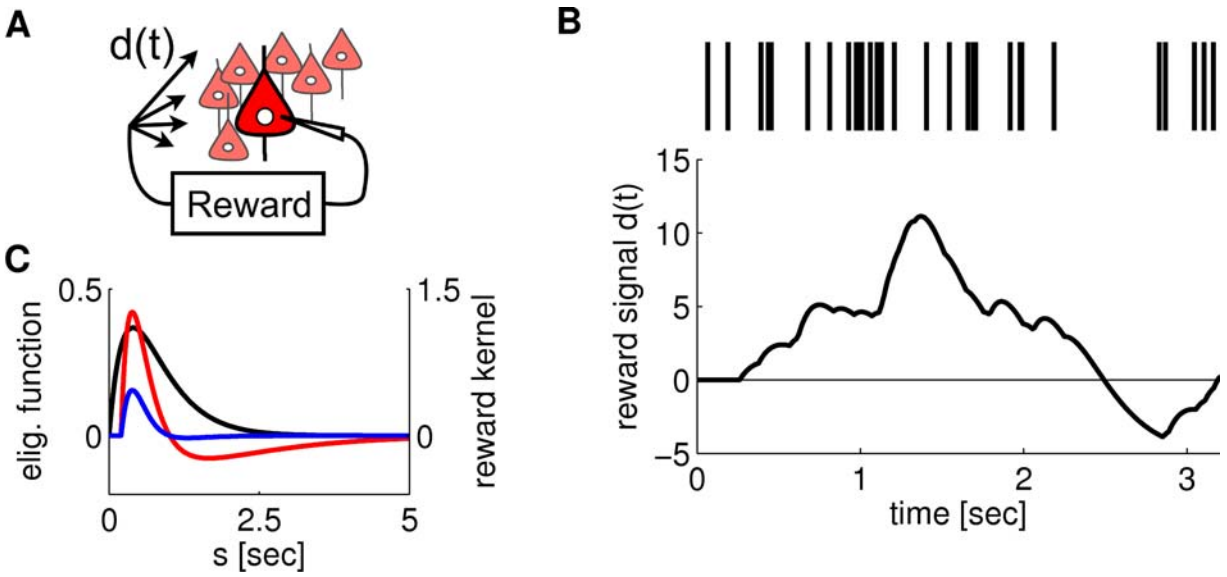

Figure 3. Setup of the model for the experiment by Fetz and Baker [17]. (A) Schema of the model: The activity of a single neuron in the circuit determines the amount of reward delivered to all synapses between excitatory neurons in the circuit. (B) The reward signal $d(t)$ in response to a spike train (shown at the top) of the arbitrarily selected neuron (which was selected from a recurrently connected circuit consisting of 4000 neurons). The level of the reward signal $d(t)$ follows the firing rate of the spike train. (C) The eligibility function $f_c(s)$ (black curve, left axis), the reward kernel $\varepsilon_r(s)$ delayed by 200 ms (red curve, right axis), and the product of these two functions (blue curve, right axis) as used in our computer experiment. The integral of $f_c(s+d_r)\varepsilon_r(s)$ is positive, as required according to Equation 10 in order to achieve a positive learning rate for the synapses to the selected neuron.
doi:10.1371/journal.pcbi.1000180.g003

an Ornstein-Uhlenbeck process) was injected into each neuron according to data from [26], in order to model synaptic background activity of neocortical neurons in-vivo (More precisely, for 50% of the excitatory neurons the amplitude of the noise injection was reduced to 20%, and instead their connection probabilities from other excitatory neurons were chosen to be larger, see Methods and Figure S1 and Figure S2 for details. The reinforced neuron had to be chosen from the latter population, since reward-modulated STDP does not work properly if the postsynaptic neuron fires too often because of directly injected noise). This background noise elicited spontaneous firing in the circuit at about 4.6 Hz. Reward-modulated STDP was applied continuously to all synapses which had excitatory presynaptic and postsynaptic neurons, and all these synapses received the same reward signal. The reward signal was modeled according to Equation 9. Figure 3C shows one reward pulse caused by a single postsynaptic spike at time $t=0$ with the parameters used in the experiment. For several postsynaptic spikes, the amplitude of the reward signal follows the firing rate of the reinforced neuron, see Figure 3B.

This model was simulated for 20 minutes of biological time. Figure 4A, 4B, and 4D show that the firing rate of the reinforced neuron increases within a few minutes (like in the experiment of [17]), while the firing rates of the other neurons remain largely unchanged. The increase of weights to the reinforced neuron shown in Figure 4C can be explained by the correlations between its presynaptic and postsynaptic spikes shown in panel E. This panel shows that pre-before-post spike pairings (black curve) are in general more frequent than post-before-pre spike pairings. The reinforced neuron increases its rate from around 4 Hz to 12 Hz, which is comparable to the measured firing rates in [15] before and after learning.

In Figure 9 of [17] and Figure 1 of [19] the results of another experiment were reported where the activity of two adjacent neurons was recorded, and high firing rates of the first neuron and low firing rates of the second neuron were reinforced simultaneously. This kind of differential reinforcement resulted in an increase and decrease of the firing rates of the two neurons correspondingly. We implemented this type of reinforcement by letting the reward signal in our model depend on the spikes of the two randomly chosen neurons (we refer to these neurons as neuron A and neuron B), i.e. $d(t)=d_+^A(t)+d_-^B(t)$, where $d_+^A(t)$ is the component that positively rewards spikes of neuron A, and $d_-^B(t)$ negatively rewards spikes of neuron B. Both parts of the reward signal, $d_+^A(t)$ and $d_-^B(t)$, were defined as in Equation 9 for the corresponding neuron. For $d_+^A(t)$ we used the reward kernel ε_r as defined in Equation 29, whereas for $d_-^B(t)$ we used $\varepsilon_{r-}=-\varepsilon_r$ (note that the integral over ε_{r-} is still zero). At the middle of the simulation (simulation time $t=10$ min), we changed the direction of the reinforcements by negatively rewarding the firing of neuron A and positively rewarding the firing of neuron B (i.e., $d(t)=d_-^A(t)+d_+^B(t)$). The results are summarized in Figure 2. With a reward signal modeled in this way, we were able to independently increase and decrease the firing rates of the two neurons according to the reinforcements, while the firing rates of the other neurons remained unchanged. Changing the type of reinforcement during the simulation from positive to negative for neuron A and from negative to positive for neuron B resulted in a corresponding shift in their firing rate change in the direction of the reinforcement.

The dynamics of a network where STDP is applied to all synapses between excitatory neurons is quite sensitive to the specific choice of the STDP-rule. The preceding theoretical analysis (see Equations 10 and 11) predicts that reward-modulated STDP affects in the long run only those excitatory synapses where the firing of the postsynaptic neuron is correlated with the reward signal. In other words: the reward signal gates the effect of STDP in a recurrent network, and thereby can keep the network within a

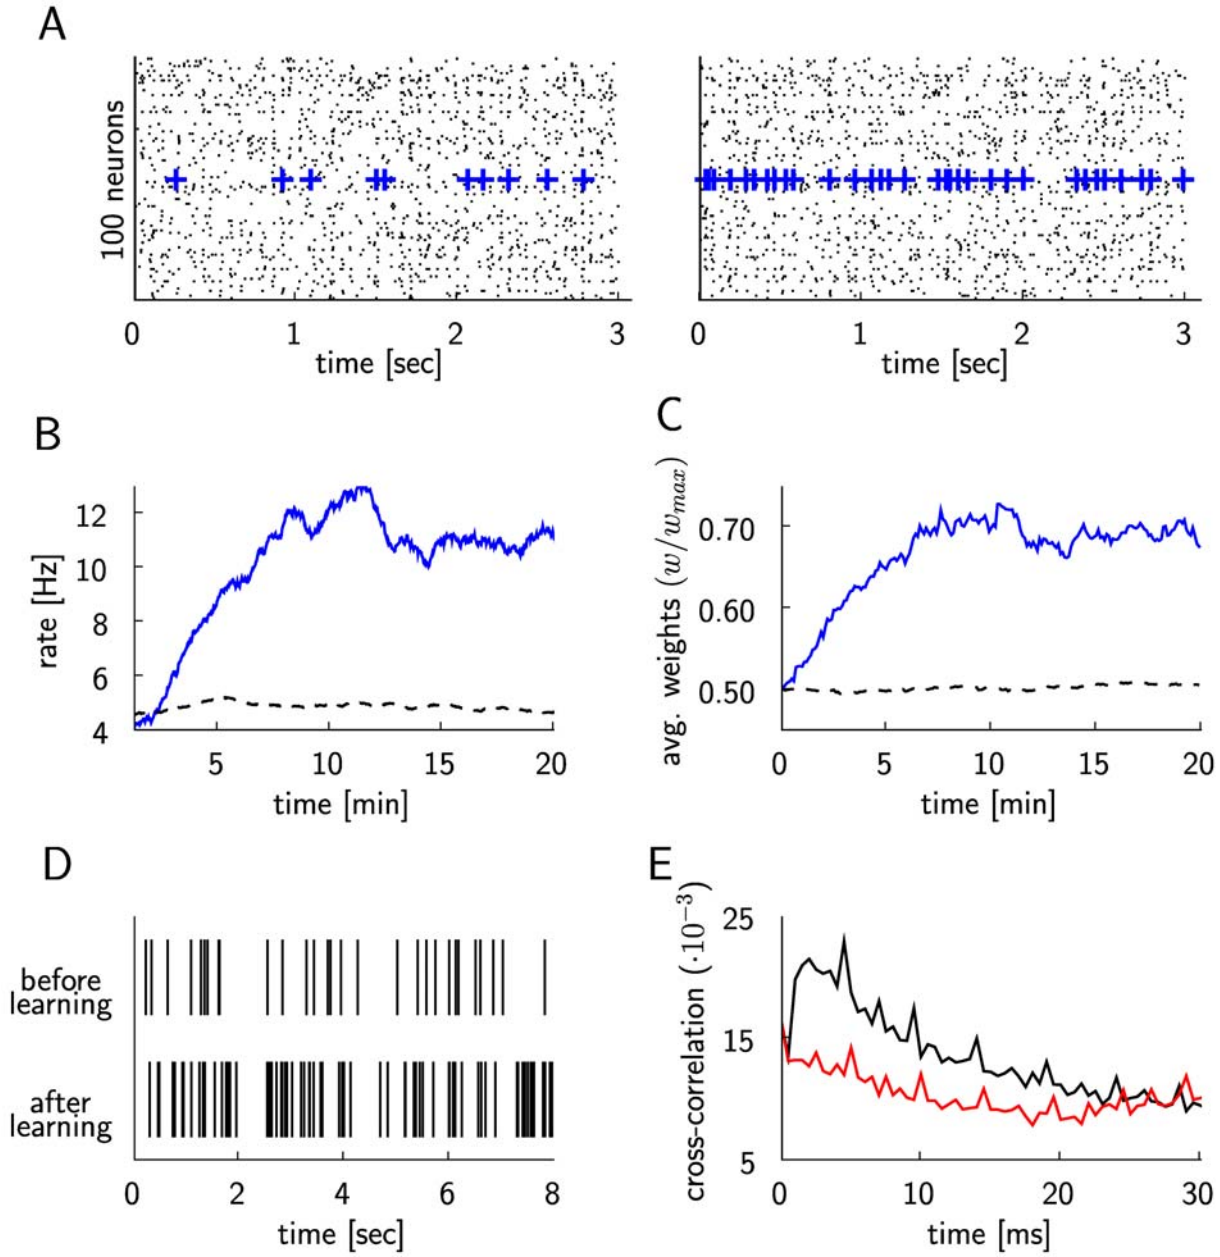

Figure 4. Simulation of the experiment by Fetz and Baker [17] for the case where an arbitrarily selected neuron triggers global rewards when it increases its firing rate. (A) Spike response of 100 randomly chosen neurons within the recurrent network of 4000 neurons at the beginning of the simulation (20 sec–23 sec, left plot), and at the end of the simulation (the last 3 seconds, right plot). The firing times of the reinforced neuron are marked by blue crosses. (B) The firing rate of the positively rewarded neuron (blue line) increases, while the average firing rate of 20 other randomly chosen neurons (dashed line) remains unchanged. (C) Evolution of the average weight of excitatory synapses to the reinforced neuron (blue line), and of the average weight of 1663 randomly chosen excitatory synapses to other neurons in the circuit (dashed line). (D) Spike trains of the reinforced neuron before and after learning. (E) Histogram of the time-differences between presynaptic and postsynaptic spikes (bin size 0.5 ms), averaged over all excitatory synapses to the reinforced neuron. The black curve represents the histogram values for positive time differences (when the presynaptic spike precedes the postsynaptic spike), and the red curve represents the histogram for negative time differences.
doi:10.1371/journal.pcbi.1000180.g004

given dynamic regime. This prediction is confirmed qualitatively by the two panels of Figure 4A, which show that even after all excitatory synapses in the recurrent network have been subject to 20 minutes (in simulated biological time) of reward-modulated STDP, the network stays within the asynchronous irregular firing regime. It is also confirmed quantitatively through Figure 5. These figures show results for the simple additive version of STDP (according to Equation 3). Very similar results (see Figure S3 and Figure S4) arise from an application of the more complex STDP-rule proposed in [22] where the weight-change depends on the current weight value.

Rewarding Spike-Times

The preceding model for the biofeedback experiment of Fetz and Baker focused on learning of firing rates. In order to explore the capabilities and limitations of reward-modulated STDP in

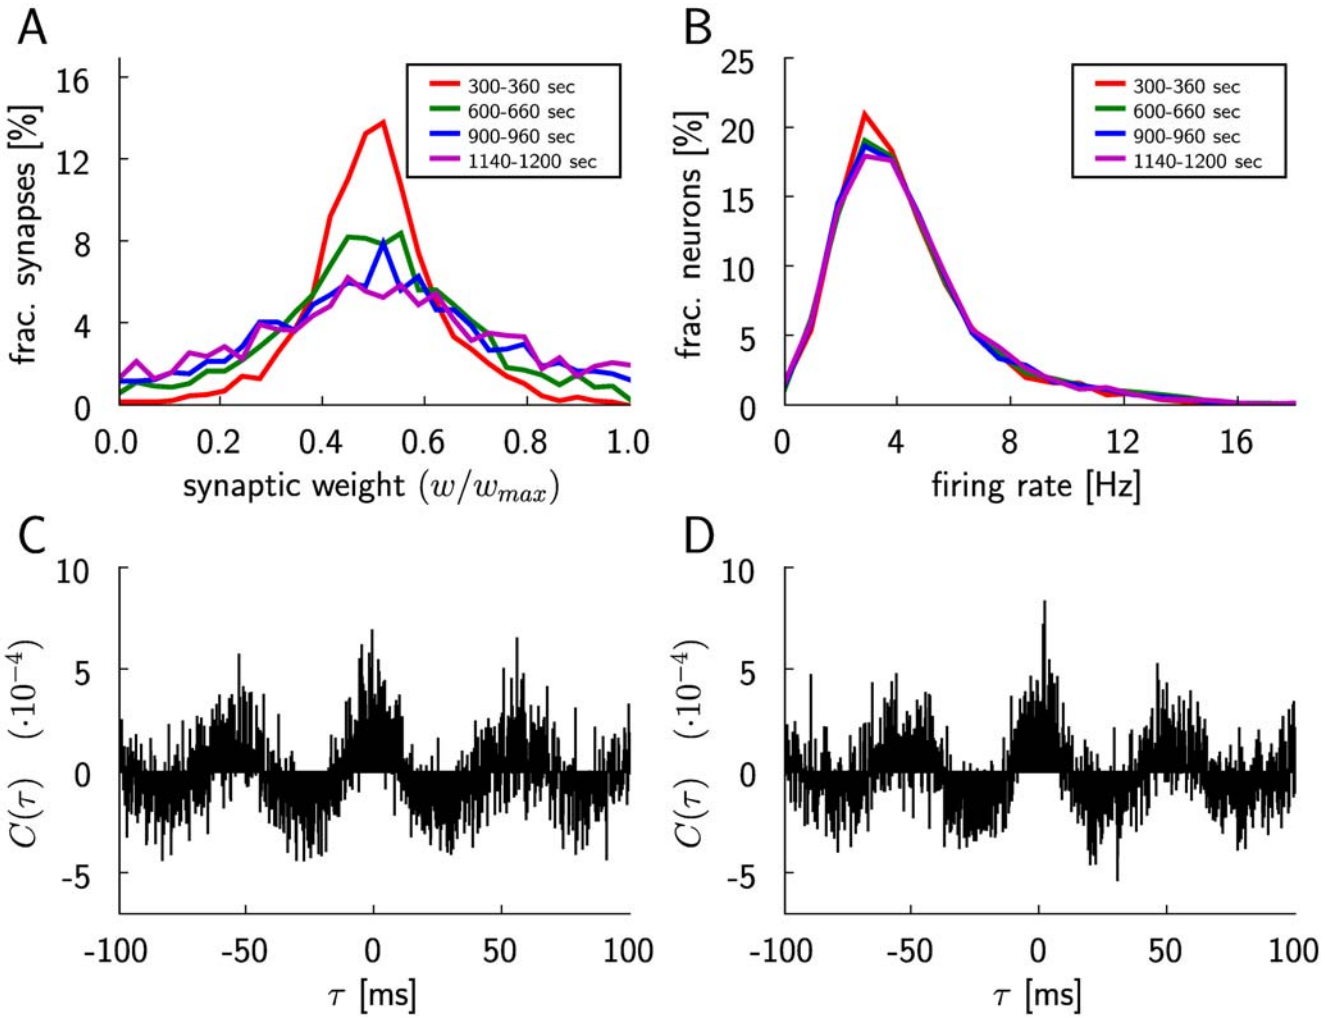

Figure 5. Evolution of the dynamics of a recurrent network of 4000 LIF neurons during application of reward-modulated STDP. (A) Distribution of the synaptic weights of excitatory synapses to 50 randomly chosen non-reinforced neurons, plotted for 4 different periods of simulated biological time during the simulation. The weights are averaged over 10 samples within these periods. The colors of the curves and the corresponding intervals are as follows: red (300–360 sec), green (600–660 sec), blue (900–960 sec), magenta (1140–1200 sec). (B) The distribution of average firing rates of the non-reinforced excitatory neurons in the circuit, plotted for the same time periods as in (A). The colors of the curves are the same as in (A). The distribution of the firing rates of the neurons in the circuit remains unchanged during the simulation, which covers 20 minutes of biological time. (C) Cross-correlogram of the spiking activity in the circuit, averaged over 200 pairs of non-reinforced neurons and over 60 s, with a bin size of 0.2 ms, for the period between 300 and 360 seconds of simulated biological time. It is calculated as the cross-covariance divided by the square root of the product of variances. (D) As in (C), but between seconds 1140 and 1200. (Separate plots of (B), (C), and (D) for two types of excitatory neurons that received different amounts of noise currents are given in Figure S1 and Figure S2.)
doi:10.1371/journal.pcbi.1000180.g005

contexts where the temporal structure of spike trains matters, we investigated another reinforcement learning scenario where a neuron should learn to respond with particular temporal spike patterns. We first apply analytical methods to derive conditions under which a neuron subject to reward-modulated STDP can achieve this.

In this model, the reward signal $d(t)$ is given in dependence on how well the output spike train S_j^{post} of a neuron j matches some rather arbitrary spike train S^* (which might for example represent spike output from some other brain structure during a developmental phase). S^* is produced by a neuron μ^* that receives the same n input spike trains $S_1,\ldots,S_n$ as the trained neuron j, with some arbitrarily chosen weights $\mathbf{w}^* = (w_1^*,\ldots,w_n^*)^T$, $w_i^* \in \{0, w_{max}\}$. But in addition the neuron μ^* receives $n'-n$ further spike trains $S_{n+1},\ldots,S_{n'}$ with weights $w_{n+1}^*,\ldots,w_{n'}^* = w_{max}$. The setup is illustrated in Figure 6A. It provides a generic reinforcement learning scenario, when a quite arbitrary (and not perfectly realizable) spike output is reinforced, but simultaneously the performance of the learner can be evaluated clearly according to how well its weights $w_{j1},\ldots,w_{jn}$ match those of the neuron μ^* for those n input spike trains which both of them have in common. The reward $d(t)$ at time t depends in this task on both the timing of action potentials of the trained neuron and spike times in the target spike train S^*

$$d(t) = \int_{-\infty}^{\infty} dr\ \kappa(r) S_j^{post}(t-d_r) S^*(t-d_r-r), \qquad (12)$$

where the function $\kappa(r)$ with $\overline{\kappa} = \int_{-\infty}^{\infty} ds\ \kappa(s) > 0$ describes how the reward signal depends on the time difference r between a

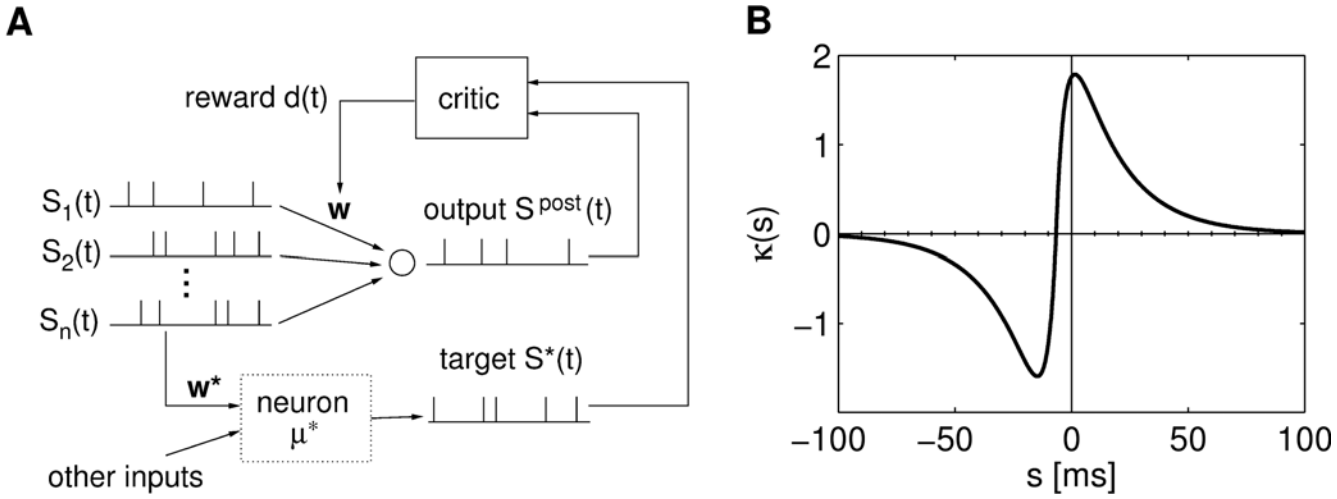

Figure 6. Setup for reinforcement learning of spike times. (A) Architecture. The trained neuron receives *n* input spike trains. The neuron μ^* receives the same inputs plus additional inputs not accessible to the trained neuron. The reward is determined by the timing differences between the action potentials of the trained neuron and the neuron μ^*. (B) A reward kernel with optimal offset from the origin of $t_\kappa = -6.6$ ms. The optimal offset for this kernel was calculated with respect to the parameters from computer simulation 1 in Table 1. Reward is positive if the neuron spikes around the target spike or somewhat later, and negative if the neuron spikes much too early.
doi:10.1371/journal.pcbi.1000180.g006

postsynaptic spike and a target spike, and $d_r > 0$ is the delay of the reward.

Our theoretical analysis (see Methods) predicts that under the assumption of constant-rate uncorrelated Poisson input statistics this reinforcement learning task can be solved by reward-modulated STDP for arbitrary initial weights if three constraints are fulfilled:

$$-\nu_{min}^{post}\overline{W} > w_{max}\overline{W}_\varepsilon \tag{13}$$

$$\int_{-\infty}^{\infty} dr\, W(r)\varepsilon(r)\varepsilon_\kappa(r) \geq -\nu_{max}^{post}\overline{W}\int_0^{\infty} dr\, \varepsilon(r)\varepsilon_\kappa(r) \tag{14}$$

$$\int_{-\infty}^{\infty} dr\, W(r)\varepsilon_\kappa(r) > -\overline{W}\overline{\kappa}\left[\frac{\nu^*\nu_{max}^{post}}{w_{max}}\frac{\bar{f}_c}{f_c(d_r)} + \frac{\nu^*}{w_{max}} + \nu^* + \nu_{max}^{post}\right] \tag{15}$$

The following parameters occur in these equations: ν^* is the output rate of neuron μ^*, ν_{min}^{post} is the minimal output rate, ν_{max}^{post} is the maximal output rate of the trained neuron, $\overline{f}_c = \int_0^\infty dr\, f_c(r)$ is the integral over the eligibility trace, $\overline{W} = \int_{-\infty}^{\infty} dr\, W(r)$ is the integral over the STDP learning curve (see Equation 2), $\varepsilon_\kappa(r) = \int_{-\infty}^{\infty} dr'\, \kappa(r')\varepsilon(r-r')$ is the convolution of the reward kernel with the shape of the postsynaptic potential (PSP) $\varepsilon(s)$, *and* $\overline{W}_\varepsilon = \int_{-\infty}^{\infty} dr\, \varepsilon(r)W(r)$ *is the integral over the PSP weighted by the learning window.*

If these inequalities are fulfilled and input rates are larger than zero, then the weight vector of the trained neuron converges on average from any initial weight vector to $\mathbf{w}^*$ (i.e., it mimics the weight distribution of neuron μ^* for those *n* inputs which both have in common). To get an intuitive understanding of these inequalities, we first examine the idea behind Constraint 13. This constraint assures that weights of synapses *i* with $w_i^* = 0$ decay to zero in expectation. First note that input spikes from a spike train S_i with $w_i^* = 0$ have no influence on the target spike train S^*. In the linear Poisson neuron model, this leads to weight changes similar to STDP which can be described by two terms. First, all synapses are subject to depression stemming from the negative part of the learning curve W and random pre-post spike pairs. This weight change is bounded from below by $\alpha\nu_i^{pre}\nu_{min}^{post}\overline{W}$ for some positive constant α. On the other hand, the positive influence of input spikes on postsynaptic firing leads to potentiation of the synapse bounded from above by $\alpha\nu_i^{pre}w_{max}\overline{W}_\varepsilon$. Hence the weight decays to zero if $-\alpha\nu_i^{pre}\nu_{min}^{post}\overline{W} > \alpha\nu_i^{pre}w_{max}\overline{W}_\varepsilon$, leading to Inequality 13. For synapses *i* with $w_i = w_{max}$, there is an additional drive, since each presynaptic spike increases the probability of a closely following spike in the target spike train S^*. Therefore, the probability of a delayed reward signal after a presynaptic spike is larger. This additional drive leads to positive weight changes if Inequalities 14 and 15 are fulfilled (see Methods).

Note that also for the learning of spike times spontaneous spikes (which might be regarded as "noise") are important, since they may lead to reward signals that can be exploited by the learning rule. It is obvious that in reward-modulated STDP, a silent neuron cannot recover from its silent state, since there will be no spikes which can drive STDP. But in addition, Condition 13 shows that in this learning scenario, the minimal output rate ν_{min}^{post}—which increases with increasing noise—has to be larger than some positive constant, such that depression is strong enough to weaken synapses if needed. On the other hand, if the noise is too strong also synapses *i* with $w_i = w_{max}$ will be depressed and may not converge correctly. This can happen when the increased noise leads to a maximal postsynaptic rate ν_{max}^{post} such that Constraints 14 and 15 are not satisfied anymore.

Conditions 13–15 also reveal how parameters of the model influence the applicability of this setup. For example, the eligibility trace enters the equations only in the form of its integral and its value at the reward delay in Equation 15. In fact, the exact shape of the eligibility trace is not important. The important property of an ideal eligibility trace is that it is high at the reward delay and low at other times as expressed by the fraction in Condition 15. Interestingly, the formulas also show that one has quite some freedom in choosing the form of the STDP window, as long as the reward kernel ε_κ is adjusted accordingly. For example, instead of a standard STDP learning window W with $W(r) \geq 0$ for $r > 0$ and

$W(r) \leq 0$ for $r<0$ and a corresponding reward kernel κ, one can use a reversed learning window W' defined by $W'(r) \equiv W(-r)$ and a reward kernel κ' such that $\varepsilon_{\kappa'}(r) = \varepsilon_{\kappa}(-r)$. If Condition 15 is satisfied for W and κ, then it is also satisfied for W' and κ' (and in most cases also Condition 14 will be satisfied). This reflects the fact that in reward modulated STDP the learning window defines the weight changes in combination with the reward signal.

For a given STDP learning window, the analysis reveals what reward kernels κ are suitable for this learning setup. From Condition 15, we can deduce that the integral over κ should be small (but positive), whereas the integral $\int_{-\infty}^{\infty} dr\, W(r)\varepsilon_{\kappa}(r)$ should be large. Hence, for a standard STDP learning window W with $W(r) \geq 0$ for $r>0$ and $W(r) \leq 0$ for $r<0$, the convolution $\varepsilon_{\kappa}(r)$ of the reward kernel with the PSP should be positive for $r>0$ and negative for $r<0$. In the computer simulation we used a simple kernel depicted in Figure 6B, which satisfies the aforementioned constraints. It consists of two double-exponential functions, one positive and one negative, with a zero crossing at some offset t_{κ} from the origin. The optimal offset t_{κ} is always negative and in the order of several milliseconds for usual PSP-shapes ε. We conclude that for successful learning in this scenario, a positive reward should be produced if the neuron spikes around the target spike or somewhat later, and a negative reward should be produced if the neuron spikes much too early.

Computer Simulation 2: Learning Spike Times

In order to explore this learning scenario in a biologically more realistic setting, we trained a LIF neuron with conductance based synapses exhibiting short term facilitation and depression. The trained neuron and the neuron μ^* which produced the target spike train S^* both received inputs from 100 input neurons emitting spikes from a constant rate Poisson process of 15 Hz. The synapses to the trained neuron were subject to reward-modulated STDP. The weights of neuron μ^* were set to $w_i^* = w_{max}$ for $0 \leq i < 50$ and $w_i^* = 0$ for $50 \leq i < 100$. In order to simulate a non-realizable target response, neuron μ^* received 10 additional synaptic inputs (with weights set to $w_{max}/2$). During the simulations we observed a firing rate of 18.2 Hz for the trained neuron, and 25.2 Hz for the neuron μ^*. The simulations were run for 2 hours simulated biological time.

We performed 5 repetitions of the experiment, each time with different randomly generated inputs and different initial weight values for the trained neuron. In each of the 5 runs, the average synaptic weights of synapses with $w_i^* = w_{max}$ and $w_i^* = 0$ approached their target values, as shown in Figure 7A. In order to test how closely the trained neuron reproduces the target spike train S^* after learning, we performed additional simulations where the same spike input was applied to the trained neuron before and after the learning. Then we compared the output of the trained neuron before and after learning with the output S^* of neuron μ^*. Figure 7B shows that the trained neuron approximates the part of S^* which is accessible to it quite well. Figure 7C–F provide more detailed analyses of the evolution of weights during learning. The computer simulations confirmed the theoretical prediction that the neuron can learn well through reward-modulated STDP only if a certain level of noise is injected into the neuron (see preceding discussion and Figure S6).

Both the theoretical results and these computer simulations demonstrate that a neuron can learn quite well through reward-modulated STDP to respond with specific spike patterns.

Computer Simulation 3: Testing the Analytically Derived Conditions

Equations 13–15 predict under which relationships between the parameters involved the learning of particular spike responses through reward-modulated STDP will be successful. We have tested these predictions by selecting 6 arbitrary settings of these parameters, which are listed in Table 1. In 4 cases (marked by light gray shading in Figure 8) these conditions were not met (either for the learning of weights with target value w_{max}, or for the learning of weights with target value 0. Figure 8 shows that the derived learning result is not achieved in exactly these 4 cases. On the other hand, the theoretically predicted weight changes (black bar) predict in all cases the actual weight changes (gray bar) that occur for the chosen simulation times (listed in the last column of Table 1) remarkably well.

Pattern Discrimination with Reward-Modulated STDP

We examine here the question whether a neuron can learn through reward-modulated STDP to discriminate between two spike patterns P and N of its presynaptic neurons, by responding with more spikes to pattern P than to pattern N. Our analysis is based on the assumption that there exist internal rewards $d(t)$ that could guide such pattern discrimination. This reward based learning architecture is biologically more plausible than an architecture with a supervisor which provides for each input pattern a target output and thereby directly produces the desired firing behavior of the neuron (since the question becomes then how the supervisor has learnt to produce the desired spike outputs).

We consider a neuron that receives input from n presynaptic neurons. A pattern X consists of n spike trains, each of time length T, one for each presynaptic neuron. There are two patterns, P and N, which are presented in alternation to the neuron, with some reset time between presentations. For notational simplicity, we assume that each of the n presynaptic spike trains consists of exactly one spike. Hence, each pattern can be defined by a list of spike times: $P = (t_1^P, \ldots, t_n^P)$, $N = (t_1^N, \ldots, t_n^N)$, where t_i^X is the time when presynaptic neuron i spikes for pattern $X \in \{P,N\}$. A generalization to the easier case of learning to discriminate spatio-temporal presynaptic firing patterns (where some presynaptic neurons produce different numbers of spikes in different patterns) is straightforward, however the main characteristics of the learning dynamics are better accessible in this conceptually simpler setup. It had already been shown in [12] that neurons can learn through reward-modulated STDP to discriminate between different *spatial* presynaptic firing patterns. But in the light of the analysis of [27] it is still open whether neurons can learn with simple forms of reward-modulated STDP, such as the one considered in this article, to discriminate *temporal* presynaptic firing patterns.

We assume that the reward signal $d(t)$ rewards—after some delay d_r—action potentials of the trained neuron if pattern P was presented, and punishes action potentials of the neuron if pattern N was presented. More precisely, we assume that

$$d(t) = \begin{cases} \alpha^P \int_0^{\infty} dr\, \varepsilon_r(r) S^{post}(t - d_r - r), \\ \text{if a pattern } P \text{ was presented} \\ \alpha^N \int_0^{\infty} dr\, \varepsilon_r(r) S^{post}(t - d_r - r), \\ \text{if a pattern } N \text{ was presented} \end{cases} \quad (16)$$

with some reward kernel ε_r and constants $\alpha^N < 0 < \alpha^P$. The goal of this learning task is to produce many output spikes for pattern P, and few or no spikes for pattern N.

The main result of our analysis is an estimate of the expected weight change of synapse i of the trained neuron for the presentation of pattern P, followed after a sufficiently long time T' by a presentation of pattern N

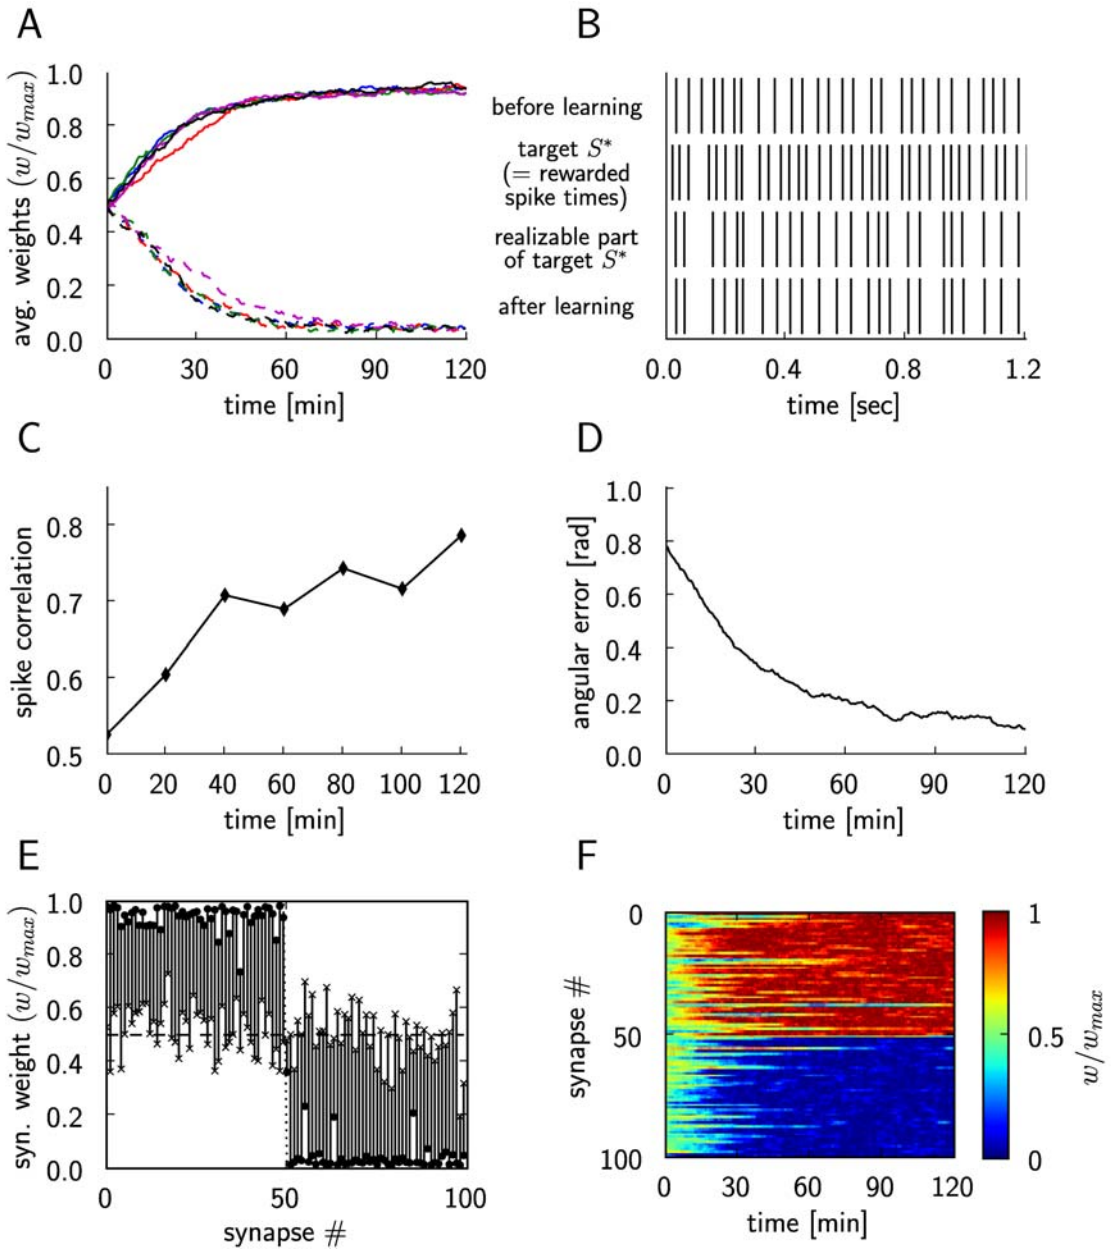

Figure 7. Results for reinforcement learning of exact spike times through reward-modulated STDP. (A) Synaptic weight changes of the trained LIF neuron, for 5 different runs of the experiment. The curves show the average of the synaptic weights that should converge to $w_i^* = 0$ (dashed lines), and the average of the synaptic weights that should converge to $w_i^* = w_{max}$ (solid lines) with different colors for each simulation run. (B) Comparison of the output of the trained neuron before (top trace) and after learning (bottom trace). The same input spike trains and the same noise inputs were used before and after training for 2 hours. The second trace from above shows those spike times S^* which are rewarded, the third trace shows the realizable part of S^* (i.e. those spikes which the trained neuron could potentially learn to reproduce, since the neuron μ^* produces them without its 10 extra spike inputs). The close match between the third and fourth trace shows that the trained neuron performs very well. (C) Evolution of the spike correlation between the spike train of the trained neuron and the realizable part of the target spike train S^*. (D) The angle between the weight vector **w** of the trained neuron and the weight vector **w*** of the neuron μ^* during the simulation, in radians. (E) Synaptic weights at the beginning of the simulation are marked with $\times$, and at the end of the simulation with •, for each plastic synapse of the trained neuron. (F) Evolution of the synaptic weights w/w_{max} during the simulation (we had chosen $w_i^* = w_{max}$ for $i<50$, $w_i^* = 0$ for $i \geq 50$).
doi:10.1371/journal.pcbi.1000180.g007

$$\Delta w_i = \int_0^{T'} dt \left[\langle \frac{dw_i(t)}{dt} \rangle_{E|P} + \langle \frac{dw_i(t)}{dt} \rangle_{E|N} \right],$$

where $\langle \cdot \rangle_{E|X}$ is the expectation over the ensemble given that pattern X was presented. This weight change can be estimated as (see Methods)

$$\Delta w_i = \int_{-\infty}^{\infty} dr W(r) \left[v^P \left(t_i^P + r \right) A_i^P + v^N \left(t_i^N + r \right) A_i^N \right], \quad (17)$$

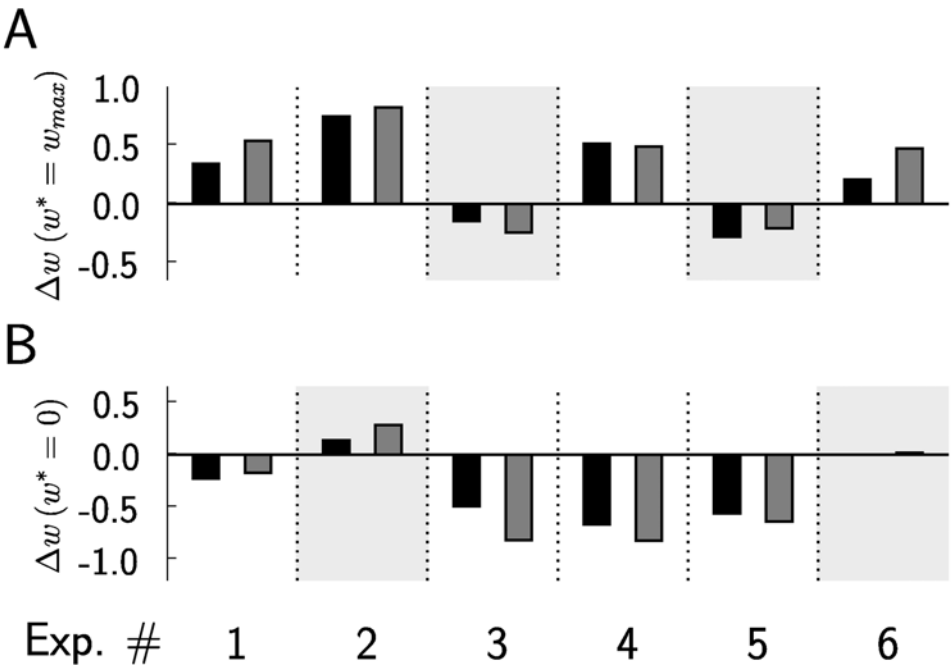

Figure 8. Test of the validity of the analytically derived conditions 13–15 on the relationship between parameters for successful learning with reward-modulated STDP. Predicted average weight changes (black bars) calculated from Equation 22 match in sign and magnitude the actual average weight changes (gray bars) in computer simulations, for 6 different experiments with different parameter settings (see Table 1). (A) Weight changes for synapses with $w_i^* = w_{max}$. (B) Weight changes for synapses with $w_i^* = 0$. Four cases where constraints 13–15 are not fulfilled are shaded in light gray. In all of these four cases the weights move into the opposite direction, i.e., a direction that decreases rewards.
doi:10.1371/journal.pcbi.1000180.g008

where $v^X(t)$ is the postsynaptic rate at time t for pattern X, and the constants A_i^X for $X \in \{P,N\}$ are given by

$$A_i^X = \alpha^X \int_0^\infty dr' \varepsilon_r(r') \left[f_c(d_r + r') + \int_0^{T'} dt f_c\left(t - t_i^X\right) v^X(t - d_r - r') \right]. \quad (18)$$

As we will see shortly, an interesting learning effect is achieved if A_i^P is positive and A_i^N is negative. Since $f_c(r)$ is non-negative, a natural way to achieve this is to choose a positive reward kernel $\varepsilon_r(r) \geq 0$ for $r > 0$ and $\varepsilon_r(r) = 0$ for $r < 0$ (also, $f_c(r)$ and $\varepsilon_r(r)$ must not be identical to zero for all r).

We use Equation 17 to provide insight on when and how the classification of temporal spike patterns can be learnt with reward-modulated STDP. Assume for the moment that $A_i^N = -A_i^P$. We first note that it is impossible to achieve through any synaptic plasticity rule that the time integral over the membrane potential of the trained neuron has after training a larger value for input pattern P than for input pattern N. The reason is that each presynaptic neuron emits the same number of spikes in both patterns (namely one spike). This simple fact implies that it is impossible to train a linear Poisson neuron (with any learning method) to respond to pattern P with more spikes than to pattern N. But Equation 17 implies that reward-modulated STDP increases the variance of the membrane potential for pattern P, and reduces the variance for pattern N. This can be seen as follows. Because of the specific form of the STDP learning curve $W(r)$, which is positive for (small) positive r, negative for (small) negative r, and zero for large r, $\Delta w_i = \int_{-\infty}^{\infty} dr W(r) v^P\left(t_i^P + r\right) A_i^P$ has a potentiating effect on synapse i if the postsynaptic rate for pattern P is larger (because of a higher membrane potential) shortly after the presynaptic spike at this synapse i than before that spike. This tends to further increase the membrane potential after that spike. On the other hand, since A_i^N is negative, the same situation for pattern N has a depressing effect on synapse i, which

Table 1. Parameter values used for computer simulation 3 (see Figure 8).

Ex.	τ_ε [ms]	w_{max}	$\nu^{post}{}_{min}$ [Hz]	$A_+ 10^6$	A_-/A_+	τ_+ [ms]	$A^{\kappa}{}_+$, $A^{\kappa}{}_-$	$\tau^{\kappa}{}_2$ [ms]	t_{sim} [h]
1	10	0.012	10	16.62	1.05	20	3.34, −3.12	20	5
2	7	0.020	5	11.08	1.02	15	4.58, −4.17	16	10
3	20	0.010	6	5.54	1.10	25	1.50, −1.39	40	19
4	7	0.020	5	11.08	1.07	25	4.67, −4.17	16	13
5	10	0.015	6	20.77	1.10	25	3.75, −3.12	20	2
6	25	0.005	3	13.85	1.01	25	3.34, −3.12	20	18

doi:10.1371/journal.pcbi.1000180.t001

counteracts the increased membrane potential after the presynaptic spike. Dually, if the postsynaptic rate shortly after the presynaptic spike at synapse i is lower than shortly before that spike, the effect on synapse i is depressing for pattern P. This leads to a further decrease of the membrane potential after that spike. In the same situation for pattern N, the effect is potentiating, again counteracting the variation of the membrane potential. The total effect on the postsynaptic membrane potential is that the fluctuations for pattern P are increased, while the membrane potential for pattern N is flattened.

For the LIF neuron model, and most reasonable other nonlinear spiking neuron models, as well as for biological neurons in-vivo and in-vitro [28–30], larger fluctuations of the membrane potential lead to more action potentials. As a result, reward-modulated STDP tends to increase the number of spikes for pattern P for these neuron models, while it tends to decrease the number of spikes for pattern N, thereby enabling a discrimination of these purely temporal presynaptic spike patterns.

Computer Simulation 4: Learning Pattern Classification

We tested these theoretical predictions through computer simulations of a LIF neuron with conductance based synapses exhibiting short-term depression and facilitation. Both patterns, P and N, had 200 input channels, with 1 spike per channel (hence this is the extreme where *all* information lies in the timing of presynaptic spikes). The spike times were drawn from an uniform distribution over a time interval of 500 ms, which was the duration of the patterns. We performed 1000 training trials where the patterns P and N were presented to the neuron in alternation. To introduce exploration for this reinforcement learning task, the neuron had injected 20% of the Ornstein-Uhlenbeck process conductance noise (see Methods for further details).

The theoretical analysis predicted that the membrane potential will have after learning a higher variance for pattern P, and a lower variance for pattern N. When in our simulation of a LIF neuron the firing of the neuron was switched off (by setting the firing threshold potential too high) we could observe the membrane potential fluctuations undisturbed by the reset mechanism after each spike (see Figure 9C and 9D). The variance of the membrane potential did in fact increase for pattern P from 2.49 $(\mathrm{mV})^2$ to 5.43 $(\mathrm{mV})^2$ (Figure 9C), and decrease for pattern N (Figure 9D), from 2.34 $(\mathrm{mV})^2$ to 1.33 $(\mathrm{mV})^2$. The corresponding plots with the firing threshold included are given in panels E and F, showing an increased member of spikes of the LIF neuron for pattern P, and a decreased number of spikes for pattern N. Furthermore, as Figure 9A and 9B show, the increased variance of the membrane potential for the positively reinforced pattern P led to a stable temporal firing pattern in response to pattern P.

We repeated the experiment 6 times, each time with different randomly generated patterns P and N, and different random initial synaptic weights of the neuron. The results in Figure 9G and 9H show that the learning of temporal pattern discrimination through reward-modulated STDP does not depend on the temporal patterns that are chosen, nor on the initial values of synaptic weights.

Computer Simulation 5: Training a Readout Neuron with Reward-Modulated STDP To Recognize Isolated Spoken Digits

A longstanding open problem is how a biologically realistic neuron model can be trained in a biologically plausible manner to extract information from a generic cortical microcircuit. Previous work [31–35] has shown that quite a bit of salient information about recent and past inputs to the microcircuit can be extracted by a non-spiking linear readout neuron (i.e., a perceptron) that is trained by linear regression or margin maximization methods. Here we examine to what extent a LIF readout neuron with conductance based synapses (subject to biologically realistic short term synaptic plasticity) can learn through reward-modulated STDP to extract from the response of a simulated cortical microcircuit (consisting of 540 LIF neurons), see Figure 10A, the information which spoken digit (transformed into spike trains by a standard cochlea model) is injected into the circuit. In comparison with the preceding task in simulation 4, this task is easier because the presynaptic firing patterns that need to be discriminated differ in temporal and spatial aspects (see Figure 10B; Figure S10 and S11 show the spike trains that were injected into the circuit). But this task is on the other hand more difficult, because the circuit response (which creates the presynaptic firing pattern for the readout neuron) differs also significantly for two utterances of the same digit (Figure 10C), and even for two trials for the same utterance (Figure 10D) because of the intrinsic noise in the circuit (which was modeled according to [26] to reflect in-vivo conditions during cortical UP-states). The results shown in Figure 10E–H demonstrate that nevertheless this learning experiment was successful. On the other hand we were not able to achieve in this way speaker-independent word recognition, which had been achieved in [31] with a linear readout. Hence further work will be needed in order to clarify whether biologically more realistic models for readout neurons can be trained through reinforcement learning to reach the classification capabilities of perceptrons that are trained through supervised learning.

Methods

We first describe the simple neuron model that we used for the theoretical analysis, and then provide derivations of the equations that were discussed in the preceding section. After that we describe the models for neurons, synapses, and synaptic background activity ("noise") that we used in the computer simulations. Finally we provide technical details to each of the 5 computer simulations that we discussed in the preceding section.

Linear Poisson Neuron Model

In our theoretical analysis, we use a linear Poisson neuron model whose output spike train $S_j^{post}(t)$ is a realization of a Poisson process with the underlying instantaneous firing rate $R_j(t)$. The effect of a spike of presynaptic neuron i at time t' on the membrane potential of neuron j is modeled by an increase in the instantaneous firing rate by an amount $w_{ji}(t')\varepsilon(t-t')$, where ε is a response kernel which models the time course of a postsynaptic potential (PSP) elicited by an input spike. Since STDP according to [12] has been experimentally confirmed only for excitatory synapses, we will consider plasticity only for excitatory connections and assume that $w_{ji} \geq 0$ for all i and $\varepsilon(s) \geq 0$ for all s. Because the synaptic response is scaled by the synaptic weights, we can assume without loss of generality that the response kernel is normalized to $\int_0^\infty ds\ \varepsilon(s) = 1$. In this linear model, the contributions of all inputs are summed up linearly:

$$R_j(t) = \sum_{i=1}^{n} \int_0^\infty ds\ w_{ji}(t-s)\varepsilon(s)S_i(t-s), \qquad (19)$$

where $S_1,\ldots,S_n$ are the n presynaptic spike trains. Since the instantaneous firing rate $R(t)$ is analogous to the membrane potential of other neuron models, we occasionally refer to $R(t)$ as the "membrane potential" of the neuron.

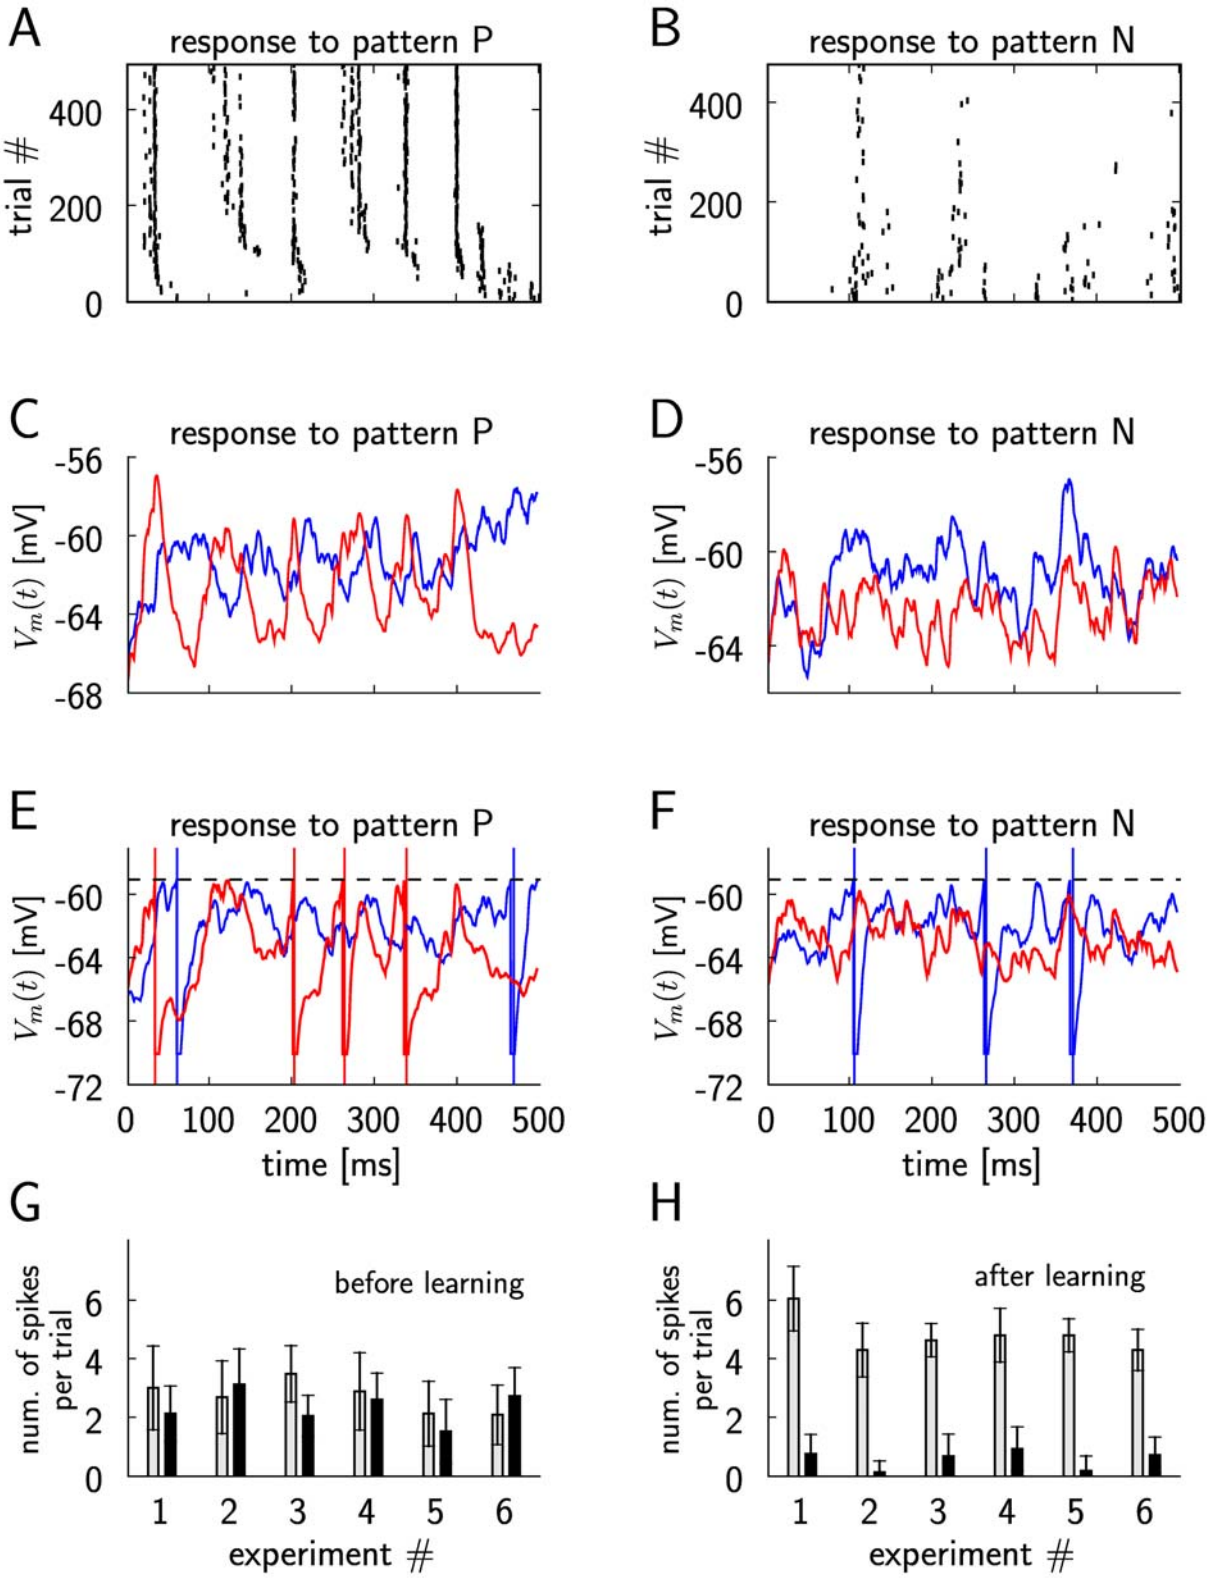

Figure 9. Training a LIF neuron to classify purely temporal presynaptic firing patterns: a positive reward is given for firing of the neuron in response to a temporal presynaptic firing pattern *P*, and a negative reward for firing in response to another temporal pattern *N*. (A) The spike response of the neuron for individual trials, during 500 training trials when pattern *P* is presented. Only the spikes from every 4-th trial are plotted. (B) As in (A), but in response to pattern *N*. (C) The membrane potential $V_m(t)$ of the neuron during a trial where pattern *P* is presented, before (blue curve) and after training (red curve), with the firing threshold removed. The variance of the membrane potential increases during learning, as predicted by the theory. (D) As in (C), but for pattern *N*. The variance of the membrane potential for pattern *N* decreases during learning, as predicted by the theory. (E) The membrane potential $V_m(t)$ of the neuron (including action potentials) during a trial where pattern *P* is presented before (blue curve) and after training (red curve). The number of spikes increases. (F) As in (E), but for trials where pattern *N* is given as input. The number of spikes decreases. (G) Average number of output spikes per trial before learning, in response to pattern *P* (gray bars) and pattern *N* (black bars), for 6 experiments with different randomly generated patterns *P* and *N*, and different random initial synaptic weights of the neuron. (H) As in (G), for the same experiments, but after learning. The average number of spikes per trial increases after training for pattern *P*, and decreases for pattern *N*.
doi:10.1371/journal.pcbi.1000180.g009

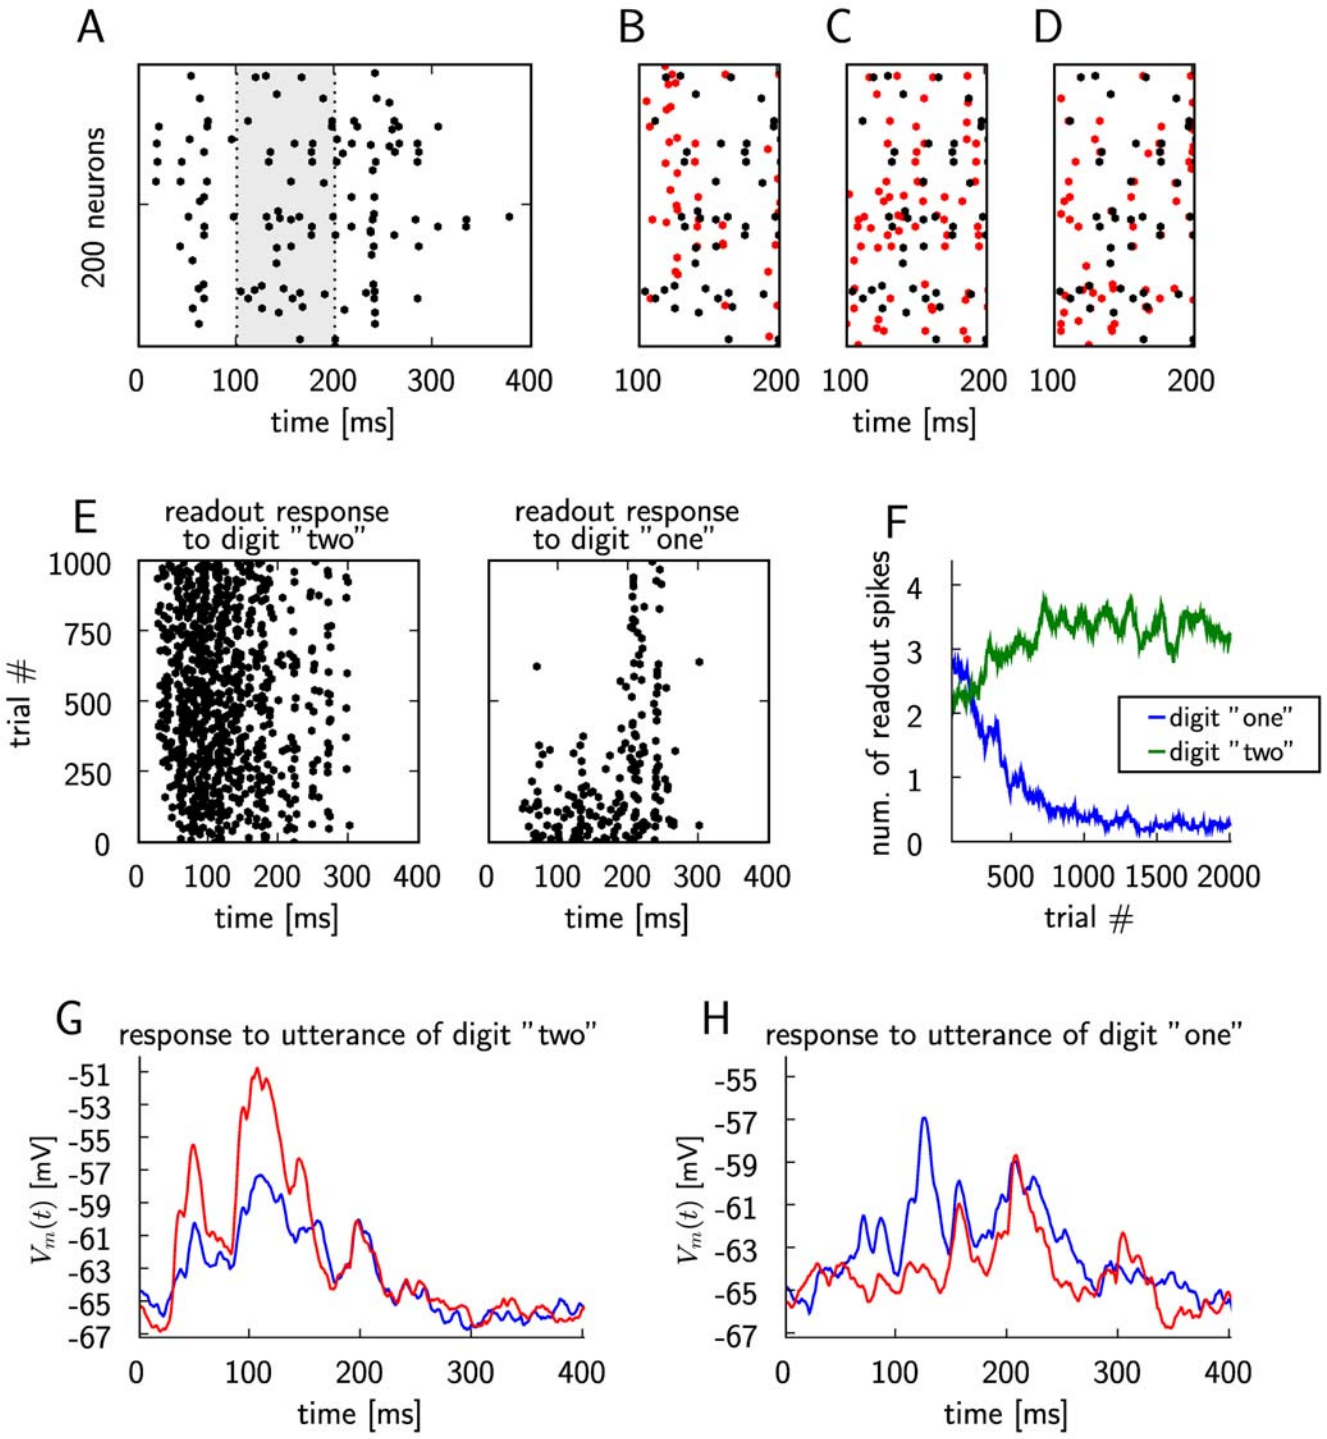

Figure 10. A LIF neuron is trained through reward-modulated STDP to discriminate as a ''readout neuron'' responses of generic cortical microcircuits to utterances of different spoken digits. (A) Circuit response to an utterance of digit ''one'' (spike trains of 200 out of 540 neurons in the circuit are shown). The response within the time period from 100 to 200 ms (marked in gray) is used as a reference in the subsequent 3 panels. (B) The circuit response from (A) (black) for the period between 100 and 200 ms, and the circuit response to an utterance of digit ''two'' (red). (C) The circuit spike response from (A) (black) and a circuit response for another utterance of digit ''one'' (red), also shown for the period between 100 and 200 ms. (D) The circuit spike response from (A) (black), and another circuit response to the same utterance in another trial (red). The responses differ due to the presence of noise in the circuit. (E) Spike response of the LIF readout neuron for different trials during learning, for trials where utterances of digit ''two'' (left plot) and digit ''one'' (right plot) are presented as circuit inputs. The spikes from each 4th trial are plotted. (F) Average number of spikes in the response of the readout during training, in response to digit ''one'' (blue) and digit ''two'' (green). The number of spikes were averaged over 40 trials. (G) The membrane potential $V_m(t)$ of the neuron during a trial where an input pattern corresponding to an utterance of digit ''two'' is presented, before (blue curve) and after training (red curve), with the firing threshold removed. (H) As in (G), but for an input pattern corresponding to an utterance of digit ''one''. The variance of the membrane potential increases during learning for utterances of the rewarded digit, and decreases for the non-rewarded digit.

doi:10.1371/journal.pcbi.1000180.g010

Learning Equations

In the following, we denote by $\langle x\rangle_{E|S_k^{post}(t),S_i^{pre}(t')}$ the ensemble average of a random variable x given that neuron k spikes at time t and neuron i spikes at time t'. We will also sometimes indicate the variables $Y_1,Y_2,\ldots$ over which the average of x is taken by writing $\langle x\rangle_{Y_1,Y_2,\ldots|\ldots}$.

Derivation of Equation 6. Using Equations 5, 1, and 4, we obtain the expected weight change between time t and $t+T$

$$\begin{aligned}
&\frac{\langle w_{ji}(t+T)-w_{ji}(t)\rangle_E}{T}=\\
&\int_0^\infty ds f_c(s)\int_0^\infty dr W(r)\langle\langle d(t)S_j^{post}(t-s)S_i^{pre}(t-s-r)\rangle_T\rangle_E+\\
&\int_0^\infty ds\, f_c(s)\int_{-\infty}^0 dr\, W(r)\langle\langle d(t)S_j^{post}(t-s+r)S_i^{pre}(t-s)\rangle_T\rangle_E\\
&=\int_0^\infty dr\, W(r)\int_0^\infty ds\, f_c(s)\langle\langle d(t)S_j^{post}(t-s)S_i^{pre}(t-s-r)\rangle_E\rangle_T+\\
&\int_{-\infty}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle\langle d(t)S_j^{post}(t-s)S_i^{pre}(t-s-r)\rangle_E\rangle_T\\
&=\int_0^\infty dr\, W(r)\int_0^\infty ds\, f_c(s)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T+\\
&\int_{-\infty}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T,
\end{aligned}$$

with $D_{ji}(t,s,r)=\langle d(t)\,|\,\text{Neuron } j \text{ spikes at } t-s, \text{ and neuron } i \text{ spikes at } t-s-r\rangle_E$, and the joint firing rate $v_{ji}(t,r)=\langle S_j(t)S_i(t-r)\rangle_E$ describes correlations between spike timings of neurons j and i. The joint firing rate $v_{ji}(t-s,r)$ depends on the weight at time $t-s$. If the learning rate defined by the magnitude of $W(r)$ is small, the synaptic weights can be assumed constant on the time scale of T. Thus, the time scales of neuronal dynamics are separated from the slow time scale of learning. For slow learning, synaptic weights integrate a large number of small changes. We can then expect that averaged quantities enter the learning dynamics. In this case, we can argue that fluctuations of a weight w_{ji} about its mean are negligible and it can well be approximated by its average $\langle w_{ji}\rangle_E$ (it is "self-averaging", see [21,36]). To ensure that average quantities enter the learning dynamics, many presynaptic and postsynaptic spikes as well as many independently delivered rewards at varying delays have to occur within T. Hence, in general, the time scale of single spike occurrences and the time scale of the eligibility trace is required to be much smaller than the time scale of learning. If time scales can be separated, we can drop the expectation on the left hand side of the last equation and write

$$\begin{aligned}
\frac{\langle w_{ji}(t+T)-w_{ji}(t)\rangle_E}{T}&=\frac{w_{ji}(t+T)-w_{ji}(t)}{T}\\
&=\frac{1}{T}\int_t^{t+T}\frac{d}{dt}w_{ji}(t')dt'=\frac{d}{dt}\langle w_{ji}(t)\rangle_T.
\end{aligned}$$

We thus obtain Equation 6:

$$\begin{aligned}
\frac{d}{dt}\langle w_{ji}(t)\rangle_T=&\int_0^\infty dr\, W(r)\int_0^\infty ds\, f_c(s)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T\\
&+\int_{-\infty}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T.
\end{aligned}$$

Simplification of Equation 6. In order to simplify this equation, we first observe that $W(r)$ is vanishing for large $|r|$. Hence we can approximate the integral over the learning window by a bounded integral $\int_{-\infty}^{\infty} dr\, W(r)\approx\int_{-T_W}^{T_W} dr\, W(r)$ for some $T_W>0$ and $T_W\ll T$. In the analyzes of this article, we consider the case where reward is delivered with a relatively large temporal delay. To be more precise, we assume that a pre-post spike pair has an effect on the reward signal only after some minimal delay d_r and that we can write $D_{ji}(t,s,r)=d_0+D_{ji}^{pre,post}(t,s,r)$ for some baseline reward d_0 and a part which depends on the timing of pre-post spike pairs with $D_{ji}^{pre,post}(t,s,r)=0$ for $s<d_r$ and $d_r>T_W$. We can then approximate the second term of Equation 6:

$$\begin{aligned}
&\int_{-\infty}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T\\
&\approx\int_{-T_W}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle\left(d_0+D_{ji}^{pre,post}(t,s,r)\right)v_{ji}(t-s,r)\rangle_T\\
&\approx\int_{-T_W}^0 dr\, W(r)\left[\int_0^\infty ds\, f_c(s)d_0\langle v_{ji}(t-s,r)\rangle_T\right.\\
&\left.+\int_{|r|}^\infty ds\, f_c(s+r)\langle D_{ji}^{pre,post}(t,s,r)v_{ji}(t-s,r)\rangle_T\right]
\end{aligned}$$

because $\langle v_{ji}(t-s-r,r)\rangle_T\approx\langle v_{ji}(t-s,r)\rangle_T$ for $r\in[-T_W,T_W]$ and $T_W\ll T$. Since $D_{ji}^{pre,post}(t,s,r)=0$ for $s\le T_W$, the second term in the brackets is equivalent to $\int_0^\infty ds\, f_c(s+r)\langle D_{ji}^{pre,post}(t,s,r)\,v_{ji}(t-s,r)\rangle_T$ which in turn is approximately given by $\int_0^\infty ds\, f_c(s)\langle D_{ji}^{pre,post}(t,s,r)v_{ji}(t-s,r)\rangle_T$ if we assume that $f_c(s+r)\approx f_c(s)$ for $s\ge d_r$ and $|r|<T_W$. We can thus approximate the second term of Equation 6 as

$$\begin{aligned}
&\int_{-\infty}^0 dr\, W(r)\int_{|r|}^\infty ds\, f_c(s+r)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T\\
&\approx\int_{-T_W}^0 dr\, W(r)\left[\int_0^\infty ds\, f_c(s)d_0\langle v_{ji}(t-s,r)\rangle_T\right.\\
&\left.+\int_0^\infty ds\, f_c(s)\langle D_{ji}^{pre,post}(t,s,r)v_{ji}(t-s,r)\rangle_T\right]\\
&\approx\int_{\infty}^0 dr\, W(r)\int_0^\infty ds\, f_c(s)\langle D_{ji}(t,s,r)v_{ji}(t-s,r)\rangle_T.
\end{aligned}$$

With this approximation, the first and second term of Equation 6 can be combined in a single integral to obtain Equation 8.

Derivations for the Biofeedback Experiment

We assume that a reward with the functional form ε_r is delivered for each postsynaptic spike with a delay d_r. The reward as time t is therefore

$$d(t)=\int_0^\infty dr\, S_k^{post}(t-d_r-r)\varepsilon_r(r).$$

Weight change for the reinforced neuron (derivation of Equation 10)

The reward correlation for a synapse ki afferent to the reinforced neuron is

$$
\begin{aligned}
&D_{ki}(t,s,r) = \langle d(t)\rangle_{E|S_k^{post}(t-s),S_i^{pre}(t-s-r)} \\
&= \int_0^\infty dr'\ \varepsilon_r(r')\langle S_k^{post}(t-d_r-r')\rangle_{E|S_k^{post}(t-s),S_i^{pre}(t-s-r)} \\
&= \int_0^\infty dr'\ \varepsilon_r(r')[v_k(t-d_r-r') + w_{ki}\varepsilon(s+r-d_r-r') + \delta(s-d_r-r')] \\
&= \int_0^\infty dr'\varepsilon_r(r')v_k(t-d_r-r') + w_{ki}\int_0^\infty dr'\ \varepsilon_r(r')\varepsilon(s+r-d_r-r') + \varepsilon_r(s-d_r).
\end{aligned}
$$

If we assume that the output firing rate is constant on the time scale of the reward function, the first term vanishes. We rewrite the result as

$$
D_{ki}(t,s,r) = \varepsilon_r(s-d_r) + w_{ki}\int_{-\infty}^{\infty} dr'\ \varepsilon_r(s-d_r+r')\varepsilon(r-r').
$$

The mean weight change for weights to the reinforced neuron is therefore

$$
\begin{aligned}
\frac{d}{dt}w_{ki}(t) = \int_{-\infty}^{\infty} dr\ W(r)\Bigg(&\int_0^\infty ds\, f_c(s)\varepsilon_r(s-d_r)\langle v_{ki}(t-s,r)\rangle_T + \\
&w_{ki}\int_{-\infty}^{\infty} dr'\ \varepsilon(r-r')\int_0^\infty ds\, f_c(s)\varepsilon_r(s-d_r+r')\langle v_{ki}(t-s,r)\rangle_T\Bigg).
\end{aligned}
\tag{20}
$$

We show that the second term in the brackets is very small compared to the first term:

$$
\begin{aligned}
&w_{ki}\int_{-\infty}^{\infty} dr'\ \varepsilon(r-r')\int_0^\infty ds\, f_c(s)\varepsilon_r(s-d_r+r')\langle v_{ki}(t-s,r)\rangle_T = \\
&w_{ki}\int_{-\infty}^{\infty} dr'\ \varepsilon(r-r')\int_0^\infty ds\, f_c(s-r')\varepsilon_r(s-d_r)\langle v_{ki}(t-s-r',r)\rangle_T \approx \\
&w_{ki}\int_{-\infty}^{\infty} dr'\ \varepsilon(r-r')\int_0^\infty ds\, f_c(s)\varepsilon_r(s-d_r)\langle v_{ki}(t-s,r)\rangle_T.
\end{aligned}
$$

The last approximation is based on the assumption that $f_c(s)\approx f_c(s-r')$ and $\langle v_{ki}(t-r',r)\rangle_T \approx \langle v_{ki}(t,r)\rangle_T$ for $r'\in[-T_W-T_\varepsilon, T_W]$. Here, T_W is the time scale of the learning window (see above), and T_ε is time scale of the PSP, i.e., we have $\varepsilon(s)\approx 0$ for $s \geq T_\varepsilon$. Since $\int_{-\infty}^{\infty} dr\ \varepsilon(r) = 1$ by definition, we see that this is the first term in the brackets of Equation 20 scaled by w_{ki}. For neurons with many input synapses we have $w_{ki} \ll 1$. Thus the second term in the brackets of Equation 20 is small compared to the first term. We therefore have

$$
\frac{d}{dt}w_{ki}(t) \approx \int_0^\infty ds\, f_c(s+d_r)\varepsilon_r(s)\int_{-\infty}^{\infty} dr\ W(r)\langle v_{ki}(t-d_r-s,r)\rangle_T.
$$

Weight change for non-reinforced neurons (derivation of Equation 11)

The reward correlation of a synapse ji to a non-reinforced neuron j is given by

$$
\begin{aligned}
D_{ji}(t,s,r) &= \langle d(t)\rangle_{E|S_j^{post}(t-s),S_i^{pre}(t-s-r)} \\
&= \int_0^\infty dr'\ \varepsilon_r(r')\langle S_k^{post}(t-d_r-r')\rangle_{E|S_j^{post}(t-s),S_i^{pre}(t-s-r)}.
\end{aligned}
$$

We have

$$
\begin{aligned}
&\langle S_k^{post}(t-d_r-r')\rangle_{E|S_j^{post}(t-s),S_i^{pre}(t-s-r)} \\
&= \frac{\langle S_k^{post}(t-d_r-r')S_j^{post}(t-s)\rangle_{E|S_i^{pre}(t-s-r)}}{\langle S_j^{post}(t-s)\rangle_{E|S_i^{pre}(t-s-r)}} \\
&= \frac{v_{ki}(t-d_r-r',s-d_r-r') + w_{ki}w_{ji}\varepsilon(s+r-d_r-r')\varepsilon(r)}{v_j(t-s) + w_{ji}\varepsilon(r)},
\end{aligned}
$$

for which we obtain

$$
\begin{aligned}
&D_{ji}(t,s,r) = \\
&\int_0^\infty dr'\ \varepsilon_r(r')\frac{v_{kj}(t-d_r-r',s-d_r-r') + w_{ki}w_{ji}\varepsilon(s+r-d_r-r')\varepsilon(r)}{v_j(t-s) + w_{ji}\varepsilon(r)}.
\end{aligned}
$$

In analogy to the previous derivation, we assume here that the firing rate $v_j(t-s)$ in the denominator results from many PSPs. Hence, the single PSP $w_{ji}\varepsilon(r)$ is small compared to $v_j(t-s)$. Similarly, we assume that with weights w_{ki}, $w_{ji} \ll 1$, the second term in the nominator is small compared to the joint firing rate $v_{kj}(t-d_r-r',s-d_r-r')$. We therefore approximate the reward correlation by

$$
D_{ji}(t,s,r) \approx \int_0^\infty dr'\ \varepsilon_r(r')\frac{v_{kj}(t-d_r-r',s-d_r-r')}{v_j(t-s)}.
$$

Hence, the reward correlation of a non-reinforced neuron depends on the correlation of this neuron with the reinforced neuron. The mean weight change for a non-reinforced neuron $j \neq k$ is therefore

$$
\frac{d}{dt}w_{ji}(t) \approx \int_0^\infty ds\, f_c(s)\int_{-\infty}^{\infty} dr\ W(r)\int_0^\infty dr'\varepsilon_r(r')\langle\frac{v_{kj}(t-d_r-r',s-d_r-r')}{v_j(t-s)}v_{ji}(t-s,r)\rangle_T
$$

This equation deserves a remark for the case that $v_j(t-s)$ is zero, since it appears in the denominator of the fraction. Note that in this case, both $v_{kj}(t-d_r-r',s-d_r-r')$ and $v_{ji}(t-s,r)$ are zero. In fact, if we take the limit $v_j(t-s)\to 0$, then both of these factors approach zero at least as fast. Hence, in the limit of $v_j(t-s)\to 0$, the term in the angular brackets evaluates to zero. This reflects the fact that since STDP is driven by pre- and postsynaptic spikes, there is no weight change if no postsynaptic spikes occur.

For uncorrelated neurons, Equation 11 evaluates to zero. For uncorrelated neurons k, j, $v_{kj}(t-d_r-r',s-d_r-r')$ can be factorized into $v_k(t-d_r-r')v_j(t-s)$, and we obtain

$$
\frac{d}{dt}w_{ji}(t) \approx \int_0^\infty ds\, f_c(s)\int_{-\infty}^{\infty} dr\ W(r)\int_0^\infty dr'\varepsilon_r(r')\langle v_k(t-d_r-r')v_{ji}(t-s,r)\rangle_T.
$$

This evaluates approximately to zero if the mean output rate of neuron k is constant on the time scale of the reward kernel.

Analysis of Spike-Timing-Dependent Rewards (Derivation of Conditions 13–15)

Below, we will indicate the variables $Y_1, Y_2, \ldots$ over which the average of x is taken by writing $\langle x\rangle_{Y_1,Y_2,\ldots|\ldots}$. From Equation 12, we can determine the reward correlation for synapse i

$$D_{ji}(t,s,r)=\int_{-\infty}^{\infty}dr'\kappa(r')\langle S_j^{post}(t-d_r)S^*(t-d_r-r')\rangle_{E|S_j^{post}(t-s),S_i^{pre}(t-s-r)}$$
$$=\int_{-\infty}^{\infty}dr'\kappa(r')\left[v_j^{post}(t-d_r)+\delta(s-d_r)+w_{ji}(s+r-d_r)\varepsilon(s+r-d_r)\right]$$
$$\left[v^*(t-d_r-r')+w_i^*\varepsilon(s+r-d_r-r')\right], \quad (21)$$

where $v_j^{post}(t)=\langle S_j^{post}(t)\rangle_E$ denotes the instantaneous firing rate of the trained neuron at time t, and $v^*(t)=\langle S^*(t)\rangle_E$ denotes the instantaneous rate of the target spike train at time t. Since weights are changing very slowly, we have $w_{ji}(t-s-r)\approx w_{ji}(t)$. In the following, we will drop the dependence of w_{ji} on t for brevity. For simplicity, we assume that input rates are stationary and uncorrelated. In this case (since the weights are changing slowly), also the correlations between inputs and outputs can be assumed stationary, $v_{ji}(t,r)=v_{ji}(r)$. With constant input rates, we can rewrite Equation 21 as

$$D_{ji}(t,s,r)=\overline{\kappa}v^*v_j^{post}+\overline{\kappa}v^*\delta(s-d_r)+\overline{\kappa}v^*w_{ji}\varepsilon(s+r-d_r)$$
$$+w_i^*\int_{-\infty}^{\infty}dr'\kappa(r')\varepsilon(s+r-d_r-r')$$
$$\left[v_j^{post}(t-d_r)+\delta(s-d_r)+w_{ji}(s+r-d_r)\varepsilon(s+r-d_r)\right],$$

with $\overline{\kappa}=\int_{-\infty}^{\infty}ds\ \kappa(s)$. We use this results to obtain the temporally smoothed weight change for synapse ji. With stationary correlations, we can drop the dependence of v_{ji} on t and write $v_{ji}(t,r)=v_{ji}(r)$. Furthermore, we define $v_{ji}^W(r)=v_{ji}(r)W(r)$ and obtain

$$\frac{d}{dt}w_{ji}(t)=\int_{-\infty}^{\infty}dr\ W(r)v_{ji}(r)\int_0^{\infty}ds\,f_c(s)\langle D_{ji}(t,s,r)\rangle_T$$
$$=\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)\overline{\kappa}\left[v^*v_j^{post}\overline{f}_c+v^*f_c(d_r)+v^*w_{ji}\int_0^{\infty}ds\,f_c(s)\varepsilon(s+r-d_r)\right]$$
$$+\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)w_i^*v^{post}\int_{-\infty}^{\infty}dr'\kappa(r')\int_0^{\infty}ds\,f_c(s)\varepsilon(s+r-d_r-r')$$
$$+\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)w_i^*\int_{-\infty}^{\infty}dr'\kappa(r')f_c(d_r)\varepsilon(r-r')$$
$$+\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)w_i^*\int_{-\infty}^{\infty}dr'\kappa(r')w_{ji}\int_0^{\infty}ds\,f_c(s)\varepsilon(s+r-d_r)\varepsilon(s+r-d_r-r').$$

We assume that the eligibility function $f_c(d_r)\approx f_c(d_r+r)$ if $|r|$ is on the time scale of a PSP, the learning window, or the reward kernel, and that d_r is large compared to these time scales. Then, we have

$$\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)\int_{-\infty}^{\infty}dr'\ \kappa(r')f_c(d_r)\varepsilon(r-r')=f_c(d_r)\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)\varepsilon_\kappa(r)$$

where $\varepsilon_\kappa(r)=\int_{-\infty}^{\infty}dr'\ \kappa(r')\varepsilon(r-r')$ is the convolution of the reward kernel with the PSP. Furthermore, we find

$$\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)\int_{-\infty}^{\infty}dr'\ \kappa(r')\int_0^{\infty}ds\,f_c(s)\varepsilon(s+r-d_r)\varepsilon(s+r-d_r-r')$$
$$\approx f_c(d_r)\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)\int_{-\infty}^{\infty}dr'\ \kappa(r')\int_0^{\infty}ds\ \varepsilon(s+r-d_r)\varepsilon(s+r-d_r-r')$$
$$=f_c(d_r)\int_{-\infty}^{\infty}dr\ v_{ji}^W(r)\int_0^{\infty}ds\ \varepsilon(s)\varepsilon_\kappa(s).$$

With these simplifications, and the abbreviation $\overline{v}_{ji}^W=\int_{-\infty}^{\infty}dr\,v_{ji}^W(r)$ we obtain the weight change at synapse ji

$$\frac{d}{dt}w_{ji}(t)\approx\overline{\kappa}v^*v_j^{post}\overline{v}_{ji}^W\overline{f}_c+f_c(d_r)\overline{\kappa}\overline{v}_{ji}^W\left[v^*+v^*w_{ji}+w_i^*v_j^{post}\right]$$
$$+f_c(d_r)w_i^*\int_{-\infty}^{\infty}dr W(r)v_{ji}(r)\varepsilon_\kappa(r)+f_c(d_r)w_{ji}w_i^*\overline{v}_{ji}^W\int_{-\infty}^{\infty}dr\ \varepsilon(r)\varepsilon_\kappa(r),$$

where $\overline{v}_{ji}^W=\int_{-\infty}^{\infty}dr\,W(r)v_{ji}(r)$.

For uncorrelated Poisson input spike trains of rate v_i^{pre} and the linear Poisson neuron model, the input-output correlations are $v_{ji}(r)=v_i^{pre}v_j^{post}+w_{ji}v_i^{pre}\varepsilon(r)$. With these correlations, we obtain $\overline{v}_{ji}^W=v_i^{pre}v_j^{post}\overline{W}+w_{ji}v_i^{pre}\overline{W}_\varepsilon$ where $\overline{W}=\int_{-\infty}^{\infty}dr\ W(r)$, and $\overline{W}_\varepsilon=\int_{-\infty}^{\infty}dr\ \varepsilon(r)W(r)$. The weight change at synapse ji is then

$$\frac{d}{dt}w_{ji}(t)\approx\overline{\kappa}\overline{f}_cv^*v_i^{pre}v_j^{post}\left[v_j^{post}\overline{W}+w_{ji}\overline{W}_\varepsilon\right]$$
$$+\overline{\kappa}f_c(d_r)v_i^{pre}\left[v_j^{post}\overline{W}+w_{ji}\overline{W}_\varepsilon\right]\left[v^*+v^*w_{ji}+w_i^*v_j^{post}\right]$$
$$+f_c(d_r)w_i^*v_i^{pre}\left[v_j^{post}\int_{-\infty}^{\infty}dr\ W(r)\varepsilon_\kappa(r)+w_{ji}\int_{-\infty}^{\infty}dr\ W(r)\varepsilon(r)\varepsilon_\kappa(r)\right] \quad (22)$$
$$+f_c(d_r)w_i^*w_{ji}v_i^{pre}\left[v_j^{post}\overline{W}+w_{ji}\overline{W}_\varepsilon\right]\int_0^{\infty}dr\ \varepsilon(r)\varepsilon_\kappa(r),$$

We will now bound the expected weight change for synapses ji with $w_i^*=w_{max}$ and for synapses jk with $w_k^*=0$. In this way we can derive conditions for which the expected weight change for the former synapses is positive, and that for the latter type is negative. First, we assume that the integral over the reward kernel is positive. In this case, the weight change given by Equation 22 is negative for synapses i with $w_i^*=0$ if and only if $v_i^{pre}>0$, and $-v_j^{post}\overline{W}>w_{ji}\overline{W}_\varepsilon$. In the worst case, w_{ji} is w_{max} and v_j^{post} is small. We have to guarantee some minimal output rate v_{min}^{post} such that even if $w_{ji}=w_{max}$, this inequality is fulfilled. This could be guaranteed by some noise current. Given such minimal output rate, we can state the first inequality which guarantees convergence of weights w_{ji} with $w_i^*=0$

$$-v_{min}^{post}\overline{W}>w_{max}\overline{W}_\varepsilon.$$

For synapses ji with $w_i^*=w_{max}$, we obtain two more conditions. The approximate weight change is given by

$$\frac{d}{dt}w_{ji}(t)\frac{1}{v_i^{pre}}\approx\overline{\kappa}\left[v_j^{post}\overline{W}+w_{ji}\overline{W}_\varepsilon\right]$$
$$\left[v^*v_j^{post}\overline{f}_c+f_c(d_r)v^*+f_c(d_r)v^*w_{ji}+f_c(d_r)v_j^{post}w_{max}\right]$$
$$+f_c(d_r)w_{max}v_j^{post}\int_{-\infty}^{\infty}dr\ W(r)\varepsilon_\kappa(r)$$
$$+f_c(d_r)w_{max}w_{ji}\int_{-\infty}^{\infty}dr\ W(r)\varepsilon(r)\varepsilon_\kappa(r)$$
$$+f_c(d_r)w_{max}w_{ji}v_j^{post}\overline{W}\int_0^{\infty}dr\ \varepsilon(r)\varepsilon_\kappa(r)$$
$$+f_c(d_r)w_{max}w_{ji}^2\overline{W}_\varepsilon\int_0^{\infty}dr\ \varepsilon(r)\varepsilon_\kappa(r).$$

The last term in this equation is positive and small. We can ignore it in our sufficient condition. The second to last term is negative. We will include in our condition that the third to last term compensates for this negative term. Hence, the second condition is

$$\int_{-\infty}^{\infty} dr\ W(r)\varepsilon(r)\varepsilon_\kappa(r) \geq -v_j^{post}\overline{W}\int_0^{\infty} dr\ \varepsilon(r)\varepsilon_\kappa(r),$$

which should be satisfied in most setups. If we assume that this holds, we obtain

$$\begin{aligned}&\frac{d}{dt}w_{ji}(t) \geq \overline{\kappa}\left[v_j^{post}\overline{W}+w_{ji}\overline{W}_\varepsilon\right]\\&\left[v^* v_j^{post}\overline{f}_c+f_c(d_r)v^*+f_c(d_r)v^* w_{ji}+f_c(d_r)v_j^{post}w_{max}\right]\\&+f_c(d_r)w_{max}v_j^{post}\int_{-\infty}^{\infty} dr\ W(r)\varepsilon_\kappa(r).\end{aligned}$$

which should be positive. We obtain the following inequality

$$\int_{-\infty}^{\infty} dr\ W(r)\varepsilon_\kappa(r) > -\overline{W}\overline{\kappa}\left[\frac{v^* v_j^{post}}{w_{max}}\frac{\overline{f}_c}{f_c(d_r)}+\frac{v^*}{w_{max}}+v^*+v^{post}\right].$$

All three inequalities are summarized in the following:

$$\begin{aligned}&-v_{min}^{post}\overline{W} > w_{max}\overline{W}_\varepsilon\\&\int_{-\infty}^{\infty} dr\ W(r)\varepsilon(r)\varepsilon_\kappa(r) \geq -v_{max}^{post}\overline{W}\int_0^{\infty} dr\ \varepsilon(r)\varepsilon_\kappa(r)\\&\int_{-\infty}^{\infty} dr\ W(r)\varepsilon_\kappa(r) > -\overline{W}\overline{\kappa}\left[\frac{v^* v_{max}^{post}}{w_{max}}\frac{\overline{f}_c}{f_c(d_r)}+\frac{v^*}{w_{max}}+v^*+v_{max}^{post}\right],\end{aligned}$$

where v_{max}^{post} is the maximal output rate. If these inequalities are fulfilled and input rates are positive, then the weight vector converges on average from any initial weight vector to $\mathbf{w}^*$. The second condition is less severe, and should be easily fulfilled in most setups. If this is the case, the first Condition 13 ensures that weights with $w^*=0$ are depressed while the third Condition 15 ensures that weights with $w^*=w_{max}$ are potentiated.

Analysis of the Pattern Discrimination Task (Derivation of Equation 17)

We assume that a trial consists of the presentation of a single pattern starting at time $t=0$. We compute the weight change for a single trial given that pattern $X\in\{P,N\}$ was presented with the help of Equations 1, 3, and 4 as

$$\begin{aligned}&\left.\frac{d}{dt}w_i(t)\right|_X=\int_0^{\infty} ds f_c(s)\left[\int_0^{\infty} dr W(r)S^{post}(t-s)\delta\left(t-s-r-t_i^X\right)\right.\\&\left.+\int_0^{\infty} dr W(-r)S^{post}(t-s-r)\delta\left(t-s-t_i^X\right)\right]d(t)\\&=\alpha^X\int_0^{\infty} ds f_c(s)\left[\int_0^{\infty} dr W(r)S^{post}(t-s)\delta\left(t-s-r-t_i^X\right)\right.\\&\left.+\int_0^{\infty} dr W(-r)S^{post}(t-s-r)\delta\left(t-s-t_i^X\right)\right]\int_0^{\infty} dr'\varepsilon_r(r')S^{post}(t-d_r-r')\\&=\alpha^X\int_0^{\infty} dr f_c\left(t-r-t_i^X\right)W(r)\int_0^{\infty} dr'\varepsilon_r(r')S^{post}\left(r+t_i^X\right)S^{post}(t-d_r-r')\\&+\alpha^X\int_0^{\infty} dr f_c\left(t-t_i^X\right)W(-r)\int_0^{\infty} dr'\varepsilon_r(r')S^{post}\left(t_i^X-r\right)S^{post}(t-d_r-r').\end{aligned}$$

We can compute the average weight change given that pattern X was presented:

$$\begin{aligned}&\left\langle\frac{d}{dt}w_i(t)\right\rangle_{E|X}=\alpha^X\int_0^{\infty} dr f_c\left(t-r-t_i^X\right)\\&W(r)\int_0^{\infty} dr'\varepsilon_r(r')\left\langle S^{post}\left(t_i^X+r\right)S^{post}(t-d_r-r')\right\rangle_{E|X}\\&+\alpha^X\int_0^{\infty} dr f_c\left(t-t_i^X\right)\\&W(-r)\int_0^{\infty} dr'\varepsilon_r(r')\left\langle S^{post}\left(t_i^X-r\right)S^{post}(t-d_r-r')\right\rangle_{E|X}.\end{aligned}$$

If we assume that f_c is approximately constant on the time scale of the learning window W, we can simplify this to

$$\begin{aligned}&\left\langle\frac{d}{dt}w_i(t)\right\rangle_{E|X}=\int_{-\infty}^{\infty} dr f_c\left(t-t_i^X\right)W(r)\\&\int_0^{\infty} dr'\varepsilon_r(r')\left\langle S^{post}\left(t_i^X+r\right)S^{post}(t-d_r-r')\right\rangle_{E|X}\alpha^X.\end{aligned}$$

For the linear Poisson neuron, we can write the auto-correlation function as

$$\begin{aligned}&\left\langle S^{post}\left(t_i^X+r\right)S^{post}(t-d_r-r')\right\rangle_{E|X}=\left[v^X\left(t_i^X+r\right)v^X(t-d_r-r')\right.\\&\left.+v^X\left(t_i^X+r\right)\delta\left(t_i^X+r-t+d_r+r'\right)\right]\\&=v^X\left(t_i^X+r\right)\left[v^X(t-d_r-r')+\delta\left(t_i^X+r-t+d_r+r'\right)\right],\end{aligned}$$

where $v^X(t)=\langle S^{post}(t)\rangle_{E|X}$ is the ensemble average rate at time t given that pattern X was presented. If an experiment for a single pattern runs over the time interval $[0,T']$, we can compute the total average weight change Δw_i^X of a trial given that pattern X was presented as

$$\begin{aligned}\Delta w_i^X&=\int_0^{T'} dt\left\langle\frac{d}{dt}w_i(t)\right\rangle_{E|X}\\&=\alpha^X\int_{-\infty}^{\infty} dr W(r)v^X\left(t_i^X+r\right)\int_0^{T'} dt f_c\left(t-t_i^X\right)\int_0^{\infty} dr'\varepsilon_r(r')\\&\quad\left[v^X(t-d_r-r')+\delta\left(t_i^X+r-t+d_r+r'\right)\right]\\&=\alpha^X\int_{-\infty}^{\infty} dr W(r)v^X\left(t_i^X+r\right)\int_0^{\infty} dr'\varepsilon_r(r')\\&\quad\left[f_c(r+d_r+r')+\int_{d_r}^{T'} dt f_c\left(t-t_i^X\right)v^X(t-d_r-r')\right]\\&\approx\alpha^X\int_{-\infty}^{\infty} dr W(r)v^X\left(t_i^X+r\right)\int_0^{\infty} dr'\varepsilon_r(r')\\&\quad\left[f_c(d_r+r')+\int_0^{T'} dt f_c\left(t-t_i^X\right)v^X(t-d_r-r')\right]\end{aligned}\tag{23}$$

By defining

$$A_i^X=\alpha^X\int_0^{\infty} dr'\varepsilon_r(r')\left[f_c(d_r+r')+\int_0^{T'} dt f_c\left(t-t_i^X\right)v^X(t-d_r-r')\right],$$

we can write Equation 23 as

$$\Delta w_i^X = \int_{-\infty}^{\infty} dr W(r) v^X\left(t_i^X + r\right) A_i^X.$$

We assume that eligibility traces and reward signals have settled to zero before a new pattern is presented. The expected weight change for the successive presentation of both patterns is therefore

$$\Delta w_i = \int_{-\infty}^{\infty} dr W(r)\left[v^P\left(t_i^P + r\right) A_i^P + v^N\left(t_i^N + r\right) A_i^N\right].$$

The equations can easily be generalized to the case where multiple input spikes per synapse are allowed and where jitter on the templates is allowed. However, the main effect of the rule can be read off the equations given here.

Common Models and Parameters of the Computer Simulations

We describe here the models and parameter values that were used in all our computer simulations. We will specify in a subsequent section the values of other parameters that had to be chosen differently in individual computer simulations, in dependence of their different setups and requirements of each computer simulation.

LIF Neuron Model

For the computer simulations LIF neurons with conductance-based synapses were used. The membrane potential $V_m(t)$ of this neuron model is given by:

$$\begin{aligned} C_m \frac{dV_m(t)}{dt} &= -\frac{V_m(t) - V_{resting}}{R_m} \\ &- \sum_{j=1}^{K_e} g_{e,j}(t)(V_m(t) - E_e) - \sum_{j=1}^{K_i} g_{i,j}(V_m(t) - E_i) - I_{noise}(t), \end{aligned} \quad (24)$$

where C_m is the membrane capacitance, R_m is the membrane resistance, $V_{resting}$ is the resting potential, and $g_{e,j}(t)$ and $g_{i,j}(t)$ are the K_e and K_i synaptic conductances from the excitatory and inhibitory synapses respectively. The constants E_e and E_i are the reversal potentials of excitatory and inhibitory synapses. I_{noise} represents the synaptic background current which the neuron receives (see below for details).

Whenever the membrane potential reaches a threshold value V_{thresh}, the neuron produces a spike, and its membrane potential is reset to the value of the reset potential V_{reset}. After a spike, there is a refractory period of length $T_{refract}$, during which the membrane potential of the neuron remains equal to the value $V_m(t) = V_{reset}$. After the refractory period $V_m(t)$ continues to change according to Equation 24.

For a given synapse, the dynamics of the synaptic conductance $g(t)$ is defined by

$$\frac{dg(t)}{dt} = -\frac{g(t)}{\tau_{syn}} + \sum_k A\left(t^{(k)} + t_{delay}\right)\delta\left(t - t^{(k)} - t_{delay}\right), \quad (25)$$

where $A(t)$ is the amplitude of the postsynaptic response (PSR) to a single presynaptic spike, which varies over time due to the inherent short-term dynamics of the synapse, and $\{t^{(k)}\}$ are the spike times of the presynaptic neuron. The conductance of the synapse decreases exponentially with time constant τ_{syn}, and increases instantaneously by amount of $A(t)$ whenever the presynaptic neuron spikes.

In all computer simulations we used the following values for the neuron and synapse parameters. The membrane resistance of the neurons was $R_m = 100\ \mathrm{M\Omega}$, the membrane capacitance $C_m = 0.3$ nF, the resting potential, reset potential and the initial value of the membrane potential had the same value of $V_{resting} = V_{reset} = V_m(0) = -70$ mV, the threshold potential was set to $V_{thresh} = -59$ mV and the refractory period $T_{refract} = 5$ ms. For the synapses we used a time constant set to $\tau_{syn} = 5$ ms, reversal potential $E_e = 0$ mV for the excitatory synapses and $E_e = -75$ mV for the inhibitory synapses. All synapses had a synaptic delay of $t_{delay} = 1$ ms.

Short-Term Dynamics of Synapses

We modeled the short-term dynamics of synapses according to the phenomenological model proposed in [37], where the amplitude $A_k = A(t_k + t_{delay})$ of the postsynaptic response for the kth spike in a spike train with inter-spike intervals $\Delta_1, \Delta_2, \ldots, \Delta_{k-1}$ is calculated with the following equations

$$\begin{aligned} A_k &= w \cdot u_k \cdot R_k \\ u_k &= U + u_{k-1}(1 - U)e^{-\Delta_{k-1}/F} \\ R_k &= 1 + (R_{k-1} - u_{k-1}R_{k-1} - 1)e^{-\Delta_{k-1}/D}, \end{aligned} \quad (26)$$

with hidden dynamic variables $u \in [0,1]$ and $R \in [0,1]$ whose initial values for the 1st spike are $u_1 = U$ and $R = 1$ (see [38] for a justification of this version of the equations, which corrects a small error in [37]). The variable w is the synaptic weight which scales the amplitudes of postsynaptic responses. If long-term plasticity is introduced, this variable is a function of time. In the simulations, for the neurons in the circuits the values for the U, D and F parameters were drawn from Gaussian distributions with mean values which depended on whether the type of presynaptic and postsynaptic neuron of the synapse is excitatory or inhibitory, and were chosen according to the data reported in [37] and [39]. The mean values of the Gaussian distributions are given in Table 2, and the standard deviation was chosen to be 50% of its mean. Negative values were replaced with values drawn from uniform distribution with a range between 0 and twice the mean value. For the simulations involving individual trained neurons, the U, D, and F parameters of these neurons were set to the values from Table 2.

We have carried out control experiments with current-based synapses that were not subject to short-term plasticity (see Figure S5, Figure S8, and Figure S9; successful control experiments with static current-based synapses were also carried out for computer simulation 1, results not shown). We found that the results of all

Table 2. Mean values of the U, D, and F parameters in the model from [37] for the short-term dynamics of synapses, depending on the type of the presynaptic and postsynaptic neuron (excitatory or inhibitory).

Source/Dest.	Exc. (U, D, F)	Inh. (U, D, F)
Exc.	0.5, 1.1, 0.02	0.25, 0.7, 0.02
Inh.	0.05, 0.125, 1.2	0.32, 0.144, 0.06

These mean values, based on experimental data from [37,39], were used in all computer simulations.

doi:10.1371/journal.pcbi.1000180.t002

our computer simulations also hold for static current-based synapses.

Model of Background Synaptic Activity

To reproduce the background synaptic input cortical neurons receive in vivo, the neurons in our models received an additional noise process as conductance input. The noise process we used is a point-conductance approximation model, described in [26]. According to [26], this noise process models the effect of a bombardment by a large number of synaptic inputs in vivo, which causes membrane potential depolarization, referred to as "high conductance" state. Furthermore, it was shown that it captures the spectral and amplitude characteristics of the input conductances of a detailed biophysical model of a neocortical pyramidal cell that was matched to intracellular recordings in cat parietal cortex in vivo. The ratio of average contributions of excitatory and inhibitory background conductances was chosen to be 5 in accordance to experimental studies during sensory responses (see [40–42]). In this model, the noisy synaptic current I_{noise} in Equation 24 is a sum of two currents:

$$I_{noise}(t) = g_e(t)(V_m(t) - E_e) + g_i(t)(V_m(t) - E_i), \qquad (27)$$

where $g_e(t)$ and $g_i(t)$ are time-dependent excitatory and inhibitory conductances. The values of the respective reversal potentials were $E_e = 0$ mV and $E_i = -75$ mV. The conductances $g_e(t)$ and $g_i(t)$ were modeled according to [26] as a one-variable stochastic process similar to an Ornstein-Uhlenbeck process:

$$\frac{dg_e(t)}{dt} = -\frac{1}{\tau_e}[g_e(t) - g_{e0}] + \sqrt{D_e}\chi_1(t)$$
$$\frac{dg_i(t)}{dt} = -\frac{1}{\tau_i}[g_i(t) - g_{i0}] + \sqrt{D_i}\chi_2(t),$$

with mean $g_{e0} = 0.012$ μS, noise-diffusion constant $D_e = 0.003$ μS and time constant $\tau_e = 2.7$ ms for the excitatory conductance, and mean $g_{i0} = 0.057$ μS, noise-diffusion constant $D_i = 0.0066$ μS, and time constant $\tau_i = 10.5$ ms for the inhibitory conductance. $\chi_1(t)$ and $\chi_2(t)$ are Gaussian white noise of zero mean and unit standard deviation.

Since these processes are Gaussian stochastic processes, they can be numerically integrated by an exact update rule:

$$g_e(t+\Delta) = g_{e0} + [g_e(t) - g_{e0}]e^{-\frac{\Delta}{\tau_e}} + A_e N_1(0,1)$$
$$g_i(t+\Delta) = g_{i0} + [g_i(t) - g_{i0}]e^{-\frac{\Delta}{\tau_i}} + A_i N_2(0,1),$$

where $N_1(0,1)$ and $N_2(0,1)$ are normal random numbers (zero mean, unit standard deviation) and A_e, A_i are amplitude coefficients given by:

$$A_e = \sqrt{\frac{D_e\tau_e}{2}\left[1 - e^{\frac{-2\Delta}{\tau_e}}\right]}$$
$$A_i = \sqrt{\frac{D_i\tau_i}{2}\left[1 - e^{\frac{-2\Delta}{\tau_i}}\right]}.$$

Reward-Modulated STDP

For the computer simulations we used the following parameters for the STDP window function $W(r)$: $A_+ = 0.01 w_{max}$, $A_-/A_+ = 1.05$, $\tau_+ = \tau_- = 30$ ms. w_{max} denotes the hard bound of the synaptic weight of the particular plastic synapse. Note that the parameter A_+ can be given arbitrary value in this plasticity rule, since it can be scaled together with the reward signal, i.e. multiplying the reward signal by some constant and dividing A_+ by the same constant results in identical time evolution of the weight changes. We have set A_+ to be 1% of the maximum synaptic weight.

We used the α-function to model the eligibility trace kernel $f_c(t)$

$$f_c(t) = \begin{cases} \frac{t}{\tau_e}e^{\frac{t}{\tau_e}} & , \quad \text{if } t > 0 \\ 0 & , \quad \text{otherwise} \end{cases}, \qquad (28)$$

where the time constant τ_e was set to $\tau_e = 0.4$ s in all computer simulations.

For computer simulations 1 and 4 we performed control experiments (see Figure S3, Figure S4, and Figure S7) with the weight-dependent synaptic update rule proposed in [22], instead of the purely additive rule in Equation 3. We used the parameters proposed in [22], i.e. $\mu = 0.4$, $\alpha = 0.11$, $\tau_+ = \tau_- = 20$ ms. The w_0 parameter was calculated according to the formula: $w_0 = \frac{1}{2}w_{max}\alpha^{1/1-\mu}$ where w_{max} is the maximum synaptic weight of the synapse. $\frac{1}{2}w_{max}$ is equal to the initial synaptic weight for the circuit neurons, or to the mean of the distribution of the initial weights for the trained neurons.

Initial Weights of Trained Neurons

The synaptic weights of excitatory synapses to the trained neurons in experiments 2–5 were initialized from a Gaussian distribution with mean $w_{max}/2$. The standard deviation was set to $w_{max}/10$ bounded within the range $[3w_{max}/10, 7w_{max}/10]$.

Software

All computer simulations were carried out with the PCSIM software package (http://www.lsm.tugraz.at/pcsim). PCSIM is a parallel simulator for biologically realistic neural networks with a fast c++ simulation core and a Python interface. It has been developed by Thomas Natschläger and Dejan Pecevski. The time step of simulation was set to 0.1 ms.

Details to Individual Computer Simulations

For all computer simulations, both for the cortical microcircuits and readout neurons, the same parameters values for the neuron and synapse models and the reward-modulated STDP rule were used, as specified in the previous section (except in computer simulation 3, where the goal was to test the theoretical predictions for different values of the parameters). Each of the computer simulations in this article modeled a specific task or experimental finding. Consequently, the dependence of the reward signal on the behavior of the system had to be modeled in a specific way for each simulation (a more detailed discussion of the reward signal can be found in the Discussion section). The parameters for that are given below in separate subsections which address the individual simulations. Furthermore, some of the remaining parameters in the experiments, i.e. the values of the synaptic weights, the number of synapses of a neuron, number of neurons in the circuit and the Ornstein-Uhlenbeck (OU) noise levels were chosen to achieve different goals depending on the particular experiment. Briefly stated, these values were tuned to achieve a certain level of firing activity in the neurons, a suitable dynamical regime of the activity in the circuits, and a specific ratio between amount of input the neurons receive from the input synapses and the input generated by the noise process.

We carried out two types of simulations: simulations of cortical microcircuits in computer simulations 1 and 5, and training of readout neurons in computer simulations 2, 3, 4, and 5. In the following we discuss these two types of simulations in more detail.

Cortical Microcircuits

The values of the initial weights of the excitatory and inhibitory synapses for the cortical microcircuits are given in Table 3. All synaptic weights were bounded in the range between 0 and twice the initial synaptic weight of the synapse.

The cortical microcircuit was composed of 4000 neurons connected randomly with connection probabilities described in Details to computer simulation 1. The initial synaptic weights of the synapses and the levels of OU noise were tuned to achieve a spontaneous firing rate of about 4.6 Hz, while maintaining an asynchronous irregular firing activity in the circuit. 50% of all neurons (randomly chosen, 50% excitatory and 50% inhibitory) received downscaled OU noise (by a factor 0.2 from the model reported in [26]), with the subtracted part substituted by additional synaptic input from the circuit. The input connection probabilities of these neurons were scaled up, so that the firing rates remain in the same range as for the other neurons. This was done in order to observe how the learning mechanisms work when most of the input conductance in the neuron comes from a larger number of input synapses which are plastic, rather than from a static noise process. The reinforced neurons were randomly chosen from this group of neurons.

We chose a smaller microcircuit, composed of 540 neurons, for the computer simulation 5 in order to be able to perform a large number of training trials. The synaptic weights in this smaller circuit were chosen (see Table 3) to achieve an appropriate level of firing activity in the circuit that is modulated by the external input. The circuit neurons had injected an Ornstein-Uhlenbeck (OU) noise multiplied by 0.4 in order to emulate the background synaptic activity in neocortical neurons in vivo, and test the learning in a more biologically realistic settings. This produced significant trial-to-trial variability in the circuit response (see Figure 10D). A lower value of the noise level could also be used without affecting the learning, whereas increasing the amount of injected noise would slowly deteriorate the information that the circuit activity maintains about the injected inputs, resulting in a decline of the learning performance.

Readout Neurons

The maximum values of the synaptic weights of readout neurons for computer simulations 2, 4, and 5, together with the number of synapses of the neurons, are given in Table 4.

The neuron in computer simulation 2 had 100 synapses. We chose 200 synapses for the neuron in computer simulation 4, in order to improve the learning performance. Such improvement of the learning performance for larger numbers of synapses is in accordance with our theoretical analysis (see Equation 17), since for learning the classification of temporal patterns the temporal variation of the voltage of the postsynaptic membrane turns out to be of critical importance (see the discussion after Equation 17). This temporal variation depends less on the shape of a single EPSP and more on the temporal pattern of presynaptic firing when the number of synapses is increased. In computer simulation 5 the readout neuron received inputs from all 432 excitatory neurons in the circuit. The synaptic weights were chosen in accordance with the number of synapses in order to achieve a firing rate suitable for the particular task, and to balance the synaptic input and the noise injections in the neurons.

For the pattern discrimination task (computer simulation 4) and the speech recognition task (computer simulation 5), the amount of noise had to be chosen to be high enough to achieve sufficient variation of the membrane potential from trial to trial near the firing threshold, and low enough so that it would not dominate the fluctuations of the membrane potential. In the experiment where the exact spike times were rewarded (computer simulation 2), the noise had a different role. As described in the Results section, there the noise effectively controls the amount of depression. If the noise (and therefore the depression) is too weak, $w^* = 0$ synapses do not converge to 0. If the noise is too strong, $w^* = w_{max}$ synapses do not converge to w_{max}. To achieve the desired learning result, the noise level should be in a range where it reduces the correlations of the synapses with $w^* = 0$ so that the depression of STDP will prevail, but at the same time is not strong enough to do the same for the other group of synapses with $w^* = w_{max}$, since they have stronger pre-before-post correlations. For our simulations, we have set the noise level to the full amount of OU noise.

Table 3. Specific parameter values for the cortical microcircuits in computer simulation 1 and 5.

Simulation No.	Neurons	p_{ee}, p_{ei}, p_{ie}, p_{ii}	$w_{exc}(0)$ [nS]	w_{inh} [nS]	C_{OU}
1	4000	0.02,0.02,0.024,0.016	10.7	211.6	1.0, 0.2
5	540	0.1	0.784	5.1	0.4

p_{conn} is the connection probability, $w_{exc}(0)$ and $w_{inh}(0)$ are the initial synaptic weights for the excitatory and inhibitory synapses respectively, and C_{OU} is the scaling factor for the Ornstein-Uhlenbeck noise injected in the neurons.
doi:10.1371/journal.pcbi.1000180.t003

Table 4. Specific parameter values for the trained (readout) neurons in computer simulation 2, 4, and 5.

Simulation No.	Num. Synapses	w_{max} [nS]	C_{OU}
2	100	11.9	1.0
4	200	5.73	0.2
5	432	2.02	0.2

w_{max} is the upper hard bound of the synaptic weights of the synapses. C_{OU} is the scaling factor for the Ornstein-Uhlenbeck noise injected in the neurons.
doi:10.1371/journal.pcbi.1000180.t004

Details to Computer Simulation 1: Model for Biofeedback Experiment

The cortical microcircuit model consisted of 4000 neurons with twenty percent of the neurons randomly chosen to be inhibitory, and the others excitatory. The connections between the neurons were created randomly, with different connectivity probabilities depending on whether the postsynaptic neuron received the full amount of OU noise, or downscaled OU noise with an additional compensatory synaptic input from the circuit. For neurons in the latter sub-population, the connection probabilities were $p_{ee} = 0.02$, $p_{ei} = 0.02$, $p_{ie} = 0.024$ and $p_{ii} = 0.016$ where the ee, ei, ie, ii indices designate the type of the presynaptic and postsynaptic neurons (e = excitatory or i = inhibitory). For the other neurons the corresponding connection probabilities were downscaled by 0.4. The resulting firing rates and correlations for both types of excitatory neurons are plotted in Figure S1 and Figure S2.

The shape of the reward kernel $\varepsilon_r(t)$ was chosen as a difference of two α-functions

$$\varepsilon_r(t) = A_r^+ \frac{t}{\tau_r^+} e^{1-\frac{t}{\tau_r^+}} - A_r^- \frac{t}{\tau_r^-} e^{1-\frac{t}{\tau_r^-}} \quad (29)$$

one positive α-pulse with a peak at 0.4 sec after the corresponding spike, and one long-tailed negative α-pulse which makes sure that the integral over the reward kernel is zero. The parameters for the reward kernel were $A_r^+ = 1.379$, $A_r^- = 0.27$, $\tau_r^+ = 0.2$ s, $\tau_r^- = 1$ s, and $d_r = 0.2$ s, which produced a peak value of the reward pulse 0.4 s after the spike that caused it.

Details to Computer Simulation 2: Learning Spike Times

We used the following function for the reward kernel $\kappa(r)$

$$\kappa(r) = \begin{cases} A_+^\kappa \left(e^{-\frac{t-t_\kappa}{\tau_1^\kappa}} - e^{-\frac{t-t_\kappa}{\tau_2^\kappa}} \right), & \text{if } t - t_\kappa \geq 0 \\ -A_-^\kappa \left(e^{-\frac{t-t_\kappa}{\tau_1^\kappa}} - e^{-\frac{t-t_\kappa}{\tau_2^\kappa}} \right), & \text{otherwise} \end{cases} \quad (30)$$

where A_+^κ and A_-^κ are positive scaling constants, τ_1^κ and τ_2^κ define the shape of the two double-exponential functions the kernel is composed of, and t_κ defines the offset of the zero-crossing from the origin. The parameter values used in our simulations were $A_+^\kappa = 0.1457$, $A_-^\kappa = -0.1442$, $\tau_1^\kappa = 30$ ms, $\tau_2^\kappa = 4$ ms and $t_\kappa = -1$ ms. The reward delay was equal to $d_r = 0.4$ s.

Details to Computer Simulation 3: Testing the Analytically Derived Conditions

We used a linear Poisson neuron model as in the theoretical analysis with static synapses and exponentially decaying postsynaptic responses $\varepsilon(s) = e^{(-s/\tau_\varepsilon)}/\tau_\varepsilon$. The neuron had 100 excitatory synapses, except in experiment #6, where we used 200 synapses. In all experiments the target neuron received additional 10 excitatory synapses with weights set to w_{max}. The input spike trains were Poisson processes with a constant rate of $r_{pre} = 6$ Hz, except in experiment # 6 where the rate was $r_{pre} = 3$ Hz. The weights of the target neuron were set to $w_i^* = w_{max}$ for $0 \leq i < 50$ and $w_i^* = 0$ for $50 \leq i < 100$.

The time constants of the reward kernel were $\tau_2^\kappa = 4$ ms, whereas τ_1^κ had different values in different experiments (reported in table 1). The value of t_κ was always set to an optimal value such that the $\varepsilon_\kappa(0) = \int_0^\infty \kappa(-s)\varepsilon(s) = 0$. The time constant τ_- of the negative part of the STDP window function $W(r)$ was set to τ_+. The reward signal was delayed by $\tau_d = 0.4$ s. The simulations were performed for varying durations of simulated biological time (see the t_{sim}-column in Table 1).

Details to Computer Simulation 4: Learning Pattern Classification

We used the reward signal from Equation 16, with an α-function for the reward kernel $\varepsilon_r(r) = \frac{e}{\tau} t e^{-t/\tau}$, and the reward delay d_r set to 300 ms. The amplitudes of the positive and negative pulses were $\alpha_P = -\alpha_N = 1.435$ and the time constant of the reward kernel was $\tau = 100$ ms.

Details to Computer Simulation 5: Training a Readout Neuron with Reward-Modulated STDP To Recognize Isolated Spoken Digits

Spike representations of speech utterances. The speech utterances were preprocessed by the cochlea model described in [43], which captures the filtering properties of the cochlea and hair cells in the human inner ear. The resulting analog signals were encoded by spikes with the BSA spike encoding algorithm described in [44]. We used the same preprocessing to generate the spikes as in [45]. The spike representations had a duration of about 400 ms and 20 input channels. The input channels were connected topographically to the cortical microcircuit model. The neurons in the circuit were split into 20 disjunct subsets of 27 neurons, and each input channel was connected to the 27 neurons in its corresponding subsets. The readout neuron was trained with 20 different spike inputs to the circuit, where 10 of them resulted from utterances of digit "one", and the other 10 resulted from utterances of digit "two" by the same speaker.

Training procedure. We performed 2000 training trials, where for each trial a spike representation of a randomly chosen utterance out of 10 utterances for one digit was injected into the circuit. The digit changed from trial to trial. Whenever the readout neuron spiked during the presentation of an utterance of digit "two", a positive pulse was generated in the reward signal, and accordingly, for utterances of digit "one", a negative pulse in the reward was generated. We used the reward signal from Equation 16. The amplitudes of the positive and negative pulses were $\alpha_P = -\alpha_N = 0.883$. The time constant of the reward kernel $\varepsilon_r(r)$ was $\tau = 100$ ms. The pulses in the reward were delayed $d_r = 300$ ms from the spikes that caused them.

Cortical microcircuit details. The cortical microcircuit model consisted of 540 neurons with twenty percent of the neurons randomly chosen to be inhibitory, and the others excitatory. The recurrent connections in the circuit were created randomly with a connection probability of 0.1. Long-term plasticity was not modeled in the circuit synapses.

The synapses for the connections from the input neurons to the circuit neurons were static, current based with axon conduction delay of 1 ms, and exponentially decaying PSR with time constant $\tau_e = 3$ ms and amplitude $w_{input} = 0.715$ nA.

Discussion

We have presented in this article analytical tools which make it possible to predict under which conditions reward-modulated STDP will achieve a given learning goal in a network of neurons. These conditions specify relationships between parameters and auxiliary functions (learning curves for STDP, eligibility traces, reward signals etc.) that are involved in the specification of the reward-modulated STDP learning rule. Although our analytical results are based on some simplifying assumptions, we have shown that they predict quite well the outcomes of computer simulations of quite complex models for cortical networks of neurons.

We have applied this learning theory for reward-modulated STDP to a number of biologically relevant learning tasks. We have shown that the biofeedback result of Fetz and Baker [17] can in principle be explained on the basis of reward-modulated STDP. The underlying credit assignment problem was extremely difficult, since the monkey brain had no direct information about the identity of the neuron whose firing rate was relevant for receiving rewards. This credit assignment problem is even more difficult from the perspective of a single synapse, and hence for the application of a local synaptic plasticity rule such as reward-modulated STDP. However our theoretical analysis (see Equations 10 and 11) has shown that the longterm evolution of synaptic weights depended only on the correlation of pairs of pre- and postsynaptic spikes with the reward signal. Therefore the firing rate of the rewarded neuron increased (for a computer simulation of a recurrent network consisting of 4000 conductance based LIF neurons with realistic background noise typical for in-vivo conditions, and 228954 synapses that exhibited data-based

short term synaptic plasticity) within a few minutes of simulated biological time, like in the experimental data of [17], whereas the firing rates of the other neurons remained invariant (see Figure 4B). We were also able to model differential reinforcement of two neurons in this way (Figure 2). These computer simulations demonstrated a remarkable stability of the network dynamics (see Figures 2A, 4A, and 5) in spite of the fact that all excitatory synapses were continuously subjected to reward-modulated STDP. In particular, the circuit remained in the asynchronous irregular firing regime, that resembles spontaneous firing activity in the cortex [9]. Other STDP-rules (without reward modulation) that maintain this firing regime have previously been exhibited in [22]. It was also reported in [17], and further examined in [46], that bursts of the reinforced neurons were often accompanied by activations of specific muscles in the biofeedback experiment by Fetz and Baker. But the relationship between bursts of the recorded neurons in precentral motor cortex and muscle activations was reported to be quite complex and often dropped out after continued reinforcement of the neuron alone. Furthermore in [46] it was shown that all neurons tested in that study could be dissociated from their correlated muscle activity by differentially reinforcing simultaneous suppression of EMG activity. These results suggest that the solution of the credit assignment problem by the monkeys (to stronger activate that neuron out of billions of neurons in their precentral gyrus that was reinforced) may have been supported by large scale exploration strategies that were associated with muscle activations. But the previously mentioned results on differential reinforcements of two nearby neurons suggest that this large scale exploration strategy had to be complemented by exploration on a finer spatial scale that is difficult to explain on the basis of muscle activations (see [19] for a detailed discussion).

Whereas this learning task focused on firing rates, we have also shown (see Figure 7) that neurons can learn via reward-modulated STDP to respond to inputs with particular spike trains, i.e., particular temporal output patterns. It has been pointed out in [27] that this is a particularly difficult learning task for reward-modulated STDP, and it was shown there that it can be accomplished with a modified STDP rule and more complex reward prediction signals without delays. We have complemented the results of [27] by deriving specific conditions (Equations 13–15) under which this learning task can be solved by the standard version of reward-modulated STDP. Extensive computer simulations have shown that these analytically derived conditions for a simpler neuron model predict also for a LIF neuron with conductance based synapses whether it is able to solve this learning task. Figure 8 shows that this learning theory for reward-modulated STDP is also able to predict quite well *how fast* a neuron can learn to produce a desired temporal output pattern. An interesting aspect of [27] is that there also the utility of third signals that provide information about changes in the expectation of reward was explored. We have considered in this article only learning scenarios where reward prediction is not possible. A logical next step will be to extend our learning theory for reward-modulated STDP to scenarios from classical reinforcement learning theory that include reward prediction.

We have also addressed the question to what extent neurons can learn via reward-modulated STDP to respond with different firing rates to different spatio-temporal presynaptic firing patterns. It had already been shown in [12] that this learning rule enables neurons to classify spatial firing patterns. We have complemented this work by deriving an analytic expression for the expected weight change in this learning scenario (see Equation 17), which clarifies to what extent a neuron can learn by reward-modulated STDP to distinguish differences in the temporal structure of presynaptic firing patterns. This theoretical analysis showed that in the extreme case, where all incoming information is encoded in the relative timing of presynaptic spikes, reward-modulated STDP is not able to produce a higher average membrane potential for selected presynaptic firing patterns, even if that would be rewarded. But it is able to increase the variance of the membrane potential, and thereby also the number of spikes of any neuron model that has (unlike the simple linear Poisson neuron) a firing threshold. The simulation results in Figure 9 confirm that in this way a LIF neuron can learn with the standard version of reward-modulated STDP to discriminate even purely temporal presynaptic firing patterns, by producing more spikes in response to one of these patterns.

A surprising feature is, that although the neuron was rewarded here only for responding with a higher firing rate to one presynaptic firing pattern P, it automatically started to respond to this pattern P with a specific temporal spike pattern, that advanced in time during training (see Figure 9A).

Finally, we have shown that a spiking neuron can be trained by reward-modulated STDP to read out information from a simulated cortical microcircuit (see Figure 10). This is insofar of interest, as previous work [31,34,47] had shown that models of generic cortical microcircuits have inherent capabilities to serve as preprocessors for such readout neurons, by combining in diverse linear and nonlinear ways information that was contained in different time segments of spike inputs to the circuit ("liquid computing model"). The classification of spoken words (that were first transformed into spike trains) had been introduced as a common benchmark task for the evaluation of different approaches towards computing with spiking neurons [31–33,45,48]. But so far all approaches that were based on learning (rather than on clever constructions) had to rely on supervised training of a simple linear readout. This gave rise to the question whether also biologically more realistic models for readout neurons can be trained through a biologically more plausible learning scenario to classify spoken words. The results of Figure 10 may be interpreted as a tentative positive answer to this question. We have demonstrated that LIF neurons with conductance based synapses (that are subject to biologically realistic short term plasticity) can learn without a supervisor through reward-modulated STDP to classify spoken digits. In contrast to the result of Figure 9, the output code that emerged here was a rate code. This can be explained through the significant in-class variance of circuit responses to different utterances of the same word (see Figure 10C and 10D). Although the LIF neuron learnt here without a supervisor to respond with different firing rates to utterances of different words by the same speaker (whereas the rate output was very similar for both words at the beginning of learning, see Figure 10E), the classification capability of these neurons has not yet reached the level of linear readouts that are trained by a supervisor (for example, speaker independent word classification could not yet be achieved in this way). Further work is needed to test whether the classification capability of LIF readout neurons can be improved through additional preprocessing in the cortical microcircuit model, through a suitable variation of the reward-modulated STDP rule, or through a different learning scenario (mimicking for example preceding developmental learning that also modifies the presynaptic circuit).

The new learning theory for reward-modulated STDP will also be useful for biological experiments that aim at the clarification of details of the biological implementation of synaptic plasticity in different parts of the brain, since it allows to make predictions which types and time courses of signals would be optimal for a

particular range of learning tasks. For each of the previously discussed learning tasks, the theoretical analysis provided conditions on the structure of the reward signal $d(t)$ which guaranteed successful learning. For example, in the biofeedback learning scenario (Figure 4), every action potential of the reinforced neuron led—after some delay—to a change of the reward signal $d(t)$. The shape of this change was defined by the reward kernel $\varepsilon(r)$. Our analysis revealed that this reward kernel can be chosen rather arbitrarily as long as the integral over the kernel is zero, and the integral over the product of the kernel and the eligibility function is positive. For another learning scenario, where the goal was that the output spike train S_j^{post} of some neuron j approximates the spike timings of some target spike train S^* (Figure 7), the reward signal has to depend on both, S_j^{post} and S^*. The dependence of the reward signal on these spike timings was defined by a reward kernel $\kappa(r)$. Our analysis showed that the reward kernel has to be chosen for this task so that the synapses receive positive rewards if the postsynaptic neuron fires close to the time of a spike in the target spike train S^* or somewhat later, and negative rewards when an output spike occurs in the order of ten milliseconds too early. In the pattern discrimination task of Figure 9 each postsynaptic action potential was followed—after some delay—by a change of the reward signal which depended on the pattern presented. Our theoretical analysis predicted that this learning task can be solved if the integrals A_i^P and A_i^N defined by Equation 18 are such that $A_i^P > 0$ and $A_i^N \approx -A_i^P$. Again, this constraints are fulfilled for a large class of reward kernels, and a natural choice is to use a non-negative reward kernel ε_r. There are currently no data available on the shape of reward kernels in biological neural systems. The previous sketched theoretical analysis makes specific prediction for the shape of reward kernels (depending on the type of learning task in which a biological neural system is involved) which can potentially be tested through biological experiments.

An interesting general aspect of the learning theory that we have presented in this article is that it requires substantial trial-to-trial variability in the neural circuit, which is often viewed as "noise" of imperfect biological implementations of theoretically ideal circuits of neurons. This learning theory for reward-modulated STDP suggests that the main functional role of noise is to maintain a suitable level of spontaneous firing (since if a neuron does not fire, it cannot find out whether this will be rewarded), which should vary from trial to trial in order to explore which firing patterns are rewarded (It had been shown in [31,34,47] that such highly variable circuit activity is compatible with a stable performance of linear readouts). On the other hand if a neuron fires primarily on the basis of a noise current that is directly injected into that neuron, and not on the basis of presynaptic activity, then STDP does not have the required effect on the synaptic connections to this neuron (see Figure S6). This perspective opens the door for subsequent studies that compare for concrete biological learning tasks the theoretically derived optimal amount and distribution of trial-to-trial variability with corresponding experimental data.

Related Work

The theoretical analysis of this model is directly applicable to the learning rule considered in [12]. There, the network behavior of reward-modulated STDP was also studied some situations different from the ones in this article. The computer simulations of [12] operate apparently in a different dynamic regime, where LTD dominates LTP in the STDP-rule, and most weights (except those that are actively increased through reward-modulated STDP) have values close to 0 (see Figure 1b and 1d in [12], and compare with Figure 5 in this article). This setup is likely to require for successful learning a larger dominance of pre-before-post over post-before-pre pairs than the one shown in Figure 4E. Furthermore, whereas a very low spontaneous firing rate of 1 Hz was required in [12], computer simulation 1 shows that reinforcement learning is also feasible at spontaneous firing rates which correspond to those reported in [17] (the preceding theoretical analysis had already suggested that the success of the model does not depend on particularly low firing rates). The articles [15] and [13] investigate variations of reward-modulated STDP rules that do not employ learning curves for STDP that are based on experimental data, but modified curves that arise in the context of a very interesting top-down theoretical approach (distributed reinforcement learning [14]). The authors of [16] arrive at similar learning rules in a supervised scenario which can be reinterpreted in the context of reinforcement learning. We expect that a similar theory as we have presented in this article for the more commonly discussed version of STDP can also be applied to their modified STDP rules, thereby making it possible to predict under which conditions their learning rules will succeed. Another reward based learning rule for spiking neurons was recently presented in [49]. This rule exploits correlations of a reward signal with noisy perturbations of the neuronal membrane conductance in order to optimize some objective function. One crucial assumption of this approach is that the synaptic plasticity mechanism "knows" which contributions to the membrane potential arise from synaptic inputs, and which contributions are due to internal noise. Such explicit knowledge of the noise signal is not needed in the reward-modulated STDP rule of [12], which we have considered in this article. The price one has to pay for this potential gain in biological realism is a reduced generality of the learning capabilities. While the learning rule in [49] approximates gradient ascent on the objective function, this cannot be stated for reward-modulated STDP at present. Timing-based pattern discrimination with a spiking neuron, as discussed in the section "Pattern discrimination with reward-modulated STDP" of this article, was recently tackled in [50]. The authors proposed the tempotron learning rule, which increases the peak membrane voltage for one class of input patterns (if no spike occurred in response to the input pattern) while decreasing the peak membrane voltage for another class of input patterns (if a spike occurred in response to the pattern). The main difference between this learning rule and reward-modulated STDP is that the tempotron learning rule is sensitive to the peak membrane voltage, whereas reward-modulated STDP is sensitive to local fluctuations of the membrane voltage. Since the time of the maximal membrane voltage has to be determined for each pattern by the synaptic plasticity mechanism, the basic tempotron rule is perhaps not biologically realistic. Therefore, an approximate and potentially biologically more realistic learning rule was proposed in [50], where plasticity following error trials is induced at synapse i only if the voltage within the postsynaptic integration time after their activation exceeds a plasticity threshold κ. One potential problem of this rule is the plasticity threshold κ, since a good choice of this parameter strongly depends on the mean membrane voltage after input spikes. This problem is circumvented by reward-modulated STDP, which considers instead the local change in the membrane voltage. Further work is needed to compare the advantages and disadvantages of these different approaches.

Conclusion

Reward-modulated STDP is a very promising candidate for a synaptic plasticity rule that is able to orchestrate local synaptic modifications in such a way that particular functional properties of larger networks of neurons can be achieved and maintained (we refer to [12] and [27] for discussion of potential biological

implementations of this plasticity rule). We have provided in this article analytical tools which make it possible to evaluate this rule and variations of this rule not just through computer simulations, but through theoretical analysis. In particular we have shown that successful learning is only possible if certain relationships hold between the parameters that are involved. Some of these predicted relationships can be tested through biological experiments.

Provided that these relationships are satisfied, reward-modulated STDP turns out to be a powerful rule that can achieve self-organization of synaptic weights in large recurrent networks of neurons. In particular, it enables us to explain seemingly inexplicable experimental data on biofeedback in monkeys. In addition reward-modulated STDP enables neurons to distinguish complex firing patterns of presynaptic neurons, even for data-based standard forms of STDP, and without the need for a supervisor that tells the neuron when it should spike. Furthermore reward-modulated STDP requires substantial spontaneous activity and trial-to-trial variability in order to support successful learning, thereby providing a functional explanation for these ubiquitous features of cortical networks of neurons. In fact, not only spontaneous activity but also STDP itself may be seen in this context as a mechanism that supports the exploration of different firing chains within a recurrent network, until a solution is found that is rewarded because it supports a successful computational function of the network.

Supporting Information

Figure S1 Variations of Figure 5B–D for those excitatory neurons which receive the full amount of Ornstein-Uhlenbeck noise. (B) The distribution of the firing rates of these neurons remains unchanged during the simulation. The colors of the curves and the corresponding intervals are as follows: red (300–360 sec), green (600–660 sec), blue (900–960 sec), magenta (1140–1200 sec). (C) Cross-correlogram of the spiking activity of these neurons, averaged over 200 pairs of neurons and over 60 s, with a bin size of 0.2 ms, for the period between 300 and 360 seconds of simulation time. It is calculated as the cross-covariance divided by the square root of the product of variances. (D) As in (C), but for the last 60 seconds of the simulation. The correlation statistics in the circuit is stable during learning.
Found at: doi:10.1371/journal.pcbi.1000180.s001 (0.06 MB PDF)

Figure S2 Variations of Figure 5B–D for those excitatory neurons which receive a reduced amount of Ornstein-Uhlenbeck noise, but receive more synaptic inputs from other neurons. (B) The distribution of the firing rates of these neurons remains unchanged during the simulation. The colors of the curves and the corresponding intervals are as follows: red (300–360 sec), green (600–660 sec), blue (900–960 sec), magenta (1140–1200 sec). (C) Cross-correlogram of the spiking activity in the circuit, averaged over 200 pairs of these neurons and over 60 s, with a bin size of 0.2 ms, for the period between 300 and 360 seconds of simulation time. It is calculated as the cross-covariance divided by the square root of the product of variances. (D) As in (C), but for the last 60 seconds of the simulation. The correlation statistics in the circuit is stable during learning.
Found at: doi:10.1371/journal.pcbi.1000180.s002 (0.06 MB PDF)

Figure S3 Variation of Figure 4 from computer simulation 1 with results from a simulation where the weight-dependent version of STDP proposed in [22] was used. This STDP rule is defined by the following equations: $\Delta w_+ = \lambda w_0^{1-\mu} w^{\mu} e^{-|\Delta t|/\tau_+}$ and $\Delta w_- = \lambda \alpha w e^{-|\Delta t|/\tau_-}$. We used the parameters proposed in [36], i.e. $\mu = 0.4$, $\alpha = 0.11$, $\tau_+ = \tau_- = 20$ ms, $\lambda = 0.1$ and $w_0 = 272.6$ pS. The w_0 parameter was calculated according to the formula: $w_0 = \frac{1}{2} w_{max} \alpha^{\frac{1}{1-\mu}}$ where w_{max} is the maximum synaptic weight of the synapse. The amplitude parameters A_r^+, A_r^- for the reward kernel were set to $A_r^+ = 1.104$ and $A_r^- = 0.221$. All other parameter values were the same as in computer simulation 1.
Found at: doi:10.1371/journal.pcbi.1000180.s003 (0.09 MB PDF)

Figure S4 Variation of Figure 5 for the weight-dependent STDP rule from [22] (as in Figure S3).
Found at: doi:10.1371/journal.pcbi.1000180.s004 (0.06 MB PDF)

Figure S5 Variation of Figure 7 (i.e., of computer simulation 2) for a simulation where we used current-based synapses without short-term plasticity. The post-synaptic response had an exponentially decaying form $\varepsilon(s) = e^{-s/\tau_\varepsilon}/\tau_\varepsilon$, with $\tau_\varepsilon = 5$ ms. The value of the maximum synaptic weight was $w_{max} = 32.9$ pA. All other parameter values were the same as in computer simulation 2.
Found at: doi:10.1371/journal.pcbi.1000180.s005 (0.17 MB PDF)

Figure S6 Dependence of the learning performance on the noise level in computer simulation 2. The angular error (defined as the angle between the weight vector $\mathbf{w}$ of the trained neuron at the end of the simulation and the weight vector $\mathbf{w}^*$ of the neuron μ^*) is taken as measure for the learning performance, and plotted for 9 simulations with different noise levels that are given on the X axis (in term of multiples of the noise level chosen for Figure 7). All other parameters values were the same as in computer simulation 2. The figure shows that the learning performance declines both for too little and for too much noise.
Found at: doi:10.1371/journal.pcbi.1000180.s006 (0.02 MB PDF)

Figure S7 Variation of Figure 9 (i.e., of computer simulation 4) with the weight-dependent STDP rule proposed in [22]. This rule is defined by the following equations: $\Delta w_+ = \lambda w_0^{1-\mu} w^{\mu} e^{-|\Delta t|/\tau_+}$ and $\Delta w_- = \lambda \alpha w e^{-|\Delta t|/\tau_-}$. We used the parameters proposed in [22], i.e. $\mu = 0.4$, $\alpha = 0.11$, $\tau_+ = \tau_- = 20$ ms, $\lambda = 0.1$ and $w_0 = 72.4$ pS. The w_0 parameter was calculated according to the formula: $w_0 = \frac{1}{2} w_{max} \alpha^{\frac{1}{1-\mu}}$ where w_{max} is the maximum synaptic weight of the synapse. The amplitude parameters of the reward kernel were set to $\alpha_P = -\alpha_N = 1.401$. All other parameter values were the same as in computer simulation 4. The variance of the membrane potential increased for pattern P from 2.35 $(\text{mV})^2$ to 3.66 $(\text{mV})^2$ (C), and decreased for pattern N (D), from 2.27 $(\text{mV})^2$ to 1.54 $(\text{mV})^2$.
Found at: doi:10.1371/journal.pcbi.1000180.s007 (0.31 MB PDF)

Figure S8 Variation of Figure 9 for a simulation where we used current-based synapses without short-term plasticity. The post-synaptic response had an exponentially decaying form $\varepsilon(s) = e^{-s/\tau_\varepsilon}/\tau_\varepsilon$, with $\tau_\varepsilon = 5$ ms. The value of the maximum synaptic weight was $w_{max} = 106.2$ pA All other parameter values were the same as in computer simulation 4. The variance of the membrane potential increased for pattern P from 2.84 $(\text{mV})^2$ to 5.89 $(\text{mV})^2$ (C), and decreased for pattern N (D), from 2.57 $(\text{mV})^2$ to 1.22 $(\text{mV})^2$.
Found at: doi:10.1371/journal.pcbi.1000180.s008 (0.31 MB PDF)

Figure S9 Variation of Figure 10 (i.e., of computer simulation 5) for a simulation where we used current-based synapses without short-term plasticity. The post-synaptic response had an exponentially decaying form $\varepsilon(s) = e^{-s/\tau_\varepsilon}/\tau_\varepsilon$, with $\tau_\varepsilon = 5$ ms. The synaptic weights of the excitatory and inhibitory synapses in the cortical microcircuit were set to $w_{exc} = 65.4$ pA and $w_{inh} = 238$ pA respectively. The maximum synaptic weight of the synapses to the readout neuron was $w_{max} = 54.3$ pA. All other parameter values were the same as in computer simulation 5.
Found at: doi:10.1371/journal.pcbi.1000180.s009 (0.27 MB PDF)

Figure S10 Spike encodings of 10 utterances of digit "one" by one speaker with the Lyon cochlea model [43], which were used as circuit inputs for computer simulation 5.
Found at: doi:10.1371/journal.pcbi.1000180.s010 (0.05 MB PDF)

Figure S11 Spike encodings of 10 utterances of digit "two" by one speaker with the Lyon cochlea model [43], which were used as circuit inputs for computer simulation 5.
Found at: doi:10.1371/journal.pcbi.1000180.s011 (0.05 MB PDF)

Acknowledgments

We would like to thank Markus Diesmann, Eberhard Fetz, Razvan Florian, Yves Fregnac, Wulfram Gerstner, Nikos Logothetis, Abigail Morrison, Matthias Munk, Gordon Pipa, and Dan Shulz for helpful discussions. In addition we would like to thank Malcolm Slaney for providing a MATLAB implementation of the cochlea model from [43], as well as Benjamin Schrauwen, David Verstraeten, Michiel D'Haene, and Stefan Klampfl for additional code that we used in our speech classification task (computer simulation 5).

Author Contributions

Conceived and designed the experiments: RL DP WM. Wrote the paper: RL DP WM.

References

1. Abbott LF, Nelson SB (2000) Synaptic plasticity: taming the beast. Nat Neurosci 3: 1178–1183.
2. Jacob V, Brasier D, Erchova I, Feldman D, Shulz DE (2007) Spike timing-dependent synaptic depression in the in vivo barrel cortex of the rat. J Neuroscience 27: 1271–1284.
3. Bailey CH, Giustetto M, Huang YY, Hawkins RD, Kandel ER (2000) Is heterosynaptic modulation essential for stabilizing Hebbian plasticity and memory? Nat Rev Neurosci 1: 11–20.
4. Gu Q (2002) Neuromodulatory transmitter systems in the cortex and their role in cortical plasticity. Neuroscience 111: 815–835.
5. Schultz W (2007) Behavioral dopamine signals. Trends Neurosci 30: 203–210.
6. Reynolds JN, Hyland BI, Wickens JR (2001) A cellular mechanism of reward-related learning. Nature 413: 67–70.
7. Reynolds JN, Wickens JR (2002) Dopamine-dependent plasticity of corticostriatal synapses. Neural Netw 15: 507–521.
8. Bao S, Chan VT, Merzenich MM (2001) Cortical remodelling induced by activity of ventral tegmental dopamine neurons. Nature 412: 79–83.
9. Shulz DE, Sosnik R, Ego V, Haidarliu S, Ahissar E (2000) A neuronal analogue of state-dependent learning. Nature 403: 549–553.
10. Thiel CM, Friston KJ, Dolan RJ (2002) Cholinergic modulation of experience-dependent plasticity in human auditory cortex. Neuron 35: 567–574.
11. Shulz DE, Ego-Stengel V, Ahissar E (2003) Acetylcholine-dependent potentiation of temporal frequency representation in the barrel cortex does not depend on response magnitude during conditioning. J Physiol Paris 97: 431–439.
12. Izhikevich EM (2007) Solving the distal reward problem through linkage of STDP and dopamine signaling. Cereb Cortex 17: 2443–2452.
13. Florian RV (2007) Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. Neural Comput 6: 1468–1502.
14. Baxter J, Bartlett PL (1999) Direct gradient-based reinforcement learning: I. Gradient estimation algorithms. Technical report. Research School of Information Sciences and Engineering, Australian National University.
15. Baras D, Meir R (2007) Reinforcement learning, spike-time-dependent plasticity, and the bcm rule. Neural Comput 19: 2245–2279.
16. Pfister JP, Toyoizumi T, Barber D, Gerstner W (2006) Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. Neural Comput 18: 1318–1348.
17. Fetz EE, Baker MA (1973) Operantly conditioned patterns of precentral unit activity and correlated responses in adjacent cells and contralateral muscles. J Neurophysiol 36: 179–204.
18. Fetz EE (1969) Operant conditioning of cortical unit activity. Science 163: 955–958.
19. Fetz EE (2007) Volitional control of neural activity: implications for brain–computer interfaces. J Physiol 579: 571–579.
20. Brunel N (2000) Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. J Physiol Paris 94: 445–463.
21. Gerstner W, Kistler WM (2002) Spiking Neuron Models. Cambridge, UK: Cambridge University Press.
22. Morrison A, Aertsen A, Diesmann M (2007) Spike-timing-dependent plasticity in balanced random networks. Neural Comput 19: 1437–1467.
23. Bi G, Poo M (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J Neurosci 18: 10464–10472.
24. Song S, Miller KD, Abbott LF (2000) Competitive hebbian learning through spiketiming dependent synaptic plasticity. Nat Neurosci 3: 919–926.
25. Kempter R, Gerstner W, van Hemmen JL (2001) Intrinsic stabilization of output rates by spike-based hebbian learning. Neural Comput 13: 2709–2741.
26. Destexhe A, Rudolph M, Fellous JM, Sejnowski TJ (2001) Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. Neuroscience 107: 13–24.
27. Farries MA, Fairhall AL (2007) Reinforcement learning with modulated spike timing-dependent synaptic plasticity. J Neurophysiol 98: 3648–3665.
28. Stevens CF, Zador AM (1998) Input synchrony and the irregular firing of cortical neurons. Nat Neurosci 1: 210–217.
29. Mainen Z, Sejnowski T (1995) Reliability of spike timing in neocortical neurons. Science 268: 1503–1505.
30. Silberberg G, Bethge M, Markram H, Pawelzik K, Tsodyks M (2004) Dynamics of population rate codes in ensembles of neocortical neurons. J Neurophysiol 91: 704–709.
31. Maass W, Natschläger T, Markram H (2002) Real-time computing without stable states: a new framework for neural computation based on perturbations. Neural Comput 14: 2531–2560.
32. Maass W, Natschläger T, Markram H (2004) Fading memory and kernel properties of generic cortical microcircuit models. J Physiol Paris 98: 315–330.
33. Destexhe A, Marder E (2004) Plasticity in single neuron and circuit computations. Nature 431: 789–795.
34. Maass W, Joshi P, Sontag ED (2007) Computational aspects of feedback in neural circuits. PLoS Comput Biol 3: e165. doi:10.1371/journal.pcbi.0020165.
35. Nikolić D, Haeusler S, Singer W, Maass W (2007) Temporal dynamics of information content carried by neurons in the primary visual cortex. In: Proceedings of NIPS 2006, Advances in Neural Information Processing Systems. Volume 19. Cambridge (Massachusetts): MIT Press. pp 1041–1048.
36. Kempter R, Gerstner W, van Hemmen JL (1999) Hebbian learning and spiking neurons. Phys Rev E 59: 4498–4514.
37. Markram H, Wang Y, Tsodyks M (1998) Differential signaling via the same axon of neocortical pyramidal neurons. Proc Natl Acad Sci U S A 95: 5323–5328.
38. Maass W, Markram H (2002) Synapses as dynamic memory buffers. Neural Networks 15: 155–161.
39. Gupta A, Wang Y, Markram H (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. Science 287: 273–278.
40. Borg-Graham I,J, Monier C, Frégnac Y (1998) Visual input evokes transient and strong shunting inhibition in visual cortical neurons. Nature 393: 369–373.
41. Hirsch JA, Alonso JM, Reid RC, Martinez LM (1998) Synaptic integration in striate cortical simple cells. J Neurosci 18: 9517–9528.
42. Anderson J, Lampl I, Reichova I, Carandini M, Ferster D (2000) Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. Nature Neurosci 3: 617–621.
43. Lyon R (1982) A computational model of filtering, detection, and compression in the cochlea. In: Proceedings of IEEE International Conference on ICASSP. pp 1282–1285.
44. Schrauwen B, Campenhout JV (2003) BSA, a fast and accurate spike train encoding scheme. In: Proceedings of the International Joint Conference on Neural Networks. Volume 4. pp 2825–2830.
45. Verstraeten D, Schrauwen B, Stroobandt D, Campenhout JV (2005) Isolated word recognition with the liquid state machine: a case study. Inf Process Lett 95: 521–528.
46. Fetz EE, Finocchio DV (1975) Correlations between activity of motor cortex cells and arm muscles during operantly conditioned response patterns. Exp Brain Res 23: 217–240.
47. Häusler S, Maass W (2007) A statistical analysis of information processing properties of lamina-specific cortical microcircuit models. Cereb Cortex 17: 149–162.
48. Hopfield JJ, Brody CD (2001) What is a moment? Transient synchrony as a collective mechanism for spatio-temporal integration. Proc Natl Acad Sci U S A 98: 1282–1287.
49. Fiete IR, Seung HS (2006) Gradient learning in spiking neural networks by dynamic perturbation of conductances. Phys Rev Lett 97: 048104.
50. Gütig R, Sompolinsky H (2006) The tempotron: a neuron that learns spike timingbased decisions. Nat Neurosci 9: 420–428.

2.5 Appendix 5

M. Helias, S. Rotter, M.-O. Gewaltig, and Markus Diesmann, Structural Plasticity controlled by calcium based correlation detection Frontiers in Computational Neuroscience December 2008 2: 7

frontiers in
COMPUTATIONAL NEUROSCIENCE

ORIGINAL RESEARCH ARTICLE
published: 24 December 2008
doi: 10.3389/neuro.10.007.2008

Structural plasticity controlled by calcium based correlation detection

Moritz Helias[1,2]*, Stefan Rotter[1,2], Marc-Oliver Gewaltig[1,3] and Markus Diesmann[1,4,5]

[1] *Bernstein Center for Computational Neuroscience, Freiburg, Germany*
[2] *Computational Neuroscience, Faculty of Biology, Albert-Ludwig University, Freiburg, Germany*
[3] *Honda Research Institute Europe, Offenbach, Germany*
[4] *Theoretical Neuroscience Group, RIKEN Brain Science Institute, Wako City, Japan*
[5] *Brain and Neural Systems Team, Computational Science Research Program, RIKEN, Wako City, Japan*

Edited by:
Nicolas Brunel, CNRS, France

Reviewed by:
Wulfram Gerstner, EPFL, Switzerland
Nicolas Brunel, CNRS, France

****Correspondence:***
Moritz Helias, Bernstein Center for Computational Neuroscience, Hansastrasse 9a, 79104 Freiburg, Germany.
e-mail: helias@bccn.uni-freiburg.de

Hebbian learning in cortical networks during development and adulthood relies on the presence of a mechanism to detect correlation between the presynaptic and the postsynaptic spiking activity. Recently, the calcium concentration in spines was experimentally shown to be a correlation sensitive signal with the necessary properties: it is confined to the spine volume, it depends on the relative timing of pre- and postsynaptic action potentials, and it is independent of the spine's location along the dendrite. NMDA receptors are a candidate mediator for the correlation dependent calcium signal. Here, we present a quantitative model of correlation detection in synapses based on the calcium influx through NMDA receptors under realistic conditions of irregular pre- and postsynaptic spiking activity with pairwise correlation. Our analytical framework captures the interaction of the learning rule and the correlation dynamics of the neurons. We find that a simple thresholding mechanism can act as a sensitive and reliable correlation detector at physiological firing rates. Furthermore, the mechanism is sensitive to correlation among afferent synapses by cooperation and competition. In our model this mechanism controls synapse formation and elimination. We explain how synapse elimination leads to firing rate homeostasis and show that the connectivity structure is shaped by the correlations between neighboring inputs.

Keywords: structural plasticity, correlation detection, synaptic death, synaptogenesis, silent synapses, STDP, Calcium calmodulin dependent kinase II (CaMKII), synaptic cooperativity

INTRODUCTION

The connectivity structure of the cortex was found to be surprisingly dynamic *in vitro* and *in vivo* (Bonhoeffer and Yuste, 2002). Synapse formation and elimination exhibit a marked dependence on spiking activity, where higher activity promotes synapse formation (Le Be and Markram, 2006). Based on geometric considerations, Stepanyants et al. (2002), Chklovskii et al. (2004) and Stepanyants et al. (2007) suggest as a basic design principle of the cortex the potential of any pair of neurons to form a connection on small length scales of a few hundred micrometers together with an activity dependent selection mechanism. This idea is supported by direct observation of spines approaching presynaptic partners in a promiscuous manner (Kalisman et al., 2005). Of these structurally possible (potential) synapses, only a small fraction (0.12–0.34) is actually realized (Stepanyants et al., 2002), and transitions from potential to actual synapses are observed *in vitro* at rates of up to 1.2 per cent per hour during increased spiking activity (Le Be and Markram, 2006). These newly formed, immature synapses are weaker than mature ones. Synapse pruning mostly affects weak, but already mature synapses. The relation of synaptic strength to synapse formation and synaptic death indicates that long term potentiation (LTP) and synapse formation may be controlled by similar mechanisms and the same may hold for long term depression (LTD) and synapse pruning. The observation of increased connectivity after prolonged spiking activity also prompts for a synapse pruning mechanism, which restores the typical synapse density. In support of this hypothesis, not only in the developing cortex, but also in adult monkeys *in vivo*, synaptic boutons emerge and disappear with rates of 7 per cent per week (Stettler et al., 2006).

The functional role of synapse formation and elimination is not well understood, but theoretical studies have exposed the benefits of the capability of structural remodeling (Chklovskii et al., 2004; Stepanyants et al., 2002). Due to the fact that the number of potential presynaptic partners of a neuron easily exceeds the actual number at any point in time by an order of magnitude, each synapse carries three to four bits of structural information. In the framework of spiking associative memories this beneficial functional role of structural plasticity can be exploited to increase the memory capacity (Knoblauch, 2006; Knoblauch et al., 2007). Therefore, synaptic formation and elimination is a candidate process for long term information storage in cortical networks during development and adulthood alike.

The microscopic mechanisms leading to the formation of new synapses and to the elimination of existing ones have not yet been completely revealed. However, there is some evidence, that newly formed synapses are created in an intermediate, silent state (Cohen-Cory, 2002; Kalisman et al., 2005). These frequently encountered silent synapses lack AMPA receptors but have NMDA receptors (Atwood and Wojtowicz, 2004). They bear a high potential for remodeling the neural circuit, since they can easily be converted into

active synapses. The most probable mechanism is the translocation of AMPA receptors into the postsynaptic density (PSD). This is also observed (Cohen-Cory, 2002; Shi et al., 1999) during LTP for which NMDA receptor activation is a necessary precondition.

There is strong indication that synapse formation and synapse pruning are controlled in a similar way as LTP and LTD (Lüscher et al., 2000). Furthermore, Le Be and Markram (2006) found that the same antagonists prevent LTP and synapse formation, and conclude that the underlying mechanisms may be similar. Calcium entering the postsynaptic site through NMDA receptors is a probable messenger causing synapse maturation (Atwood and Wojtowicz, 2004). For the induction of LTP and LTD, which in parts come about by modulating the number of postsynaptic AMPA receptors, Cormier et al. (2001) showed that both depend in a threshold like fashion on the intracellular calcium peak amplitude: higher amplitudes lead to LTP and lower ones to LTD.

Spike pairing experiments showed the calcium signal in spines to depend on the correlated activation of the pre- and postsynaptic site: Nevian and Sakmann (2004, 2006) found a marked dependency on the relative timing and independence on the position of the spine along the dendrite, making the NMDA receptor mediated calcium level in a spine an attractive candidate substrate to convey information about the correlation between presynaptic and postsynaptic activity.

For the calcium signal in a spine to control synaptic maturation or synaptic death it must activate a downstream signaling cascade. Calcium/calmodulin dependent kinase II (CaMKII) is a calcium activated kinase, which is crucial for the LTP of a synapse. In its activated state it can phosphorylate several structures among them AMPA receptors which respond with increased conductivity. There is also evidence, that CaMKII is involved in the insertion process of new AMPA receptors into the PSD. This causes LTP or turns a silent synapse (only having NMDA receptors) into an active one having AMPA and NMDA receptors. There is recent evidence from detailed biophysical modeling studies (Graupner and Brunel, 2007) that the activity dependent calcium influx can activate CaMKII in a bistable fashion and hereby explain spike timing dependent synaptic plasticity (STDP). For a recent review of phenomenological models of STDP see Morrison et al. (2008).

CaMKII forms holoenzymes of two ring molecules consisting of six subunits each. A subunit can either be active or inactive. Transitions between the inactive and the active state are triggered by calcium signals of different amplitudes. Short and weak calcium signals typically lead to activation of a single subunit by binding calcium or calmodulin to it. After the calcium level has dropped, unbinding of calcium and hence deactivation occurs within 0.1–0.2 s. At larger calcium concentrations an active subunit of the molecule (to which calcium is already bound) can phosphorylate the neighboring subunit. This only requires one additional calcium molecule to bind to the second subunit to expose its phosphorylation site. So phosphorylation can propagate along the ring and the molecule remains active even after calcium has returned to the resting level. At resting calcium concentrations, protein phosphatase 1 (PP1) can dephosphorylate an active subunit, but a neighboring active site can immediately rephosphorylate it again. This regenerating effect explains the long time scales of several minutes for the deactivation of groups of active CaMKII molecules. Even longer time scales of hours of persistent activity are found at resting calcium concentrations in the special chemical environment of the PSD, where the concentration of PP1 is low compared to the number of CaMKII subunits. For a comprehensive review see Lisman et al. (2002). Detailed biophysical simulations (Miller et al., 2005) have confirmed bistability between an active and an inactive state of whole populations of approximately 20 holoenzymes. This effect is due to saturation of the phosphatase in the active state. The study found life times of both states in the range of 100 years. The time scale of the attractor dynamics is on the order of tens of minutes, but for strong fluctuations of the calcium signal the bistability vanishes. Calcium can not only activate CaMKII (either directly or via calmodulin), but also protein phosphatases like calcineurin and protein phosphatase 1, which dephosphorylate CaMKII. These phosphatases have a higher affinity to calcium than CaMKII. Therefore, they become active at lower calcium concentrations and counteract the phosphorylation of CaMKII. This is in agreement with the finding that LTD in CA1 dendrites can be induced if the calcium concentration is below 180 nM, while LTP requires it to exceed 540 nM (Cormier et al., 2001). For a review see Cavazzini et al. (2005).

To our knowledge, previous models for structural plasticity either used simplified neuron models (Butz et al., 2008; Dammasch et al., 1986) or plasticity rules depending on the firing rate alone and not taking into account formation and death of individual synapses (van Ooyen et al., 1995). Consequently correlation dependent structure formation is outside the scope of these works. In this modeling study, we investigate how the biologically known pathways outlined above interplay to achieve a mechanism capable of detecting correlation between the presynaptic and the postsynaptic spiking activity. We focus on the calcium control hypothesis: the calcium signal mediated by NMDA receptors is the beginning of the signaling cascade. The main features of NMDA receptors entering our model are: (1) Their fast binding to glutamate followed by slow unbinding. (2) The quasi-instantaneous removal of the magnesium block upon postsynaptic depolarization to open the channel. In these assumptions, our model is similar to previous work by Shouval et al. (2002) on a mathematical model to explain spike timing dependent plasticity (STDP) based on the properties of NMDA receptors and the calcium control hypothesis. In contrast to their work, we assume the postsynaptic depolarization by the backpropagating action potential (bpAP) to be a short event.

The next stage of the signaling pathway in our model is a calcium controlled bistable effector molecule, like e.g. CaMKII. The important properties for our model are: (1) The long time constants of sustained activation of each individual molecule by high calcium concentrations (bistability). We are interested in the regime, where calcium fluctuations dominate the activation dynamics and the slower attractor dynamics causing bistability of the whole population of molecules is negligible. (2) The ability of the kinase to influence synaptic plasticity via AMPA receptor insertion. We assume a minimum amount of the kinase to be necessary for promoting synapses from silent to functional and we assume that there a minimal amount of active kinase is required to prevent synapse death. (3) Two disjoint ranges of calcium concentration that control the transitions between the inactive and the active state of each molecule. (4) The relatively low number of molecules, which makes a

statistical description essential. We derive a model for correlation detection based on this pathway and investigate its dynamics under realistic conditions of irregular spike trains. We show that the model represents a viable mechanism to sense correlation between the pre- and postsynaptic activity. Controlling synaptic pruning, it can implement a firing rate homeostasis. Cohen-Cory (2002) already suspected, that activity-dependent remodeling selectively stabilizes coactive incoming synapses and destabilizes others. We demonstrate that such cooperation and competition between synapses naturally emerges from the microscopic model and that a neuron can learn the correlations between neighboring inputs.

In the section "Spike Time Dependence of Postsynaptic Calcium Concentration" we explain the origin of spike timing dependence of the postsynaptic calcium concentration, mention the main findings of recent imaging experiments, and develop a model for the peak amplitude of the postsynaptic calcium signal. We show it to qualitatively reproduce the experimental findings. In the section "Ca^{2+} Transients Caused by Correlated Irregular Spiking" we show that for irregular spiking activity this model predicts a distinct dependency of the observed postsynaptic calcium signal on the correlation between the presynaptic and the postsynaptic spiking events. The section titled "A Counter for Correlated Events" derives a biologically motivated model of a mechanism to "count" correlated events and therefore to assess the degree of correlation between the presynaptic and the postsynaptic activity. The section titled "Rate Homeostasis by Synaptic Pruning" shows that controlling synaptic pruning by this correlation measure can act as a rate regulation for the postsynaptic neuron at low rates. In the section "Cooperation and Competition by Spatial Input Correlation" we demonstrate that cooperation and competition between synapses depends on the correlation between neighboring inputs and that a synaptic pruning process manifests these input correlations in the resulting network structure. The last section discusses our results.

All simulations were carried out with the NEST simulation software (Gewaltig and Diesmann, 2007) using the computationally efficient implementation of synaptic maturation and death provided in the section "Algorithmic Implementation of Synapse Maturation and Synapse Death" in Appendix. Preliminary results have been presented in abstract form (Helias et al., 2007).

SPIKE TIME DEPENDENCE OF POSTSYNAPTIC CALCIUM CONCENTRATION

In this section we show, that the calcium peak amplitude in a spine in good approximation depends exponentially on the temporal difference of the presynaptic and the postsynaptic spiking and that the calcium influx is largest, if the presynaptic cell fires shortly before the postsynaptic cell. This makes the calcium signal an appropriate candidate carrier of information on causal correlation.

Figure 1A illustrates the situation at a synapse subject to a spike pairing protocol. The principal mechanism causing the dependence of the postsynaptic calcium signal can readily be understood. With each presynaptic spike a small amount of glutamate is released, increasing the glutamate concentration ρ_{glu} at time t_{glu}. Glutamate binds to a fraction of the postsynaptic NMDA receptors. There is experimental evidence (Mainen et al., 1999) that only a fraction $n(t)$ of the NMDA receptors bind to glutamate. $n(t)$ reaches its maximum n_0 after a rise time of $\tau_{rise,nmda} \simeq 5 - 10$ ms. We choose this maximum obtained for a single presynaptic release to be the unit of n. After the glutamate concentration $\rho_{glu}(t)$ has decayed back to its resting value, the receptors unbind and $n(t)$ decays back to 0. As long as the postsynaptic spine is not depolarized, the NMDA receptors are blocked by magnesium and thus have a low conductance (Jahr and Stevens, 1990a,b). Arrival of a postsynaptic back-propagating action potential (called bpAP in the following) depolarizes the spine (indicated as V_{spine} in **Figure 1A**), removes this block and Ca^{2+} can flow into the spine. The amount of NMDA receptors opening at time t_{BP} therefore equals

$$n(t_{BP}) = n(t_{BP} - t_{glu}).$$

This is illustrated in **Figure 1B**. Thus the NMDA receptor mediated conductivity depends on the relative timing $\Delta t_{syn} = t_{BP} - t_{glu}$.

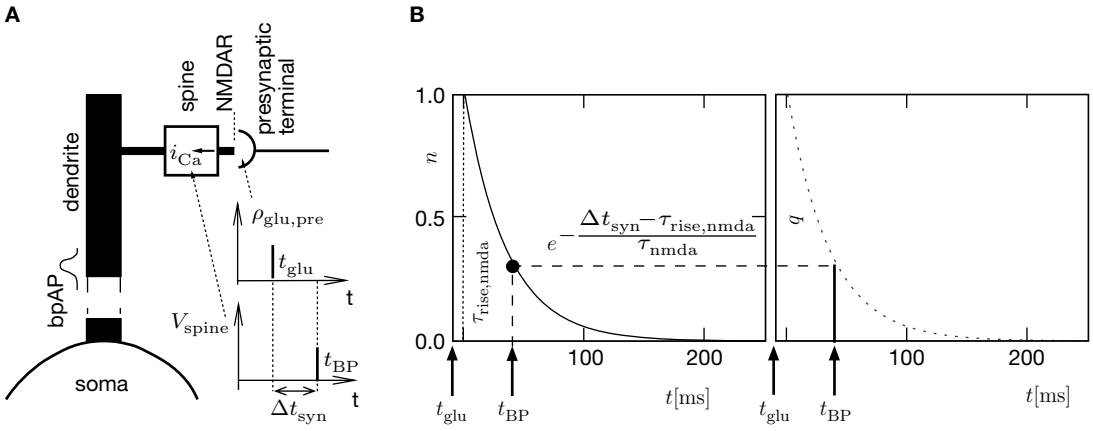

FIGURE 1 | (A) Sketch of an experiment measuring the calcium concentration in spines during a spike pairing protocol. The calcium concentration inside the spine is measured while varying the temporal difference Δt_{syn} between the onset t_{glu} of the glutamate concentration and the arrival of a postsynaptic bpAP t_{BP} at the spine. **(B)** Rise in glutamate concentration at t_{glu} in the synaptic cleft causes NMDA receptors in the spine to bind to glutamate. The fraction $n(t)$ of glutamate bound NMDARs jumps to its maximum ($n_0 = 1$ for simplicity) after time $\tau_{rise,nmda}$. As long as the spine is not depolarized, the NMDA pore is blocked by Mg^{2+}. A postsynaptic spikes causes a bpAP to arrive at the spine at t_{BP} and to unblock the currently glutamate bound NMDA receptors $n(t_{BP})$, so Ca^{2+} can enter the spine. The total amount of Ca^{2+} influx q is proportional to $n(t_{BP})$ and therefore depends on the relative timing Δt_{syn}.

This model assumes instantaneous removal of the Mg^{2+} block. However, the detailed mechanism is more complicated. Removal of the Mg^{2+} block can happen at several time scales; a very fast one of $\simeq$100 µs (Jahr and Stevens, 1990b) and several recently found longer time scales of a few 100 ms (Kampa et al., 2004). The longer time scales were found to effectively narrow the time window of substantial NMDA conductance (Kampa et al., 2004). Here, we do not intend to capture the NMDA receptor kinetics in full detail, but rather to construct a functional yet quantitative model where the experimental results constrain the model parameters. Calcium imaging studies on spines (Nevian and Sakmann, 2004) show, that the timing dependence of the Ca^{2+} peak amplitude can be fitted for positive Δt_{syn} to a function with a rise time of $\tau_{rise,nmda}$ followed by a exponential decay with a single time constant of τ_{nmda}. These studies typically measure the timing $\Delta t = t_{post} - t_{pre}$ between the spiking in the presynaptic neuron's soma t_{pre} and the spiking in the postsynaptic neuron's soma t_{post} (see Morrison et al., 2008, for the definition of delays in models of synaptic plasticity). With d_{BP} being the delay for a bpAP and d_{glu} being the delay between presynaptic spike and glutamate release, the timing difference at the synapse is

$$\begin{aligned}\Delta t_{syn} &= t_{BP} - t_{glu} \\ &= t_{post} + d_{BP} - (t_{pre} + d_{glu}) \\ &= \Delta t + d_{BP} - d_{glu}.\end{aligned}$$

In this work, we neglect the continuous binding process and instead assume that the number n of glutamate bound NMDA receptors jumps to a positive value after the rise time $\tau_{rise,nmda}$, like

$$n(t) = n_0 H(t - t_{glu} - \tau_{rise,nmda}) e^{-\frac{t - t_{glu} - \tau_{rise,nmda}}{\tau_{nmda}}}, \quad (1)$$

where $H(t)$ denotes the Heaviside function. Assuming the bpAP to be a point event, the total Ca^{2+} influx into the spine

$$\begin{aligned}q &= \int_{-\infty}^{\infty} q_0 \delta(t - t_{BP}) n_0 H(t - t_{glu} - \tau_{rise,nmda}) e^{-\frac{t - t_{glu} - \tau_{rise,nmda}}{\tau_{nmda}}} dt \\ &= q_0 n_0 H(\Delta t_{syn} - \tau_{rise,nmda}) e^{-\frac{\Delta t_{syn} - \tau_{rise,nmda}}{\tau_{nmda}}} \\ &= q_0 n_0 H(\Delta t - \overbrace{(\tau_{rise,nmda} + d_{glu} - d_{BP})}^{\tau_{rise}}) e^{-\frac{\Delta t - (\tau_{rise,nmda} + d_{glu} - d_{BP})}{\tau_{nmda}}} \\ &= q_0 n_0 H(\Delta t - \tau_{rise}) e^{-\frac{\Delta t - \tau_{rise}}{\tau_{nmda}}} \quad (2)\end{aligned}$$

depends exponentially on the timing Δt where $i_{Ca} = q_0 \delta(t - t_{BP})$ is the calcium current through a single open NMDA channel during a bpAP. Our expression for q defines the effective rise time τ_{rise} and suggests to measure q in units of $q_0 n_0$, i.e. we set $q_0 n_0 = 1$. Nevian and Sakmann (2004) measured an effective rise time $\tau_{rise} = 10$ ms and an exponential decay with a time constant $\tau_{nmda} = 32$ ms.

The postsynaptic depolarization due to the AMPA receptor activation only causes a small NMDA conductivity, as shown in measurements of the NMDA conductivity for voltage patterns caused by spike pairing experiments (Kampa et al., 2004) and also directly by observing that the calcium transient in spines in spike pairing experiments is not decreased significantly by blocking AMPA receptors (Nevian and Sakmann, 2004). In our model, we neglect the influence of AMPA mediated depolarizations on the calcium signal.

Assuming the bpAP to be a point event is obviously an approximation. As well as the assumption, that the glutamate binding state jumps from 0 to 1 at $t = t_{glu} + \tau_{rise,nmda}$ instead of showing a continuous increase on a time scale of 5–10 ms (as found experimentally by Kampa et al., 2004). However, for timings $\Delta t_{syn} \geq \tau_{rise,nmda}$, it is easy to see that this does not qualitatively change the total Ca^{2+} influx (assuming $i_{Ca}(t < 0) = 0$, because

$$\begin{aligned}q(t_{BP}) &= \int_{-\infty}^{\infty} i_{Ca}(t - t_{BP}) n(t - t_{glu}) dt \\ &= \int_{0}^{\infty} i_{Ca}(t) n(t - t_{glu} + t_{BP}) dt \\ &= \int_{0}^{\infty} i_{Ca}(t) n(t + \Delta t_{syn}) dt \\ &= \int_{0}^{\infty} i_{Ca}(t) e^{-\frac{t}{\tau_{nmda}}} dt \; e^{-\frac{\Delta t_{syn} - \tau_{rise,nmda}}{\tau_{nmda}}} \\ &= q_0' e^{-\frac{\Delta t_{syn} - \tau_{rise,nmda}}{\tau_{nmda}}}.\end{aligned}$$

As before, q depends exponentially on the relative timing Δt_{syn}. The exact shape of the bpAP only enters the timing independent factor q_0'. For $\Delta t_{syn} < \tau_{rise,nmda}$, the integral depends both on the shape of the resulting $i_{Ca}(t)$, which in turn depends on the actual shape of the bpAP, and on the rising flank of the NMDA channel activation. A typical trace of $i_{Ca}(t)$ has a half duration of $\simeq$5 ms and the rise time of NMDA activation is of the same order of magnitude (5–10 ms). Hence, the slope of $q(\Delta t_{syn})$ in this range is large compared to the following exponential decay with τ_{nmda}, as confirmed experimentally (Nevian and Sakmann, 2004). The step like approximation (Eq. 2) therefore seems adequate. Our choice to describe the bpAP to be a point event has the consequence that in this spike pairing experiment there is no calcium influx for larger negative timings $\Delta t_{syn} - \tau_{rise,nmda} < 0$. Taking into account a finite short decay time for the postsynaptic depolarization as in Shouval et al. (2002) would lead to a small calcium influx also for the post- before presynaptic timing.

The Ca^{2+} influx q leads to a transient signal which in good approximation decays in an exponential fashion (Nevian and Sakmann, 2004, 2006; Waters et al., 2003) with a decay time of 20–200 ms (reviewed in Cavazzini et al., 2005). We therefore assume the calcium peak amplitude to depend linearly on the amount q of calcium influx.

Ca^{2+} TRANSIENTS CAUSED BY CORRELATED IRREGULAR SPIKING

So far we studied the model for the case of a spike pairing protocol, where a presynaptic action potential is paired with a postsynaptic action potential. In this section we investigate how much information about the correlation between presynaptic and postsynaptic events is contained in the Ca^{2+} signal if the spiking activity is irregular. To this end, we simulate the presynaptic and the postsynaptic spiking as Poisson processes with rates ν_i and ν_o, respectively. Both processes share a fraction of spikes that appear in the spike trains with a fixed temporal distance Δt. In the following we call these events "pair events". **Figure 2A** illustrates these processes. The strength of the temporal correlation is given by the correlation coefficient ε, which is the conditional probability of a pair event (i.e. seeing the temporally correlated presynaptic partner spike at $t - \Delta t$), given a postsynaptic spike at time t.

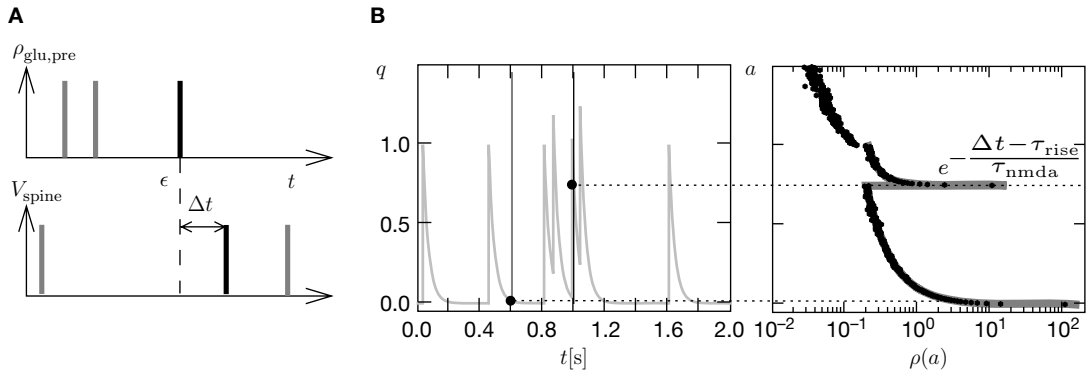

FIGURE 2 | (A) Presynaptic glutamate release treated as point events $\rho_{\mathrm{glu,pre}}(t)$. The back-propagating action potentials cause a postsynaptic depolarization $V_{\mathrm{spine}}(t)$, taken to be point events as well. Pre- and postsynaptic events occur with Poisson statistics at ν = 5 Hz, but both sequences contain a fraction of correlated events (black bars), which have a fixed temporal distance Δt = 15 ms from each other. ε = 0.1 is the probability of observing a presynaptic partner, given a postsynaptic spike. The gray bars denote postsynaptic spikes, which have no correlated partner among the presynaptic events. **(B)** The left plot shows the shot noise (light gray) produced by the presynaptic events. A postsynaptic spike t_{post} samples the shot noise at the time of occurrence, $a := q(t_{\mathrm{post}})$. Postsynaptic spikes that belong to a pair (i.e. which are preceded by a presynaptic event Δt before) are indicated by black bars. They sample $q(t)$ at high values and thus produce the peak in the histogram (right panel) at $e^{-(\Delta t-\tau_{\mathrm{rise}})/\tau_{\mathrm{nmda}}} \simeq 0.73$ (with τ_{rise} = 5 ms). The black dots result from a realization of the stochastic process ($T = 10^5$ s) using Scientific Python (Jones et al., 2001), the gray curve illustrates (Eq. 5). The peak amplitude is proportional to the correlation coefficient ε = 0.1. Uncorrelated postsynaptic events cause the peak at low values of a.

Let A be the set of presynaptic spike times and B the set of postsynaptic spike times within a finite interval $[0, T]$. The effects of the presynaptic spikes add up linearly in our model. Therefore, the number of glutamate bound NMDA receptors found at time t is

$$n(t) = n_0 \sum_{\{t_{\mathrm{pre}} \in A | t_{\mathrm{pre}} < t - \tau_{\mathrm{rise}}\}} e^{-\frac{t - t_{\mathrm{pre}} - \tau_{\mathrm{rise}}}{\tau_{\mathrm{nmda}}}} \tag{3}$$

and the total Ca^{2+} influx at the point in time t_{post} of a postsynaptic spike is accordingly

$$q_A(t_{\mathrm{post}}) = \sum_{\{t_{\mathrm{pre}} \in A | t_{\mathrm{pre}} < t_{\mathrm{post}} - \tau_{\mathrm{rise}}\}} e^{-\frac{t_{\mathrm{post}} - t_{\mathrm{pre}} - \tau_{\mathrm{rise}}}{\tau_{\mathrm{nmda}}}},$$

where we again set $q_0 n_0$ to unity. For different realizations of the presynaptic and postsynaptic spiking the elements of the set $a_{A,B} = \{q_A(t_{\mathrm{post}}) \mid t_{\mathrm{post}} \in B\}$ are random variables, since A and B are random sets. We therefore need to calculate the probability density function of q_A and show that it depends on the correlation ε between pre- and postsynaptic events. Initially assume $\varepsilon = 0$, i.e. A and B to be two independent sets of Poisson points. In this case the points in time t_{post} at which $q_A(t)$ is sampled are randomly and uniformly drawn from the interval $[0, T]$. Consequently the elements $a \in a_{A,B}$ occur according to the probability density function of $q_A(t)$, where t is a randomly and uniformly drawn point in time $t \in [0, T]$. This is the amplitude distribution $\rho_0(q)$ of a shot noise with an exponential kernel and can be calculated recursively (Gilbert and Pollak, 1960) as

$$\rho_0(q) = \begin{cases} C\, q^{r-1} & \text{for } 0 < q \leq 1 \\ q^{r-1}(C - r\int_1^q \rho_0(x-1)x^{-r}\mathrm{d}x) & \text{for } q > 1 \end{cases}$$

with $r = \nu_i \tau_{\mathrm{nmda}}$

$$C = \frac{e^{-r\gamma}}{\Gamma(r)}$$

$$\gamma = 0.577215665, \tag{4}$$

where γ is Euler's constant. In the case $0 < \varepsilon \leq 1$, however, we have to distinguish two cases: Given a postsynaptic spike, with probability $1 - \varepsilon$ this spike is not part of a pair event. The postsynaptic event is uncorrelated with respect to the presynaptic events and thus samples $q(t)$ at a random point in $[0, T]$ as discussed above. So the contribution of this event to the probability density $\rho(a)$ is $(1 - \varepsilon)\rho_0(a)$. With probability ε the postsynaptic spike is part of a pair event. In this case, we know that there is a presynaptic spike at time $t_{\mathrm{post}} - \Delta t$ of which we still see the deterministic contribution $\Delta a = e^{-(\Delta t - \tau_{\mathrm{rise}})/\tau_{\mathrm{nmda}}}$ to a. Since we assume the presynaptic spikes to have Poisson statistics, the existence of the presynaptic spike at $t_{\mathrm{post}} - \Delta t$ does not influence the statistics of the other presynaptic spikes. The latter produce a shot noise background obeying the distribution $\rho_0(a')$: the probability density to observe the value $a = a' + \Delta a$ equals the probability density $\rho_0(a')$. Hence the contribution in this case is $\varepsilon\rho_0(a - \Delta a)$. Considering both types of postsynaptic spikes we arrive at

$$\rho(a) = \varepsilon\rho_0(a - \Delta a) + (1 - \varepsilon)\rho_0(a) \tag{5}$$

with $\Delta a = e^{-(\Delta t - \tau_{\mathrm{rise}})/\tau_{\mathrm{nmda}}}$.

Figure 2B illustrates this result. The postsynaptic events belonging to spike pairs (black bars) sample the shot noise signal at high values and cause the peak in the histogram around 0.73. Its height scales linearly with the correlation coefficient ε. The uncorrelated gray events cause a background manifested in the histogram by the large peak at low values. Its amplitude scales proportional to $1 - \varepsilon$.

In conclusion, the information about correlation between the presynaptic and the postsynaptic events enters the probability distribution $\rho(a)$ and is therefore exhibited in the amplitude distribution of Ca^{2+} transients in the postsynaptic spine: correlated postsynaptic events which come in close temporal proximity to a presynaptic spike produce a high Ca^{2+} influx.

A COUNTER FOR CORRELATED EVENTS

In this section we devise a model of a postsynaptically realized counter for correlated events. The model is based on CaMKII, the most probable downstream signaling protein. CaMKII is an example of a bistable effector protein, whose transitions between its active and its inactive state are triggered by distinct calcium concentrations. Since the number of CaMKII molecules in a spine is low a statistical description is essential. **Figure 3A** shows a schematic drawing of the possible transitions of a single ring molecule in our model and their dependence on the calcium concentration: If the calcium influx exceeds the highest threshold Θ_h, the number x of active CaMKII molecules is increased. We call this a plus-event. If the influx is between Θ_l and Θ_h, the amount of active molecules remains unaffected. This region is often referred to as "no man's land" (Cormier et al., 2001). For a low calcium influx between Θ_b and Θ_l, the number of active molecules decreases. We call this a minus-event. If the influx is below Θ_b, the number of active molecules remains the same. Previous theoretical work (Shouval et al., 2002) assumes a similar dependence of the synaptic weight change on the calcium concentration, but does not take into account the intermediate concentration, where no plasticity occurs. After a high calcium event, the concentration eventually drops to levels between Θ_b and Θ_l, where it can activate the phosphatase and thus deactivates CaMKII molecules. So the increase of active molecules in our model caused by a high Ca^{2+} event is understood to be the net effect of activated minus deactivated molecules.

Cormier et al. (2001) measured the thresholds for constant calcium concentrations. It is unclear, whether these values also hold for transient calcium influx through NMDA receptors. Furthermore, for the spike pairing protocol, only relative calcium concentrations have been measured. We therefore pursue a more phenomenological approach and choose Θ_h such that a plus-event occurs, if the temporal distance between the presynaptic and the postsynaptic event Δt is in the range $0 \le \Delta t = t_{\text{post}} - t_{\text{pre}} \le \Delta t_+ = 20$ ms, consistent with the potentiation window of STDP (Bi and Poo, 1998). Furthermore, Bi and Poo (1998) showed that the transition from LTD to LTP occurs in a relatively narrow time window symmetric around $\Delta t = 0$. Since in the present work we aim at a functional model we choose the effective rise time to be $\tau_{\text{rise}} = 0$ ms, in order for the coincidence window to start at $\Delta t = 0$. Θ_h is then given by $\Theta_h = e^{-\Delta t_+ / \tau_{\text{nmda}}}$ (see **Figure 4**). Subsequently, we use the absolute values of Θ_h and Θ_l of Cormier et al. (2001) to infer from the ratios $\frac{\Theta_l}{\Theta_h} \simeq 0.75$ and $\frac{\Theta_b}{\Theta_h} \simeq 0.3$ the appropriate values of Θ_l and Θ_b for our condition.

Note that this choice of thresholds has the consequence for the spike pairing protocol that for $0 \le \Delta t \le \Delta t_+$ we obtain activation of CaMKII molecules, but we do not obtain deactivation for the reversed timing $\Delta t < 0$ as suggested by the experimentally observed LTD window of STDP. In the section "Sensitivity of Results to Model Assumptions" we discuss an extension of our model to incorporate the LTD window for negative relative timing and we argue that our results are invariant under this modification.

Knowing the probability distribution $\rho(a)$ of the calcium peak amplitudes a given a postsynaptic depolarization, we can calculate the probability of occurrence $p_- = P(\Theta_b \le a < \Theta_l \mid \text{postsynaptic spike})$ of a minus-event and $p_+ = P(a \ge \Theta_h \mid \text{postsynaptic spike})$ of a plus-event. In order to detect pairs of correlated pre- and postsynaptic events, their relative timing must satisfy $\Delta t \le \Delta t_+$. This is equivalent to the condition $\Delta a := e^{-\Delta t / \tau_{\text{nmda}}} \ge \Theta_h$, meaning, that correlated events cause an influx larger than Θ_h (compare **Figure 3B**).

Using Eqs. 5 and 4 for the probability of a plus-event we arrive at

$$\begin{aligned} p_+(\varepsilon) &= \int_{\Theta_h}^{\infty} \rho(a)\,\mathrm{d}a \\ &= (1-\varepsilon)\int_{\Theta_h}^{\infty} \rho_0(a)\,\mathrm{d}a + \varepsilon \underbrace{\int_{\Theta_h}^{\infty} \rho_0(a-\Delta a)\,\mathrm{d}a}_{=1 \text{ for } \Theta_h < \Delta a} \\ &= (1-\varepsilon)\left(1 - \int_0^{\Theta_h} \rho_0(a)\,\mathrm{d}a\right) + \varepsilon \end{aligned}$$

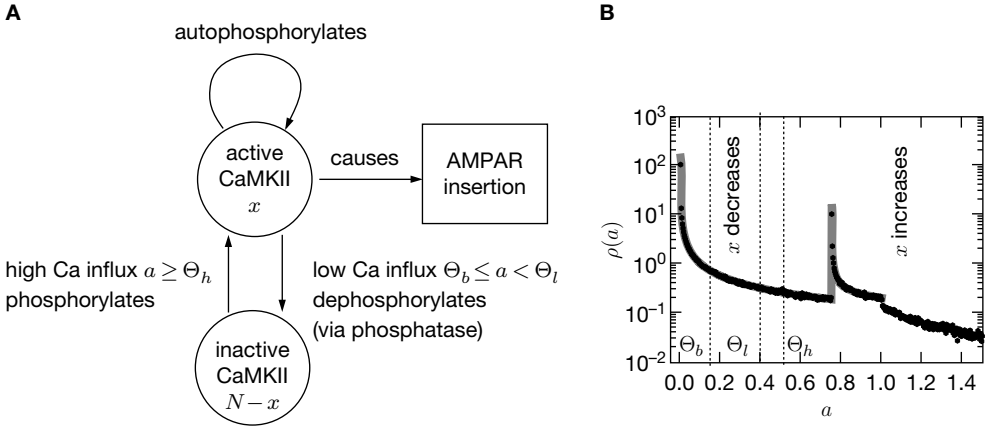

FIGURE 3 | (A) CaMKII molecules exhibit bistability: An inactive state and a highly phosphorylated, active state. Transitions between the two states are triggered by Ca^{2+} influx. If the influx exceeds a threshold Θ_h, the molecule switches from inactive to active. Low Ca^{2+} influx between Θ_b and Θ_l deactivates the molecule. **(B)** At each postsynaptic spike, Ca^{2+} can enter the spine and can change the amount x of active CaMKII molecules. The probability density function $\rho(a)$ of the Ca^{2+} influx at the points in time of a postsynaptic spike $a = q(t_{\text{post}})$ determines the direction of the effect on x (same data as in **Figure 2B**). According to this effect, it can be divided into different regions: For $\Theta_b \le a < \Theta_l$ the number of active molecules is reduced. This event occurs with probability p_- (Eq. 7). For $a \ge \Theta_h$ the number is increased, occurring which probability p_+ (Eq. 6). In the regions between Θ_l and Θ_h and below Θ_b there is no influence on x.

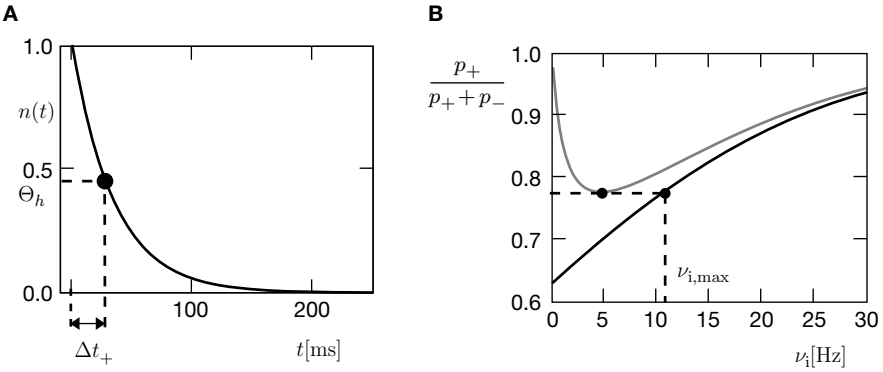

FIGURE 4 | (A) Spike pairs with timing differences $\Delta t \leq \Delta t_+ = 20$ ms produce plus-events. This determines the threshold $\Theta_h = e^{-\Delta t_+ / \tau_{nmda}}$. **(B)** Dependence of $s(\varepsilon) = \frac{p_+(\varepsilon)}{p_+(\varepsilon)+p_-(\varepsilon)}$ on the presynaptic rate ν_i for two different correlation coefficients $\varepsilon = 0$ (black) and for $\varepsilon = 0.3$ (gray). The maximum spike rate $\nu_{i,\max}$ at which the detector can discriminate these two correlation coefficients results from the condition that $s(\varepsilon) > s(0)$ for all $\nu_i \leq \nu_{i,\max}$.

$$= 1 + (\varepsilon - 1)\frac{C}{r}\Theta_h^r. \tag{6}$$

Here r and C are given as in Eq. 4. An analogous calculation, which is valid under the same assumption $\Theta_b < \Theta_l < \Theta_h \leq \Delta a$ yields

$$\begin{aligned} p_-(\varepsilon) &= \int_{\Theta_b}^{\Theta_l} \rho(a)\,\mathrm{d}a \\ &= (1-\varepsilon)\int_{\Theta_b}^{\Theta_l} \rho_0(a)\,\mathrm{d}a + \underbrace{\int_{\Theta_b}^{\Theta_l} \rho_0(a-\Delta a)\,\mathrm{d}a}_{=0 \text{ for } \Theta_b<\Theta_l<\Theta_h\leq\Delta a} \\ &= (1-\varepsilon)\frac{C}{r}\left(\Theta_l^r - \Theta_b^r\right). \end{aligned} \tag{7}$$

Each PSD of a spine contains a number of CaMKII molecules, which was found to be $N = 80$ in a typical PSD on average (Chen et al., 2005). We are only interested in the behavior of the two stable states of a molecule: On long time scales, a molecule can only be in the fully activated or in the completely inactive state. Let the number of active molecules be x and assume that the postsynaptic neuron spikes with a rate ν_o. Then x is a random variable and we are interested in the equilibrium distribution and its dependence on the correlation between the presynaptic and postsynaptic spike train. Each postsynaptic spike may lead to a plus-event with probability $p_+(\varepsilon)$. In such an event, each of the $N - x$ inactive molecules can become active. In our model, this transition happens independently for each particle with probability p. Given x active molecules, the number of molecules per time being activated is $\nu_o p_+ p(N - x)$. Analogously, the number of molecules per time being deactivated is $\nu_o p_- q x$. Here, q is the probability of a minus-event to deactivate a particular molecules. This scenario is sketched in **Figure 5**. In equilibrium, both currents must compensate, leading to the expected number of active molecules

$$\langle x \rangle_{\mathrm{eq}} = N\frac{p_+ p}{p_+ p + p_- q}. \tag{8}$$

Thus $\langle x \rangle_{\mathrm{eq}}$ depends on the relative probabilities $p_+(\varepsilon)$ for a plus-event and $p_-(\varepsilon)$ for a minus-event given by Eqs. 6 and 7. **Figure 5B** shows the probability distribution for the number of active molecules for an ensemble of synapses, where the presynaptic and the postsynaptic activity are correlated Poisson processes. The higher the correlation coefficient ε between presynaptic and postsynaptic activity is, the more the distribution is shifted to the right. For a derivation of the first and second moment of the probability distribution see section "Probability Distribution for the Number of Active CaMKII Molecules" in Appendix.

The amount of active CaMKII molecules is the signal that can trigger downstream processes in the postsynaptic spine. Synapse maturation caused by insertion of new AMPA receptors into the PSD is such a process. In our model we assume that the process requires the presence of a certain minimal amount X_m of active CaMKII molecules. The probability of a synapse to maturate is therefore the probability that the number x of active molecules exceeds the threshold X_m. In a premature synapse, the initial Ca^{2+} concentration is low and hence the amount of active CaMKII is low as well. We would like to know the mean time needed by the signal x, starting at $x = 0$, to cross the threshold X_m for the first time. This is the mean first passage time problem. We approximate the mean first passage rate by the decay rate of the slowest decaying eigenvector of the CaMKII distribution in the section "Mean First Passage Time Problem for the Number of CaMKII Molecules" in Appendix. This solution is plotted for different thresholds in **Figure 6**.

We treat synapse pruning analogously. In a mature synapse the initial amount of active CaMKII is already beyond the threshold X_m. Due to pre- and postsynaptic activity, the amount x of active CaMKII may decrease or increase, depending on the rates and the pre- and postsynaptic correlation. If eventually x falls below the minimal amount X_d, the synapse dies. We choose $X_d < X_m$ for two reasons. First, once the amount of CaMKII is high, autophosphorylation will act regeneratively (Miller et al., 2005), making the decrease of x harder. We do not model this dynamics explicitly, but rather incorporate its effect in our choice of $X_d < X_m$. Secondly, experiments by Le Be and Markram (2006) suggest, that there is a "period of grace" for newly formed synapses during which they are not pruned, preventing many synapses to be created and destructed in vain. Our choice is a possible implementation.

RATE HOMEOSTASIS BY SYNAPTIC PRUNING

In this section we employ the correlation detection mechanism to control synaptic pruning and demonstrate its capability to regulate

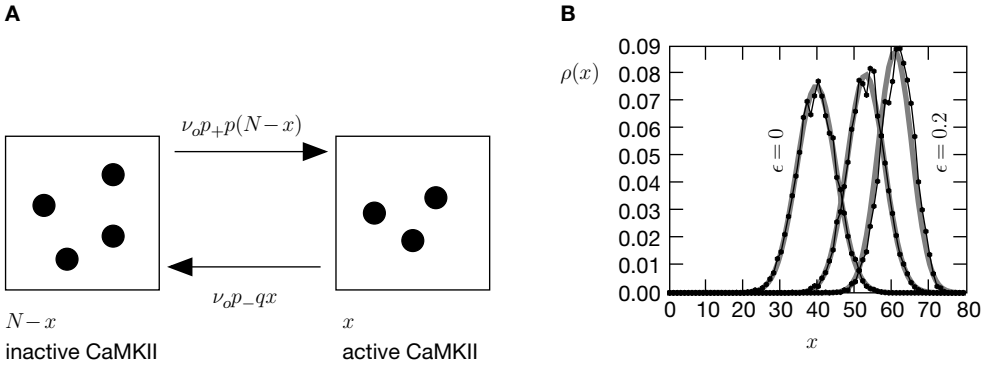

FIGURE 5 | (A) Finite reservoir model of CaMKII. Each of the $N = 80$ molecules can either be active or inactive. Given a high calcium influx $a \geq \Theta_h$ (which occurs with probability p_+), each CaMKII molecule has the probability $p = 0.01$ to be activated. Given a low calcium influx $\Theta_b \leq a < \Theta_l$ (occurring with probability p_-), the probability of an active molecule to be deactivated is $q = 0.01$. Presynaptic and postsynaptic events are Poisson with rates $\nu_i = \nu_o = 5$ Hz. The transition rates between the inactive and the active state are therefore $\nu_o p_+ p(N - x)$ and $\nu_o p_- qx$, respectively. **(B)** Equilibrium distribution of the number of active molecules x in an ensemble of synapses for different correlation coefficients $\varepsilon = \{0, 0.1, 0.2\}$. The activation rate increases with the correlation ε, whereas the deactivation rate decreases, shifting $\rho(x)$ to the right. The black curves shows simulation results (temporal resolution 0.1 ms), the gray curves are Gaussians parameterized by Eqs. 8 and 17.

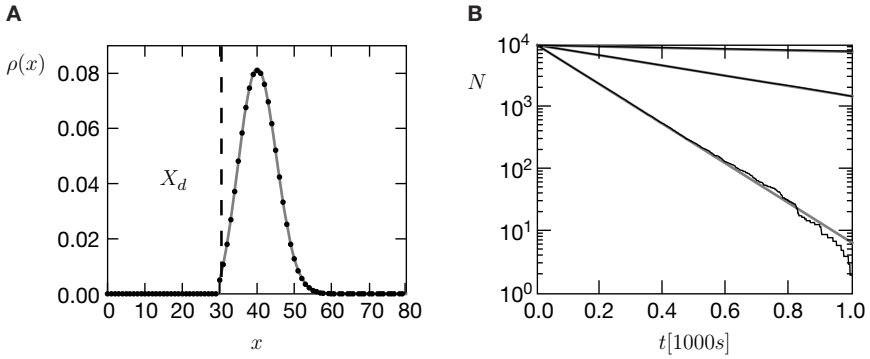

FIGURE 6 | (A) Distribution $\rho(x) = \left\langle \frac{\rho(x,t)}{\int_0^N \rho(x',t)\mathrm{d}x'} \right\rangle_t$ of the number of active molecules in an ensemble of synapses averaged over time. Synapses with $x < X_d$ are pruned. This threshold acts as an absorbing boundary. Semi-analytic expression (Eq. 19) (gray), simulation of N = 10 000 synapses subject to presynaptic Poisson activity of $\nu_i = 5$ Hz and postsynaptically the spiking activity of an integrate and fire neuron with $\nu_o = 9$ Hz (black). **(B)** The number of surviving synapses as a function of time. Simulation (black) and analytical expectation value (gray) (Eq. 29). The death rate (slope) corresponds to the eigenvalue η_1. Different rates are obtained for the thresholds $X_d = 25$, 30, 35, where the fastest decay belongs to the highest threshold $X_d = 35$. (The parameters of the integrate and fire neuron are: membrane time constant $\tau_m = 20$ ms, threshold $V_{th} = 15.0$ mV,, reset potential $V_r = 0$ mV. It receives excitatory Poisson input of $\nu_{ext} = 35400$ Hz from synapses with weight $w = 0.05$ mV and inhibitory Poisson inputs of rate $\nu_{inh} = 5600$ Hz from synapses of weight $-gw = -0.2$ mV.)

spike rate towards a state of low rate. Before doing so we generalize the correlation model studied in previous sections to the more realistic scenario of a temporally extended correlation between the presynaptic and the postsynaptic activity.

The correlation detection depends on the probabilities p_+ for a plus-event (high Ca^{2+} influx) and p_- for a minus-event (low Ca^{2+} influx). A plus-event happens, whenever the influx a is in the range $a \geq \Theta_h$. In order to compute p_+ and analogously p_- we need to specify the pair correlation function $C_{\mathrm{io}}(t_1, t_2)$ between the presynaptic and the postsynaptic spike train. Generally, the correlation function is defined as

$$C_{\mathrm{io}}(t_1, t_2) = \lim_{\mathrm{d}t \downarrow 0} \frac{\Pr(t_{\mathrm{pre}} \in [t_1, t_1 + \mathrm{d}t] \wedge t_{\mathrm{post}} \in [t_2, t_2 + \mathrm{d}t])}{\mathrm{d}t^2}.$$

Here we restrict ourselves to stationary processes, such that $C_{\mathrm{io}}(t_1, t_2) = C_{\mathrm{io}}(t_2 - t_1)$ only depends on the relative timing $\tau = t_2 - t_1$ between the presynaptic and the postsynaptic spike. Furthermore, we assume the presynaptic spikes to be Poisson events emitted at rate ν_i. The postsynaptic spikes appear with mean rate ν_o. For large τ the cross-correlation function decays to $\lim_{|\tau| \to \infty} C_{\mathrm{io}}(\tau) = \nu_i \nu_o$. Hence we can write

$$C_{\mathrm{io}}(\tau) = \nu_i \nu_o + c_{\mathrm{io}}(\tau),$$

where the cross-covariance c_{io} vanishes for large $|\tau|$. We now know the conditional probability to observe a presynaptic spike at $t_{\mathrm{pre}} \in [t - \tau, t - \tau + dt]$ provided that a postsynaptic spike occurs at time t

$$\Pr(t_{\mathrm{pre}} \in [t - \tau, t - \tau + \mathrm{d}\tau] \,|\, \text{post spike at } t) = \nu_i \, \mathrm{d}\tau + \frac{1}{\nu_o} c_{\mathrm{io}}(\tau) \, \mathrm{d}\tau. \quad (9)$$

For a presynaptic spike to cause a plus-event it must have appeared within the potentiation window (see **Figure 4A**), i.e. $0 \leq \tau \leq \Delta t_+$. So the correlation detector measures the probability

ε_{eff} exceeding chance level of finding a presynaptic spike within this potentiation window, given a postsynaptic event at $t = 0$. This probability is $\varepsilon_{\text{eff}} = \frac{1}{\nu_o}\int_0^{\Delta t_+} c_{\text{io}}(\tau)\,\mathrm{d}\tau$. The cross-covariance $c_{\text{io}}(\tau)$ decreases on a time scale comparable to the membrane time constant of the neuron, which is typically shorter than the potentiation window $\Delta t_+ = 20$ ms. An example can be seen in **Figure 7A**. In this case we can make the simplification

$$\varepsilon_{\text{eff}} = \frac{1}{\nu_o}\int_0^{\Delta t_+} c_{\text{io}}(t')\,\mathrm{d}t' \overset{c_{\text{io}}\text{decays faster than }\Delta t_+}{=} \frac{1}{\nu_o}\int_0^{\infty} c_{\text{io}}(t')\,\mathrm{d}t'. \tag{10}$$

The probabilities of plus and minus-events are then given by Eqs. 6 and 7, respectively, with $\varepsilon = \varepsilon_{\text{eff}}$. An analytic approximation for the effective correlation coefficient in the framework of linear response theory for an integrate and fire neuron model can be found in the section "Input Output Correlation of an Integrate and Fire Neuron" in Appendix. Note that the causal dependence of output spikes on the input spikes and the assumed Poisson statistics of the incoming activity leads to the input–output correlation function (**Figure 7A**) which only deviates from baseline for $\Delta t > 0$. Therefore, the position of the temporal window for minus events is uncritical in this setup, as long as c_{io} is at baseline within this window. Thus we would obtain the same results, if the time window for minus events was at negative times.

We now have the tools to investigate synaptic pruning in the scenario depicted in **Figure 7B**. A neuron initially has a number k_0 of synaptic excitatory inputs, each of which reaches the neuron via a spiny synapse with a calcium based correlation detector as described in the section "A Counter for Correlated Events" All synapses are mature and may eventually die depending on the correlation variable x: A synapse is pruned as soon as the amount x of active CaMKII undercuts a critical threshold $x < X_{\text{d}}$. The initial distribution of x over the ensemble of synapses is the eigenvector of the slowest decaying eigenmode at the initial firing frequency $\nu_o(0)$ of the neuron. The choice is justified, if we think of the initial connectivity as the outcome of a slow dynamic wiring process, during which the synaptic amount x of active CaMKII had enough time to settle in this eigenmode. In addition to the excitatory inputs, the neuron receives a static configuration of inhibitory connections. Due to the pruning process, excitatory synapses progressively die and hence the number of excitatory connections $k(t)$ decays and the neuron's firing rate ν_o decreases (see **Figures 8A,B**). The pruning process continues until the postsynaptic neuron stops spiking. Thus, pruning is a mechanism to regulate the firing rate downwards. If the process is slow compared to the dynamics of x, we can assume the distribution of x to follow the eigenstate adiabatically.

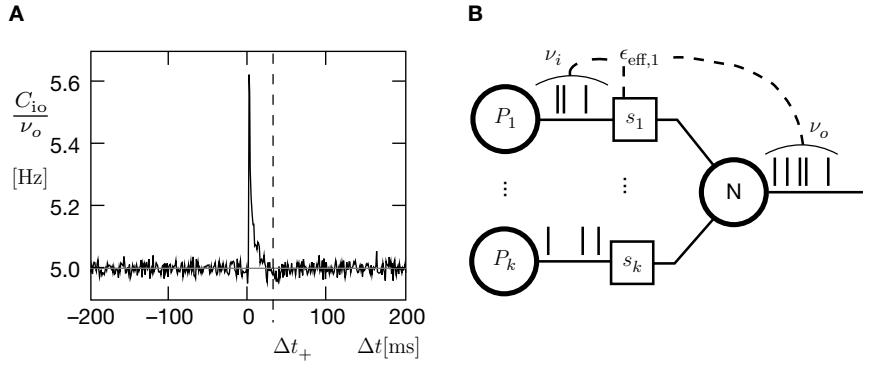

FIGURE 7 | (A) Input–output correlation function normalized to the postsynaptic spike rate. The peak drops to baseline ($\nu_i = 5$ Hz, gray) on a time scale which is shorter than the width of the potentiation window $\Delta t_+ = 20$ ms. The area below the peak is the effective correlation coefficient ε_{eff} (see Eq. 10). **(B)** A neuron *N* receiving *k* Poisson spiking inputs of rate $\nu_i = 5$ Hz via synapses $s_1 \ldots s_k$. Each synapse measures the correlation between its input spike train and the spiking activity of the neuron *N*. A synapse is pruned as soon as its number of active CaMKII has fallen below a threshold $x < X_d = 30$.

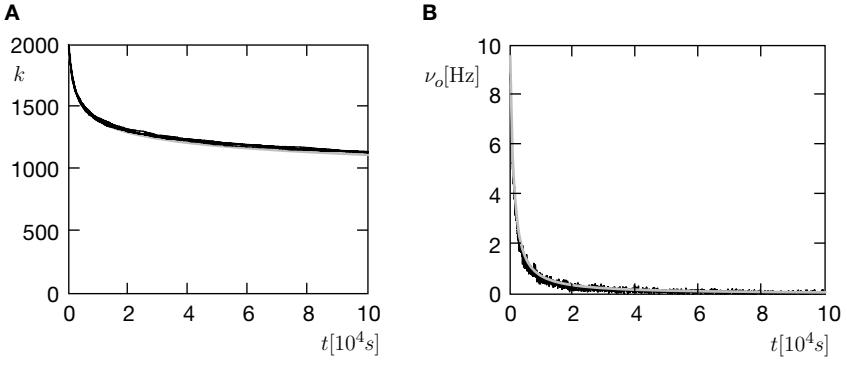

FIGURE 8 | (A) The number of incoming connections (vertical) in the course of the pruning process of **Figure 7B**. **(B)** The firing rate of the neuron (vertical) reduces with the reduction of inputs **(A)**. Black: five simulation trials. Gray: Numerical integration of Eq. 11. Initial number of incoming connections $k_0 = 2000$. In addition, the neuron receives 5080 excitatory connections and 1120 inhibitory connections, which are not pruned.

In this approximation, the development of connectivity obeys the differential equation

$$\frac{\mathrm{d}k}{\mathrm{d}t} = -k\eta_1(\nu_o(k), k)$$
$$k(0) = k_0, \tag{11}$$

where $\eta_1(\nu_o, k)$ is the eigenvalue of the slowest decaying eigenmode. All terms are accessible: a derivation of η_1 is presented in the section "Mean First Passage Time Problem for the Number of CaMKII Molecules" in Appendix and ν_o can be calculated using Eq. 34. Therefore, Eq. 11 can be numerically integrated. **Figure 8** compares this semi-analytic result with a direct simulation of the model.

The negative slope of the pruning curve in **Figure 8A** is decreasing with decreasing number of incoming synapses k. The reason is the dependence of the synapse death rate on the postsynaptic firing rate: The time scale of activation and deactivation of the CaMKII molecules is determined by the postsynaptic firing rate ν_o (compare Eq. 20). Consequently, also the death rate is proportional to ν_o. This is an interesting feature, since it facilitates a firing rate homeostasis: If new synapses are created with a constant rate the input connectivity to the neuron has a stable fixed point at k^* where synaptic death is just compensated by synapse creation. The firing rate of the neuron assumes a corresponding fixed firing rate ν_o^*. A similar example of such a homeostasis will also be shown in the section "Synaptic Maturation and Pruning".

COOPERATION AND COMPETITION BY SPATIAL INPUT CORRELATION

In the previous section we investigated a synaptic pruning process, where all excitatory inputs are uncorrelated. However, there is evidence that coactive inputs are stabilized (Cohen-Cory, 2002) and therefore less likely to be pruned. Here we show, that the calcium based correlation detection mechanism naturally leads to cooperation between synapses, which stabilizes coactive inputs.

SYNAPTIC PRUNING

In the first setup, we investigate a neuron receiving excitatory inputs from two different sources: A pool of n_p presynaptic neurons with uncorrelated Poisson spiking activity at rate ν_i. We call these inputs the "independent inputs" in the following. The second pool of n_c Poisson spiking neurons at rate ν_i, however, generates correlated spike trains. We use the multiple interaction process (Kuhn et al., 2003) to produce the spike trains and call these the "correlated inputs". The correlation coefficient $0 < c \leq 1$ is the probability of input neuron i having an input spike at time t, given neuron j has a spike at the same time. Thus, $c = 0$ results in uncorrelated Poisson processes, whereas for $c = 1$ all n_c spike trains are the same.

Initially (at $t = 0$) there are $n_p(0) = n_c(0) = 1000$ incoming synapses. Each of the synapses uses the calcium based correlation detection mechanism to determine the number of active CaMKII molecules x. Whenever x falls below the threshold X_d, the corresponding synapse dies. **Figure 9A** illustrates this scenario.

Figure 9B shows the evolution of the number of incoming connections. The synapses from independent sources exhibit a higher death rate than synapses from correlated inputs. We can readily understand this behavior: Given there is a spike at input i of the correlated pool, each of the remaining $n_c - 1$ inputs also delivers a spike with probability c at the same time. Thus, given a spike at the correlated input i, the expectation value for the sum of all inputs at this time is $\langle w \rangle_{\mathrm{mip}} = w(1 + c(n_c - 1))$, where w is the homogeneous synaptic weight. In contrast, a spike from the uncorrelated pool only carries its own weight $w < \langle w \rangle_{\mathrm{mip}}$. A higher synaptic weight results in a higher probability of the target neuron to emit a spike. Thus, the probability of the neuron to fire in response to a spike from one of the correlated inputs is higher than for an uncorrelated input. This probability is proportional to the correlation coefficient $\varepsilon_{\mathrm{eff}}$ between the presynaptic and the postsynaptic spike train (for a derivation of an analytic expression for $\varepsilon_{\mathrm{eff}}$ see Eq. 37 in the section "Correlated Poisson Input" in Appendix). With $\varepsilon_{\mathrm{eff,mip}} > \varepsilon_{\mathrm{eff,Poisson}}$, the number of active molecules x in the synapses from correlated inputs is higher than in those from independent inputs (see also **Figure 5B**) and hence their death rate is lower (see also "A Counter for Correlated Events").

Although there is no direct interaction between synapses from correlated inputs, cooperativity emerges between them and helps

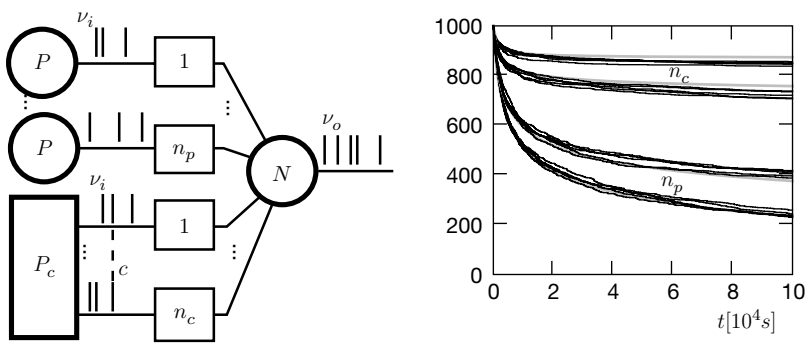

FIGURE 9 | (A) A neuron receiving two pools of excitatory inputs: n_p synapses provide uncorrelated Poisson inputs at $\nu_i = 5$ Hz, and n_c synapses correlated Poisson inputs at rate $\nu_i = 5$ Hz. The correlated events are produced by the multiple interaction process parameterized by the pairwise correlation c. A continuous pruning process eliminates synapses when the number of active CaMKII falls below a threshold $x < X_d = 30$. **(B)** Evolution of connectivity structure. Initially, both groups have the same number of synapses $n_p(0) = n_c(0) = 1000$. In addition, the neuron receives 5080 excitatory connections and 1120 inhibitory connections, which are not pruned. The pruning process mainly affects the uncorrelated inputs. The top and the bottom trace are for a pair correlation $c = 0.01$, the intermediate traces for $c = 0.005$. Black: five simulation trials. Gray: Numerical integration of Eq. 12.

to stabilize these inputs in favor of the uncorrelated inputs. There is not only cooperation among synapses of the correlated pool, but also competition between the two pools: The number of uncorrelated inputs decreases with increasing correlation among the correlated inputs. This is because the synaptic deathrate increases with the firing rate of the target neuron.

The time evolution of the number of inputs $n_p(t)$ and $n_c(t)$ can be calculated numerically completely analogous to the section "Rate Homeostasis by Synaptic Pruning". The system of differential equations governing the dynamics is

$$\begin{aligned}\frac{\mathrm{d}n_p}{\mathrm{d}t} &= -n_p \eta_{1,\,\mathrm{Poisson}}(\nu_o(n_p, n_c), n_p, n_c)\\ \frac{\mathrm{d}n_c}{\mathrm{d}t} &= -n_c \eta_{1,\,\mathrm{mip}}(\nu_o(n_p, n_c), n_p, n_c)\\ n_p(0) &= n_{p,0}\\ n_c(0) &= n_{c,0},\end{aligned} \tag{12}$$

where $\eta_{1,\mathrm{Poisson}}(\nu_o, n_p, n_c)$ and $\eta_{1,\mathrm{mip}}(\nu_o, n_p, n_c)$ are the slowest decaying eigen-modes calculated according to the section "Mean First Passage time Problem for the Number of CaMKII Molecules" in Appendix.

SYNAPTIC MATURATION AND PRUNING

In this section, we extend the scenario of **Figure 9A** by incorporating a process of synapse creation. Again we have two pools of input neurons: The first pool of N_p neurons has Poisson spiking activity with rate ν_i, the second pool of N_c neurons has Poisson activity with pair correlation c and the same rate ν_i. Excitatory synapses from both pools can exist in either of two states, premature or mature as depicted in **Figure 10A**. Newly created synapses are in the premature state, lacking AMPA receptors. Their synaptic weight is 0. A synapse becomes mature, if the number of active molecules x exceeds a threshold X_m. The mature synapse has the synaptic weight $w > 0$. This synapse dies, if the number of active molecules falls below a threshold X_d. Initially there are no premature synapses $n_{p,\mathrm{pre}}(0) = n_{c,\mathrm{pre}}(0) = 0$, and there are as many mature synapses from independent inputs as from correlated inputs $n_{p,\mathrm{mat}}(0) = n_{c,\mathrm{mat}}(0) > 0$. Premature synapses are constantly created: For each of the $N_p - n_{p,\mathrm{mat}} - n_{p,\mathrm{pre}}$ presently unestablished connections from independent sources, the rate of realization is λ_{pre}. Premature synapses from correlated sources are created analogously with the same rate λ_{pre}.

The evolution of the connectivity in the presence of maturation and pruning is shown in **Figure 10B**. The connectivity approaches an equilibrium state after $t \simeq 600$ s. The number of mature synapses from correlated inputs $n_{c,\mathrm{mat}}$ increases, while the number of synapses from uncorrelated inputs $n_{p,\mathrm{mat}}$ decreases (upper two traces). The two values saturate at different levels, such that $n_{c,\mathrm{mat}} > n_{p,\mathrm{mat}}$. Initially, the number of premature synapses increases for both input pools (lower two traces). In equilibrium, the numbers of premature synapses saturate at different levels $n_{p,\mathrm{pre}} > n_{c,\mathrm{pre}}$. The explanation for this observation is the same as in the previous subsection: Correlatedly activated synapses exhibit cooperation, they experience a higher correlation coefficient between the input spike train and the output spike train produced by the neuron. Hence, for synapses from the correlated pool, the maturation rate of premature synapses is higher and the deathrate of mature synapses is lower as compared to the uncorrelated inputs. As in the section "Synaptic Pruning", the target neuron becomes selective for the correlated pool.

A similar differential equation as Eq. 12 quantitatively describes the evolution of the connectivity in this model. Its numerical solution (**Figure 10B**, gray) corresponds well to a direct simulation of the system (black). The observed deviations are due to synapse maturation: Each synapse entering the mature state has a number x of active CaMKII just above threshold X_m. This perturbs the equilibrium distribution of x for the mature synapses and hence influences their death rate. The direction of this influence depends on the relative position of X_m with respect to the equilibrium $\langle x \rangle_{\mathrm{eq}}$: For $X_m > \langle x \rangle_{\mathrm{eq}}$ the observed death rate is lower than the analytic estimate and vice versa.

DISCUSSION

In the present work we describe a novel model of the synaptic mechanisms controlling synapse pruning and synapse maturation.

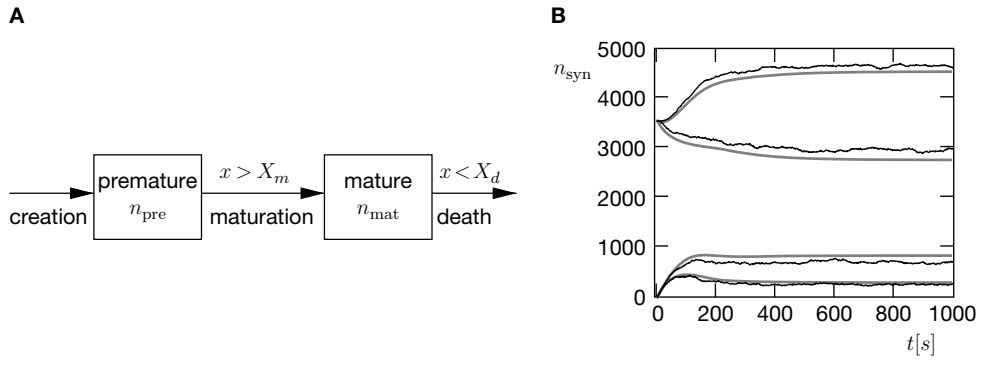

FIGURE 10 | (A) Each synapse has two states: In the premature state, its synaptic weight is $w = 0$. If the number x of active CaMKII exceeds a threshold $X_m = 45$, the synapse maturates and exhibits a synaptic weight $w = 0.05$ mV. If x falls below $X_d = 30$, the synapse dies. **(B)** Formation of input structure of a neuron during constant creation of premature synapses with rate $\lambda_{\mathrm{pre}} = 0.05\,\frac{1}{s}$ synapses per second in otherwise the same input scenario as in **Figure 9A**. Initially, there are $n_{p,\mathrm{mat}}(0) = n_{c,\mathrm{mat}}(0) = 3540$ mature excitatory synapses from each input pool and $n_{p,\mathrm{pre}}(0) = n_{c,\mathrm{pre}}(0) = 0$ premature synapses. Additionally, there are inputs from 1120 inhibitory synapses not subject to pruning. The traces show (top to bottom): mature synapses from correlated pool $n_{c,\mathrm{mat}}$, mature synapses from uncorrelated pool $n_{p,\mathrm{mat}}$, premature synapses from uncorrelated pool $n_{p,\mathrm{pre}}$, premature synapses from correlated pool $n_{c,\mathrm{pre}}$. Black: simulation, gray: analytical approximation.

To our knowledge, this is the first model of structural plasticity based on the microscopic dynamics of the single synapse. Instead of constructing a phenomenological model, we use recent experimental findings to identify plausible postsynaptic mechanisms and to constrain the choice of model parameters. We analyze the cortical scenario of irregular spiking activity and quantify the dependence of structural plasticity on the correlation between the presynaptic and the postsynaptic activity. Functionally, the dynamics acts as a Hebbian learning rule for synaptogenesis and pruning. Our model can be understood as a biological plausible implementation of a cascade synapse (Rubin and Fusi, 2007) for spiking activity, where the number x of active CaMKII molecules plays the role of internal states. However, in contrast to the cascade model, once x falls below the critical threshold the synapse is lost and cannot be reactivated by subsequent potentiating events. Here we provide a full statistical treatment of the number of active molecules. This is essential, since due to the low number of molecules ($N \simeq 80$) fluctuations cannot be neglected. Previous work on STDP (Cai et al., 2007; Shouval and Kalantzis, 2005) demonstrated, that stochastic fluctuations can change the phenomenology of a learning rule.

We show that there are two distinct rate regimes: In the low rate regime up to $\nu_i \simeq 10$ Hz, our model works as a correlation detector in the sense that synchronously activated synapses on the same dendrite stabilize. This frequency limit is a direct consequence of the NMDA unbinding time constant τ_{nmda}. Hence the model constitutes a viable mechanism for a network using temporal (correlation) codes at low rates. At higher rates ($\nu_i > 10$ Hz), synapse stabilization occurs irrespective of correlations between input and output. This is in accordance with recent experimental evidence that higher activity leads to increased synapse formation (Le Be and Markram, 2006) and to classical LTP induced by tetanic stimulation protocols. Furthermore, we show that by controlling synaptic pruning our model exhibits several desirable features for a neuronal network: The synaptic pruning rate increases with the postsynaptic firing rate. This enables the homeostasis of firing rate while synapses are continuously created, a finding obtained earlier using abstract rate based models of structural plasticity (van Ooyen et al., 1995). Moreover, the synapses targeting the same neuron naturally exhibit cooperation and competition. These emerging phenomena render the proposed microscopic mechanism relevant for the theory of learning in neuronal networks: The evolution of connectivity is sensitive to correlations in the inputs and hence neurons become selective for coactive inputs. This stabilization of coactive inputs has been proposed earlier based on experimental evidence, reviewed in (Cohen-Cory, 2002). Our treatment explains, how cooperation and competition are mediated solely by the identical postsynaptic activity experienced by synapses. Previous theoretical work (Kempter et al., 1999, 2001) showed cooperativity to emerge from the interplay of spike timing dependent learning rules with the spiking dynamics in the framework of spike response models. In general, a quantitative understanding of the interaction between a spike timing based plasticity rule and the integrate-and-fire dynamics is a hard problem. Here we provide such an analysis for our specific learning rule assuming all-to-all spike interaction. The analysis allows us to obtain semi-analytic expressions for the evolution of network structure. Specifically, we present results for the case of incoming irregular Poisson activity as well as for the case of correlated inputs generated by a multiple-interaction-process (Kuhn et al., 2003).

SENSITIVITY OF RESULTS TO MODEL ASSUMPTIONS

For analytical convenience we use an integrate and fire neuron model with δ-shaped postsynaptic currents, as commonly used in network simulations and theoretical works. These currents can cause an immediate spiking response to an incoming spike. More realistic neuron models have postsynaptic currents with finite rise times and hence also the input–output correlation function shows a finite latency of $\tau_{\mathrm{response}} \simeq 5$ ms in response to an input spike. Since we aim at a consistent theory of the interaction between learning rule and neuronal dynamics, we compensate for the lack of latency by reducing the measured glutamate binding time $\tau_{\mathrm{rise}} \simeq 5 \ldots 10$ ms to $\tau_{\mathrm{rise}} = 0$ ms. As in the natural setting the whole mass of the input–output correlation function falls into the time window τ_+ of the learning rule. A mismatch of the time constants would not change the observed phenomena qualitatively, but slightly reduce the sensitivity to correlations. We can extend our analysis to more realistic and more complicated neuron models, if an expression for the corresponding integral input–output correlation $\int_0^\infty C_{\mathrm{io}}(t)\,\mathrm{d}t$ is known.

The probabilities p and q for activation and deactivation of CaMKII molecules by large and small calcium events respectively, are set to $p = q = 0.01$ without experimental reference. Once experimental data are available, these parameters need to be reconsidered. The smaller the values, the narrower the distribution of active CaMKII molecules. This increases the sensitivity to correlation but also the latency of the distribution in following transient changes in the correlation. The threshold X_d for the minimum number of active CaMKII molecules required for survival can be chosen such that the synapse turnover in a neuronal network reaches experimentally observed values of 7% of all synapses per week (Stettler et al., 2006). As soon as an experimental value for X_d is available we can check whether our model consistently relates the two experimental measures. We use a multiple interaction process to generate incoming spiking activity with higher order correlations. The choice is motivated by analytical convenience rather than by biological realism. However, we expect other models of higher order correlation to exhibit qualitatively similar results for cooperation and competition between synapses.

Due to our assumption of a point event like bpAP (see also "Spike Time Dependence of Postsynaptic Calcium Concentration") there is no calcium influx through NMDA receptors, if the postsynaptic depolarization precedes the presynaptic release of glutamate. For our model the consequence is that post-before-pre pairings do not cause deactivation of CaMKII molecules (see also "A Counter for Correlated Events"). Hence if the intention was to explain STDP, we would not be able to reproduce the part of the rule expressing LTD. Experimentally, Nevian and Sakmann (2004) observe calcium influx through NMDA receptors only for the pre-before-post condition. Additionally, there is influx through voltage dependent calcium channels (VDCC) opened by the bpAP. Also AMPA receptors depolarize the spine and hence lead to calcium entry through coactivated NMDA receptors and VDCCs. However, the resulting NMDA conductivity is small (Kampa et al., 2004) and calcium transients in spines are not significantly decreased by

blocking AMPA receptors (Nevian and Sakmann, 2004). So here we assume the main pathway for calcium influx into the spine to be NMDA receptors and we neglect the additional sources of calcium. As long as their contribution to the calcium signal is small, they constitute a background elevation of the overall calcium level in the spine and we obtain qualitatively similar results. If, however, their contributions were of comparable size as the calcium transient through NMDA receptors, a temporally close post-before-pre timing could lead to sufficient calcium influx to deactivate CaMKII and hence create a LTD window. Previous biophysical models of STDP (Shouval et al., 2002) assumed a finite decay time for the postsynaptic depolarization. This as well leads to an intermediate calcium influx for the post-before-pre condition. Our model can be extended in the same way. However, as argued in the section "Rate Homeostasis by Synaptic Pruning", for the results presented here, the exact temporal position of the LTD window is irrelevant. The observed rate homeostasis only depends on the proportionality between the synapse death rate and the postsynaptic firing rate, resulting from the CaMKII dynamics. The cooperativity between synapses (see "Cooperation and Competition by Spatial Input Correlation") only requires a Hebbian type plasticity, i.e. plus events are caused by a postsynaptic spike following a presynaptic spike in close succession.

SCOPE AND ACCURACY OF THE ANALYTICAL TREATMENT

The treatment of the CaMKII dynamics as a Markov process implicitly assumes the high (plus) and low (minus) calcium events to occur uncorrelatedly with the respective probabilities p_+ or p_-. However, the occurrence of a plus event or a minus event depends on the current value $n(t_{\text{post}})$, the fraction of glutamate bound NMDA receptors at the time of the postsynaptic spike. $n(t)$ has a time constant $\tau_{\text{nmda}} = 32$ ms. Hence, if postsynaptic spikes occur with arbitrarily small inter-spike-intervals, the probability to observe a plus event is slightly higher after a previous plus event. Nevertheless, for realistic postsynaptic spike trains, small inter-spike-intervals are rare due to refractoriness, so the neglect of the temporal correlation of $n(t)$ is well justified (see also section "Mean First Passage Time Problem for the Number of CaMKII Molecules" in Appendix, **Figure 12B**). A more thorough treatment must take into account the actual auto-correlation function of the postsynaptic spiking activity. The dynamics of the discrete amount of active CaMKII molecules is mapped to a continuous system. By comparing the continuous analytic probability density to the discrete numerical solution of the probability mass function, we found that this approximation is sufficiently accurate as long as the total number of molecules is large enough ($N > 30$). Furthermore, we approximate the CaMKII dynamics as a diffusion with a x-independent diffusion constant, whereas the original problem leads to a x-dependent diffusion term. Again, the nearly perfect agreement of the numerical solution (taking into account x-dependent diffusion) with the analytically obtained Gaussian, proves this approximation adequate (see **Figure 11**). However, if the distribution approaches the saturation limits $x = 0$ or $x = N$, we observe pronounced deviations. This may occur for highly correlated spiking and at excessive firing rates. The results for the synaptic death and maturation rates are based on the slowest decaying eigenvalue η_1 of the CaMKII activation distribution. This component has the largest time constant. If the ensemble of synapses is initialized with a distribution containing contributions different from the slowest decaying eigenvector, we observe transient deviations in the pruning rate. These transients decay typically 15 times faster than the slowest eigenvector $\left(\frac{|\eta_1|}{|\eta_2|} \simeq \frac{1}{15}\right)$. So after a sufficiently long time, the CaMKII distribution obeys the analytical solution in good approximation. For calculating the equilibrium state of connectivity, the slowest decaying eigenvector is the exact description if the number x of active molecules of a new synapse is drawn from this eigenvector. But even for non-stationary connectivity, the distribution of activated CaMKII across synapses follows the structure quasi adiabatically and the approximation is well fulfilled as shown by comparing analytical results with direct simulations (see "Rate Homeostasis by Synaptic Pruning" and "Cooperation and Competition by Spatial Input Correlation"). The larger deviations in **Figure 10** are attributed to this approximation: Each synapse entering the mature state has a number x of active CaMKII just above threshold X_m. This perturbs the equilibrium distribution, explaining the deviation of the simulated from the analytically obtained connectivity structure.

FUNCTIONAL ROLE OF STRUCTURAL PLASTICITY

Recent experiments (Stettler et al., 2006) demonstrated synaptogenesis and synaptic death to occur even in the adult cortex. However, their functional relevance for the neural network still remains to be illuminated. Theoretical considerations (Chklovskii et al., 2004; Stepanyants et al., 2002) suggest rewiring in networks to provide the dominating substrate of information storage. In artificial neural networks optimization of wiring was already shown to contribute significantly to memory capacity (Knoblauch, 2006; Knoblauch et al., 2007). More generally, restructuring of connectivity allows biological systems to optimize their circuitry to fulfill a specific set of functions. Understanding the mechanisms of wiring optimization will also be beneficial for technical systems (e.g. integrated circuits), since due to their essentially two dimensional nature, the number of contacts existing at any point in time is a precious and limited resource. For the above purposes a phenomenological model of structural plasticity is sufficient. However, a microscopic model is required to uncover the underlying biological mechanisms and resulting limitations. Moreover, an understanding of the control mechanisms of connectivity may contribute to the development of medical protocols to promote plasticity after neural lesions, e.g. as experienced after a stroke. How a system can exhibit a plastic structure and yet acquire and maintain its functionality is still a matter of research. To answer this question it is not sufficient to investigate synaptogenesis and synaptic death under conditions of stationary activity, but we rather have to consider the close interplay between the structural dynamics and the correlation dynamics exposed by the present work. The framework presented here provides the tools for this endeavor; we are now in the position to investigate structural plasticity in recurrent neural networks.

APPENDIX

PROBABILITY DISTRIBUTION FOR THE NUMBER OF ACTIVE CaMKII MOLECULES

Numeric solution of the equilibrium distribution

Suppose there is a finite amount of CaMKII in the PSD. Each of the N molecules can either be active or inactive. Transitions between

these states are triggered by calcium influx, where a high calcium amplitude $a \geq \Theta_h$ causes each inactive molecule to be activated with probability p, and a low calcium amplitude $\Theta_b \leq a < \Theta_l$ causes each active molecule to be inactivated with probability q. The rate of high calcium events is $\lambda_+ = \nu_o p_+(\varepsilon)$, the rate of low calcium events $\lambda_- = \nu_o p_-(\varepsilon)$, where ν_o denotes the postsynaptic firing rate. Given x active molecules, the number of molecules being activated per time is $\lambda_+ p(N-x)$, the current of molecules being deactivated is $\lambda_- qx$. Since these rates depend on the state x of the system, the process has the Markov property. Thus, the system is uniquely defined by the transition probability $P(x, y)$ from state x into state y. Let $\lambda = \nu_o(p_+(\varepsilon) + p_-(\varepsilon))$ be the rate of events that change the state of the system, $\tilde{p}_+ = \frac{p_+(\varepsilon)}{p_+(\varepsilon)+p_-(\varepsilon)}$ the probability that the event was a plus-event and $\tilde{p}_- = \frac{p_-(\varepsilon)}{p_+(\varepsilon)+p_-(\varepsilon)}$ be defined analogously. Given an event (either plus or minus), the transition probability is

$$\begin{aligned} P(x,y) &= H(y-x)\tilde{p}_+ B(y-x \mid p, N-x) \\ &\quad + H(x-y)\tilde{p}_- B(x-y \mid q, x), \end{aligned} \tag{13}$$

where $B(k \mid p, N) = \binom{N}{k} p^k (1-p)^{N-k}$ is the binomial distribution and H the Heaviside function. We are interested in the equilibrium distribution $\rho(x)$, which must fulfill

$$\begin{aligned} \rho(y) &= \sum_{x=0}^{N} \rho(x) P(x,y) \\ 0 &= \sum_{x=0}^{N} \rho(x)(P(x,y) - \delta_{x,y}) \qquad \forall y \in [0, N] \end{aligned}$$

To obtain the non-trivial solution, we have to take into account the normalization condition $1 = \Sigma_{x=0}^{N} \rho(x)$. This results in the inhomogeneous linear system of $N+1$ equations

$$\delta_{y,0} = \sum_{x=0}^{N} \rho(x)(P(x,y) - \delta_{x,y} + \delta_{y,0}) \qquad \forall y \in [0, N], \tag{14}$$

which can be solved numerically. Since the transition probability $P(x, y)$ is a positive stochastic matrix, i.e. it fulfills $P(x, y) \geq 0$ and $\forall x : \Sigma_{y=0}^{N} P(x, y) = 1$, according to the Perron-Frobenius theorem (MacCluer, 2000) its largest eigenvalue is 1 and the respective eigenvector $\rho(x)$ is unique with positive entries. This guarantees a unique solution of Eq. 14 with the desired properties of a probability distribution.

Analytic approximation of the equilibrium distribution

In order to obtain information about the equilibrium density $\rho(x)$ we investigate its first and second moments. Suppose, x obeys the distribution $\rho(x, t)$ at time t with a well defined first and second moment. In this case we can determine the infinitesimal time evolution of $\mu(t) = \langle x(t) \rangle$ and $\sigma^2(t) = \langle (x(t) - \langle x \rangle)^2 \rangle$ for short times Δt. For the mean we obtain

$$\begin{aligned} \mu(t+\Delta t) &= \langle x(t+\Delta t) \rangle \\ &= \sum_{x=0}^{N} \rho(x,t) \Bigg(\underbrace{\lambda \Delta t \, \tilde{p}_+}_{\Pr(\text{plus event})} \sum_{k=0}^{N-x} \underbrace{\binom{N-x}{k} p^k (1-p)^{N-x-k}}_{\Pr(\text{activate } k \text{ molecules})} (x+k) \\ &\quad + \underbrace{\lambda \Delta t \, \tilde{p}_-}_{\Pr(\text{minus event})} \sum_{k=0}^{x} \underbrace{\binom{x}{k} q^k (1-q)^{x-k}}_{\Pr(\text{deactivate } k \text{ molecules})} (x-k) \\ &\quad + \underbrace{(1 - \lambda \Delta t (\tilde{p}_+ + \tilde{p}_-))}_{\Pr(\text{no event})} x \Bigg) \\ &= \lambda \Delta t \, \tilde{p}_+ Np + \langle x(t) \rangle (1 - \lambda \Delta t \, \tilde{p}_+ p - \lambda \Delta t \, \tilde{p}_- q), \end{aligned}$$

where the identity $\Sigma_{k=0}^{N} \binom{N}{k} p^k (1-p)^{N-k} k = Np$ (mean of the binomial distribution) was used. The first moment of the distribution fulfills the differential equation

$$\begin{aligned} \frac{d\mu(t)}{dt} &= \lim_{\Delta t \to 0} \frac{\langle x(t+\Delta t) \rangle - \langle x(t) \rangle}{\Delta t} \\ &= \lambda \tilde{p}_+ Np - \mu(t) \lambda (\tilde{p}_+ p + \tilde{p}_- q). \end{aligned} \tag{15}$$

In equilibrium, the mean value is

$$\mu_{\text{eq}} = \langle x \rangle_{\text{eq}} = N \frac{\tilde{p}_+ p}{\tilde{p}_+ p + \tilde{p}_- q}. \tag{16}$$

In addition, we can calculate its temporal evolution from Eq. 15 to be

$$\mu(t) = \mu(0) e^{-t\lambda(\tilde{p}_+ p + \tilde{p}_- q)} + \mu_{\text{eq}} (1 - e^{-t\lambda(\tilde{p}_+ p + \tilde{p}_- q)}).$$

So the time scale on which the distribution approaches its equilibrium is given by

$$\tau_{\text{relax}} = \frac{1}{\lambda(\tilde{p}_+ p + \tilde{p}_- p)}.$$

Similarly, the infinitesimal time evolution of the variance unfolds to

$$\begin{aligned} &\sigma^2(t+\Delta t) \\ &= \left\langle (x(t+\Delta t) - \langle x(t) \rangle)^2 \right\rangle \\ &= \sum_{x=0}^{N} \rho(x) \Bigg(\lambda \tilde{p}_+ \Delta t \sum_{k=0}^{N-x} \binom{N-x}{k} p^k (1-p)^{N-x-k} (x+k-\mu(t))^2 \\ &\quad + \lambda \tilde{p}_- \Delta t \sum_{k=0}^{x} \binom{x}{k} q^k (1-q)^{x-k} (x-k-\mu(t))^2 \\ &\quad + (1 - \lambda \Delta t (\tilde{p}_+ + \tilde{p}_-))(x-\mu(t))^2 \Bigg) \\ &= \lambda \Delta t \, \sigma^2(t) (\tilde{p}_+ (-2p + p^2) + \tilde{p}_- (-2q + q^2)) \\ &\quad + \lambda \Delta t \, \mu^2(t) (\tilde{p}_+ p^2 + \tilde{p}_- q^2) \\ &\quad + \lambda \Delta t \, \mu(t) (\tilde{p}_+ (-p(1-p) - 2Np^2) + \tilde{p}_- q(1-q)) \\ &\quad + \lambda \Delta t \tilde{p}_+ (Np(1-p) + N^2 p^2) + \sigma^2(t), \end{aligned}$$

where we used $\Sigma_{k=0}^{N} \binom{N}{k} p^k (1-p)^{N-k} k = Np$ and $\Sigma_{k=0}^{N} \binom{N}{k} p^k (1-p)^{N-k} k^2 = Np(1-p) + (Np)^2$. We obtain the differential equation

$$\frac{d\sigma^2(t)}{dt} = \lim_{\Delta t \to 0} \frac{\sigma^2(t+\Delta t) - \sigma^2(t)}{\Delta t}$$

$$
\begin{aligned}
&= \lambda\sigma^2(t)(\tilde{p}_+(p^2-2p)+\tilde{p}_-(q^2-2q)) \\
&\quad + \lambda\mu^2(t)(\tilde{p}_+p^2+\tilde{p}_-q^2) \\
&\quad + \lambda\mu(t)(\tilde{p}_-q(1-q)-\tilde{p}_+p((1-p)+2Np)) \\
&\quad + \lambda\tilde{p}_+N(p+(N-1)p^2)
\end{aligned}
$$

In equilibrium, we determine

$$
\sigma_{\mathrm{eq}}^2 = \frac{-\mu_{\mathrm{eq}}^2(\tilde{p}_+p^2+\tilde{p}_-q^2)+\mu_{\mathrm{eq}}(\tilde{p}_+p((1-p)+2Np)-\tilde{p}_-q(1-q))-\tilde{p}_+N(p+(N-1)p^2)}{\tilde{p}_+(p^2-2p)+\tilde{p}_-(q^2-2q)}. \tag{17}
$$

Figure 11 compares the numerical solution of Eq. 14 to a Gaussian with mean μ_{eq} and standard deviation σ_{eq}, demonstrating that the Gaussian approximation is sufficient.

Mean first passage time problem for the number of CaMKII molecules

In the model the number of active CaMKII molecules x determines whether a synapse maturates or dies. Both decisions are triggered by the crossing of different thresholds. Thus we have to calculate the mean time until x passes the threshold X for the first time. This is known as the mean first passage time problem with an absorbing boundary at $x = X$. Here we are interested in an approximation for the mean first passage time, or equivalently, for the threshold passing rate.

Figure 5 shows that the equilibrium distribution $\rho(x)$ in absence of any absorbing boundary is well described by a Gaussian distribution. This observation suggests a mapping of the discrete first passage time problem to a continuous one with $\tilde{x} \in [0, N] \subset \mathbb{R}$ and the corresponding replacement of the discrete probability distribution ρ by a continuous probability density function $\tilde{\rho}$

$$
\begin{aligned}
&\rho : [0,\ldots,N] \subset \mathbb{N} \mapsto \mathbb{R} \\
&\qquad \rho(x) = \int_{x-\frac{1}{2}}^{x+\frac{1}{2}} \tilde{\rho}(x')\,\mathrm{d}x' \\
&\text{with } \tilde{\rho} : \left[-\frac{1}{2}, N+\frac{1}{2}\right] \subset \mathbb{R} \mapsto \mathbb{R}.
\end{aligned} \tag{18}
$$

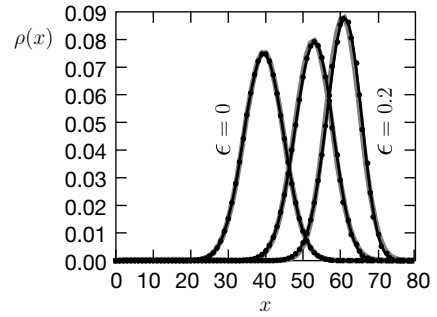

FIGURE 11 | Comparison between the numerically obtained distribution of the number of active molecules in an ensemble of synapses solving Eq. 14 (black dotted curve) and the Gaussian approximation (gray curve) using Eqs. 16 and 17 for different pair correlation coefficients $\varepsilon = \{0, 0.1, 0.2\}$ between the presynaptic and the postsynaptic spiking activity. Same parameters as in **Figure 5**.

In the following we omit the tilde; x and ρ name the continuous variables. The time evolution of the probability density $\rho(x)$ can be described by the Fokker–Planck equation

$$
\frac{\partial\rho}{\partial t} = -\frac{\partial}{\partial x}S(x)\rho, \tag{19}
$$

where $S(x)$ denotes the probability flux operator. An explicit form can be determined under the following assumptions: (1) The process can be described as a diffusion. (2) The equilibrium distribution of the process is given by a Gaussian $\rho_0(x) = \frac{1}{\sqrt{2\pi}\sigma_{\mathrm{eq}}} e^{-(x-\mu_{\mathrm{eq}})^2/(2\sigma_{\mathrm{eq}}^2)}$, where μ_{eq} and σ_{eq} are determined by Eqs. 16 and 17. The stationary probability density has to fulfill $\frac{\partial\rho}{\partial t} = 0$, so from Eq. 19 it follows that $S(x)\rho_0(x) = S_0 = \text{const}$. The constant must be $S_0 = 0$, since the probability current vanishes at $x = -\frac{1}{2}$ it vanishes for all x. We intend to describe the process as a diffusion. Therefore, S should contain only first derivatives in x and the flux operator must have ρ_0 as its stationary solution, leading to

$$
\begin{aligned}
&0 = S(x)\rho_0(x) \\
\Rightarrow\ &S(x) = D\left(-(x-\mu_{\mathrm{eq}}) - \sigma_{\mathrm{eq}}^2\frac{\partial}{\partial x}\right).
\end{aligned}
$$

The constant D controls the time scale of the process. We determine D such that the dynamics of the diffusion process Eq. 19 matches the dynamics of our process for the first moment of ρ. In doing so, we follow Ricciardi et al. (1999) and determine the infinitesimal drift term $A_1(x)$ from the master equation of the process, which yields

$$
\begin{aligned}
A_1(x) &= \lim_{\Delta t\downarrow 0}\frac{1}{\Delta t}E[x(\Delta t)-x(0)] \\
&= \underbrace{\lambda\Delta t\tilde{p}_+}_{\Pr(\text{plus event})}\sum_{k=0}^{N-x}\underbrace{\binom{N-x}{k}p^k(1-p)^{N-x-k}k}_{\Pr(k\text{ particles activated})} \\
&\quad + \underbrace{\lambda\Delta t\tilde{p}_-}_{\Pr(\text{minus event})}\sum_{k=0}^{x}\underbrace{\binom{x}{k}q^k(1-q)^{x-k}k}_{\Pr(k\text{ particles deactivated})} \\
&= \lambda((-\tilde{p}_+p-\tilde{p}_-q)x + N\tilde{p}_+p).
\end{aligned} \tag{20}
$$

Subsequently we require $A_1(x) = -D(x-\mu_{\mathrm{eq}})$, which fixes $D = \lambda(\tilde{p}_+p+\tilde{p}_-q)$.

To simplify the notation, we rescale the variable $y = \frac{x-\mu_{\mathrm{eq}}}{\sigma_{\mathrm{eq}}}$. The flux operator now reads $S(y) = D\sigma_{\mathrm{eq}}(-y-\frac{\partial}{\partial y})$. A separation ansatz for the probability density $\rho(y,t) = \phi(y)e^{\eta t}$ turns the Fokker–Planck equation 19 into the corresponding eigenvalue problem

$$
\begin{aligned}
\eta\phi(y) &= -\tfrac{1}{\sigma_{\mathrm{eq}}}\tfrac{\partial}{\partial y}S(y)\phi(y) \\
&= \underbrace{D\frac{\partial}{\partial y}(y+\frac{\partial}{\partial y})}_{=:L_{\mathrm{FP}}}\phi(y).
\end{aligned} \tag{21}
$$

To find the eigenvalues of L_{FP}, we follow Risken (1996) and transform Eq. 21 into a Hermitian operator, using the following transformation

$$U(y) = \mathrm{e}^{-\frac{1}{2}\int^y y'\,\mathrm{d}y'} = \mathrm{e}^{-\frac{y^2}{4}},$$

which satisfies

$$\frac{\partial}{\partial y} U(y) \circ = U(y)\left(-\frac{1}{2} y + \frac{\partial}{\partial y}\right) \circ .$$

Hence the transformed operator L is

$$\begin{aligned} L &= U^{-1}(y) L_{\mathrm{FP}} U(y) \\ &= D\left(\frac{1}{2} - \frac{1}{4} y^2 + \frac{\partial^2}{\partial y^2}\right). \end{aligned}$$

Let $\phi(y)$ be an eigenvector of L_{FP}, then obviously $\psi = U^{-1}(y)\phi(y)$ is an eigenvector of L with the same eigenvalue, so we have to solve

$$\begin{aligned} L\psi &= \eta\psi \\ 0 &= \left(\frac{1}{2} - \frac{1}{4} y^2 + \frac{\partial^2}{\partial y^2} - \frac{\eta}{D}\right)\psi(y). \end{aligned}$$

This is the differential equation of the harmonic oscillator known from quantum mechanics. The general solution of this differential equation is a linear combination of confluent hypergeometric functions (Abramowitz and Stegun, 1974)

$$\begin{aligned} \psi(y) &= c\left(\mathrm{e}^{-\frac{1}{4}y^2} M\left(\frac{1}{2}\frac{\eta}{D}; \frac{1}{2}; \frac{1}{2} y^2\right) + r\,\mathrm{e}^{-\frac{1}{4}y^2} y M\left(\frac{1}{2}\frac{\eta}{D} + \frac{1}{2}; \frac{3}{2}; \frac{1}{2} y^2\right)\right) \\ \phi(y) &= U(y)\psi(y) \\ &= c\left(\mathrm{e}^{-\frac{1}{2}y^2} M\left(\frac{1}{2}\frac{\eta}{D}; \frac{1}{2}; \frac{1}{2} y^2\right) + r\,\mathrm{e}^{-\frac{1}{2}y^2} y M\left(\frac{1}{2}\frac{\eta}{D} + \frac{1}{2}; \frac{3}{2}; \frac{1}{2} y^2\right)\right) \\ &= c\left(M\left(-\frac{1}{2}\frac{\eta}{D} + \frac{1}{2}; \frac{1}{2}; -\frac{1}{2} y^2\right) + r\, y M\left(-\frac{1}{2}\frac{\eta}{D} + 1; \frac{3}{2}; -\frac{1}{2} y^2\right)\right). \end{aligned} \tag{22}$$

The last equality follows from $M(a, b, z) = \mathrm{e}^z M(b-a, b, -z)$ where the constant c must be determined using the normalization condition $1 = \int \phi \mathrm{d}x$. The boundary conditions constrain the choice of possible values for η as well as for the constants r. Without loss of generality, we assume that the threshold X has to be crossed from above (synapse pruning). The case in which the threshold represents an upper boundary (synapse maturation) can be handled analogously by interchanging the roles of plus and minus-events and the associated quantities. At the upper end $x = N + \frac{1}{2}$ of the interval, the probability flux has to vanish, i.e. this is a totally reflecting boundary

$$S\phi\big|_{x=N+\frac{1}{2}} = 0. \tag{23}$$

The lower end of the interval is given by the threshold X. This is an absorbing boundary, but does not require the probability density to vanish, since the rate of minus-events is limited and hence the threshold passing rate remains finite even if $\phi(X) > 0$. The probability for the system in state x to cross the threshold X after having received a minus-event is

$$\begin{aligned} p_{\mathrm{exit}}(x) &= \sum_{k=x-X+1}^{x} \binom{x}{k} q^k (1-q)^{x-k} \\ &= \sum_{l=x-k}^{l=x-k} \underbrace{\binom{x}{x-l}}_{=\binom{x}{l}} q^{x-l} (1-q)^l \\ &= \sum_{l=0}^{X-1} B(l|1-q, x) \end{aligned}$$

where the last expression is the cumulative density function of the binomial distribution. **Figure 12** shows the typical steep increase of the exit probability in the vicinity of the threshold X.

The flux η through the boundary can be calculated as

$$S_X = \lambda_- \sum_{x=X}^{N} p_{\mathrm{exit}}(x)\rho(x) \tag{24}$$

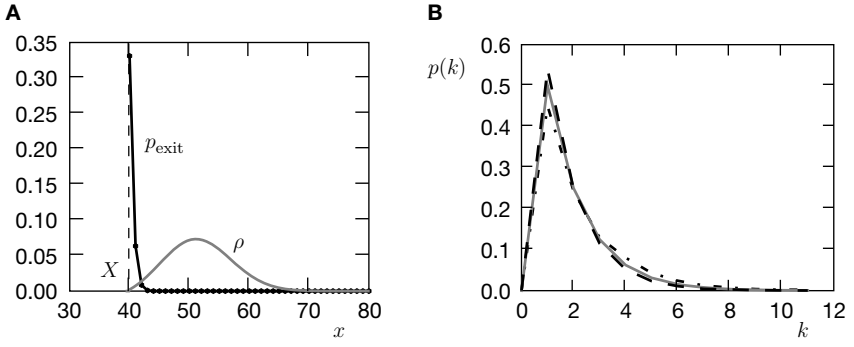

FIGURE 12 | (A) Exit probability p_{exit} given a minus-event depending on the state x (black) and probability density function ρ (gray) of the number of active molecules x. **(B)** Probability of having k consecutive plus-events given at least one plus-event. Expected probability $p(k) = \tilde{p}_+^{k-1}(1-\tilde{p}_+)$ for independently occurring events (gray). Statistics of events caused by a postsynaptic spike train of an integrate and fire neuron (black dashed line) and of events caused by a postsynaptic Poisson spike train (black dashed-dotted line).

$$\simeq \lambda_- \sum_{x=X}^{N} p_{\text{exit}}(x)\left(\rho(X)+\left.\frac{\partial\rho}{\partial x}\right|_{x=X}(x-X)\right). \tag{25}$$

The probability density can be approximated by a linear Taylor series here, since $p_{\text{exit}}(x)$ vanishes except in a small range near the threshold (see **Figure 12**). Using the flux operator S, the flux through the boundary can be expressed as

$$\begin{aligned} S_X &= S\rho\big|_X \\ &= -D\sigma_{\text{eq}}\left(\frac{X-\mu_{\text{eq}}}{\sigma_{\text{eq}}}\rho(X)+\frac{\partial\rho}{\partial y}(X)\right). \end{aligned} \tag{26}$$

Taken together we obtain the boundary condition

$$\begin{aligned} 0 &= \left(\lambda_- \sum_{x=X}^{N} p_{\text{exit}}(x) + D(X-\mu_{\text{eq}})\right)\rho(X) \\ &+ \left(\lambda_- \sum_{x=X}^{N} p_{\text{exit}}(x)\frac{x-X}{\sigma_{\text{eq}}} + D\sigma_{\text{eq}}\right)\frac{\partial\rho}{\partial y}(X). \end{aligned} \tag{27}$$

Using the explicit form of ρ given by Eq. 22, and its derivative

$$\begin{aligned} \frac{\partial\rho}{\partial y} &= y\left(\frac{\eta}{D}-1\right)M\left(-\frac{1}{2}\frac{\eta}{D}+\frac{1}{2}+1;\frac{1}{2}+1;-\frac{1}{2}y^2\right) \\ &+ r\,M\left(-\frac{1}{2}\frac{\eta}{D}+1;\frac{3}{2};-\frac{1}{2}y^2\right) \\ &+ r\,y^2\left(\frac{1}{3}\frac{\eta}{D}-\frac{2}{3}\right)M\left(-\frac{1}{2}\frac{\eta}{D}+2;\frac{3}{2}+1;-\frac{1}{2}y^2\right), \end{aligned}$$

where we exploit the identity $M'(a,b,z)=\frac{a}{b}M(a+1,b+1,z)$, the boundary condition Eq. 23 assumes the form

$$\begin{aligned} 0 &= S(y)\phi(y)\big|_{y=\frac{N+\frac{1}{2}-\mu_{\text{eq}}}{\sigma_{\text{eq}}}} \\ 0 &= \left(y+\tfrac{\partial}{\partial y}\right)\phi(y)\big|_{y=\frac{N+\frac{1}{2}-\mu_{\text{eq}}}{\sigma_{\text{eq}}}} \end{aligned}$$

which fixes the constant r as

$$r = -\left.\frac{yM\left(-\frac{1}{2}\frac{\eta}{D}+\frac{1}{2};\frac{1}{2};-\frac{1}{2}y^2\right)+y\left(\frac{\eta}{D}-1\right)M\left(-\frac{1}{2}\frac{\eta}{D}+\frac{3}{2};\frac{3}{2};-\frac{1}{2}y^2\right)}{(1+y^2)M\left(-\frac{1}{2}\frac{\eta}{D}+1;\frac{3}{2};-\frac{1}{2}y^2\right)+y^2\left(\frac{1}{3}\frac{\eta}{D}-\frac{2}{3}\right)M\left(-\frac{1}{2}\frac{\eta}{D}+2;\frac{5}{2};-\frac{1}{2}y^2\right)}\right|_{y=\frac{N+\frac{1}{2}-\mu_{\text{eq}}}{\sigma_{\text{eq}}}}. \tag{28}$$

Finally, we simultaneously solve Eqs. 28 and 27 for η by a simple numerical bisection method. By construction, the solutions $\eta_0>\eta_1>\ldots$ are eigenvalues of the Fokker–Planck operator Eq. 21. Since eventually the system has to pass the threshold, there cannot be an equilibrium eigenstate with $\eta_0=0$; all eigenstates are decaying. We are interested in the largest negative eigenvalue, which determines the slowest decaying eigenvector Eq. 22, i.e. the largest $\eta_1<0$ which solves Eq. 21 and fulfills the boundary conditions. Being initially ($t=0$) in this eigenstate, the probability of finding the system still in the interval $[0, X]$ after time t is

$$\Pr\left(x\in\left[-\frac{1}{2},X+\frac{1}{2}\right]\right)=\int_{-\frac{1}{2}}^{X+\frac{1}{2}}\rho(x,t)\,\mathrm{d}x=\mathrm{e}^{\eta_1 t}, \tag{29}$$

where η_1 is called the escape rate or the inverse mean first passage time. The eigenvector in Eq. 22 resulting from the considerations above is plotted in **Figure 6** together with results of direct simulations. The decay rate expected from Eq. 29 agrees well with the simulation. In this analytical treatment we have assumed that plus and minus-events occur independently of the history of previous events with probabilities $\tilde{p}_+$ and $\tilde{p}_-$, respectively. However, this is an approximation, since plus and minus-events depend on the signal $n(t)$, which has the time scale τ_{nmda}. Thus, given a postsynaptic spike at $t_{\text{post},1}$ and given $n(t_{\text{post},1})>\Theta_h$, a plus-event is produced. $n(t)$ then decays exponentially with the time constant τ_{nmda} or it jumps to even higher values, if a presynaptic spike arrives. Therefore a later postsynaptic spike $t_{\text{post},2}>t_{\text{post},1}$ definitely generates another plus-event as long as $n(t_{\text{post},i})\mathrm{e}^{-(t_{\text{post},2}-t_{\text{post},1})/\tau_{\text{nmda}}}>\Theta_h$, i.e. within a finite time window after the first event. Consequently, the probability $\tilde{p}_{++}$ for a plus-event to follow another plus-event is slightly higher than $\tilde{p}_+$. The same is true for minus-events. But since a neuron's spike train exhibits refractoriness and the probability is low that the second spike occurs within the time window determined by τ_{nmda}, the correlated occurrence of plus (or analogously minus) events has only a small impact. This is verified in **Figure 12B**, where the probability of k consecutive plus-events is displayed. For events caused by the spike train of an integrate and fire neuron, the curve is near its theoretically expected value $p(k)=\tilde{p}_+^{k-1}(1-\tilde{p}_+)$, while for spikes with Poisson statistics, the correlated occurrence of plus-events can be explained by $p(k)=\tilde{p}_{++}^{k-1}(1-\tilde{p}_{++})$ with $\tilde{p}_{++}>\tilde{p}_+$.

INPUT–OUTPUT CORRELATION OF AN INTEGRATE AND FIRE NEURON

Poisson input

Here we calculate the correlation coefficient of the incoming spike train at a given excitatory synapse and the output spike train of an integrate and fire neuron, where the correlation coefficient is defined by Eq. 10. If we have incoming Poisson spikes at a stationary rate ν_i, the neuron fires in a stationary fashion with rate ν_o as well. We can then rewrite the cross-correlation function as $C_{\text{io}}(\Delta t)=\nu_o(\Delta t\,|\,\text{given an input spike at } t=0)\nu_i$ and formally the correlation coefficient in Eq. 10 reads

$$\varepsilon_{\text{eff}}=\frac{\nu_i}{\nu_o}\int_0^\infty \nu_o(t'\,|\,\text{given an input spike at } t=0)-\nu_o\,\mathrm{d}t'. \tag{30}$$

For the particular case of δ-shaped postsynaptic currents we find an explicit expression for Eq. 30 as follows. The membrane potential V is governed by

$$\begin{aligned} \tau\frac{\mathrm{d}V}{\mathrm{d}t} &= -V+\tau\sum_i w_i\delta(t-t_i) \\ V(t_+) &= V_{\text{r}} \text{ if } V(t_-)>V_{\text{th}}, \end{aligned} \tag{31}$$

with reset potential V_{r} and spike threshold V_{th}. Here t_i is the time of the i-th incoming event and w_i its weight. For each event, the

weight w_i is a random variable with distribution K. This allows us to calculate Eq. 30 also for the case where a neuron receives N correlated spike trains (see section "Correlated Poisson Input"). The incoming events are assumed to be Poisson events with rate ν. Then the membrane potential is a Markov process with distribution $P(V, t)$ which obeys the evolution equation

$$P(V,t) = \int P^{\text{trans}}(V,t \mid V',t')P(V',t')\,\mathrm{d}V'.$$

We closely follow the derivation of the Fokker–Planck equation for the integrate and fire neuron given by Gerstner and Kistler (2002), but generalize it for arbitrary weight distributions K. As we assume stationary input, we set $t' = 0$ without loss of generality. For short times Δt the transition probability becomes

$$P^{\text{trans}}(V,\Delta t \mid V',0) = \underbrace{(1-\nu\Delta t)}_{\Pr(\text{no event})}\,\delta(V - V'\,\mathrm{e}^{-\frac{\Delta t}{\tau}}) + \underbrace{\nu\Delta t}_{\Pr(\text{event})}\int K(w)\delta(V-(V'+w)\,\mathrm{e}^{-\frac{\Delta t}{\tau}})\,\mathrm{d}w$$

and the membrane potential distribution after time Δt is

$$\begin{aligned}P(V,\Delta t) &= \int_{-\infty}^{\infty} P^{\text{trans}}(V,\Delta t \mid V')P(V',0)\,\mathrm{d}V' \\ &= (1-\nu\Delta t)\mathrm{e}^{\frac{\Delta t}{\tau}}P(V\mathrm{e}^{\frac{\Delta t}{\tau}},0) + \nu\Delta t\int K(w)\,\mathrm{e}^{\frac{\Delta t}{\tau_m}}P(V\mathrm{e}^{\frac{\Delta t}{\tau_m}}-w,0)\,\mathrm{d}w.\end{aligned}$$

Taking the limit $\Delta t \to 0$ yields

$$\tau\frac{\partial P(V,t)}{\partial t} = V\frac{\partial P}{\partial V}(V,t) + P(V,t) - \nu\tau P(V,t) + \nu\tau\int K(w)\,P(V-w,t)\,\mathrm{d}w.$$

A Kramers–Moyal expansion can be made under the assumption that the weight w is small and the rate ν is high. This is also known as the diffusion limit:

$$\begin{aligned}&\tau\frac{\partial P(V,t)}{\partial t} \\ &= V\frac{\partial P(V,t)}{\partial V} + P(V,t) - \nu\tau P(V,t) \\ &\quad + \nu\tau\int K(w)\left(P(V,t) - w\frac{\partial P(V,t)}{\partial V} + \frac{w^2}{2}\frac{\partial^2 P(V,t)}{\partial V^2} + O(w^3)\right)\mathrm{d}w \\ &= V\frac{\partial P(V,t)}{\partial V} + P(V,t) - \nu\tau P(V,t) + \nu\tau P(V,t) \\ &\quad - \nu\tau\langle w\rangle_K\frac{\partial P(V,t)}{\partial V} + \frac{1}{2}\nu\tau\langle w^2\rangle_K\frac{\partial^2 P(V,t)}{\partial V^2} + O\left(\langle w^3\rangle_K\right) \\ &= \frac{\partial}{\partial V}\left((V - \underbrace{\nu\tau\langle w\rangle_K}_{=:\mu}) + \underbrace{\nu\tau\langle w^2\rangle_K}_{=:\sigma^2}\frac{1}{2}\frac{\partial}{\partial V}\right)P(V,t) + O\left(\langle w^3\rangle_K\right).\end{aligned} \tag{32}$$

In the limit of small $\langle w\rangle_K$ only the first and the second moment of K enter the Fokker–Planck equation. This determines the drift constant μ and the diffusion constant σ^2 to

$$\begin{aligned}\mu &= \tau\nu\langle w\rangle_K \\ \sigma^2 &= \tau\nu\langle w^2\rangle_K.\end{aligned} \tag{33}$$

In equilibrium, the firing rate ν_o of the neuron can be calculated by the formula (Brunel, 2000; Brunel and Hakim, 1999; Ricciardi, 1977)

$$\frac{1}{\nu_\mathrm{o}} = \tau\sqrt{\pi}\int_{y_\mathrm{r}}^{y_\mathrm{th}} e^{u^2}(1+\mathrm{erf}(u))\,\mathrm{d}u$$

$$\begin{aligned}y_\mathrm{th} &= \frac{V_\mathrm{th}-\mu}{\sigma} \\ y_\mathrm{r} &= \frac{V_\mathrm{r}-\mu}{\sigma}.\end{aligned} \tag{34}$$

Finally, we employ linear response theory to calculate the input–output correlation coefficient in Eq. 30 for one particular incoming synapse in the limit of small weights w. First assume, that every event has the same weight w. Since the incoming spike rate is stationary, without loss of generality, we can assume an incoming spike to occur at time $t = 0$. This is a small perturbation of the membrane potential and shifts the distribution by weight w to the right, which can be taken into account in the Fokker–Planck equation 32 as

$$\begin{aligned}&\tau\frac{\partial P}{\partial t}(V,t) \\ &= \frac{\partial}{\partial V}\left(\frac{\sigma^2}{2}\frac{\partial}{\partial V} + V - \mu\right)P(V,t) + \tau\delta(t)\left(P(V-w,t) - P(V,t)\right) \\ &= \frac{\partial}{\partial V}\left(\frac{\sigma^2}{2}\frac{\partial}{\partial V} + V - \mu\right)P(V,t) + \tau\delta(t)\,(-w)\frac{\partial P}{\partial V}(V,t) + O(w^2) \\ &= \frac{\partial}{\partial V}\left(\frac{\sigma^2}{2}\frac{\partial}{\partial V} + V - \mu - w\tau\delta(t)\right)P(V,t) + O(w^2).\end{aligned}$$

To first order in w, the incoming spike can be described as a δ-shaped perturbation of the mean input $\mu(t) = \mu + \tau w\delta(t)$. Since the events are Poisson, their auto-correlation is flat, thus in the instant $t = 0^-$ before the spike arrives the neuron's membrane potential obeys the equilibrium distribution. The effect of the perturbation on the firing rate can be treated in the framework of linear response theory for small $w << V_\mathrm{th} - \mu$. The response of the firing rate can then reads

$$\nu_o(t \mid \text{input spike at } t = 0) = \nu_o + w\tau h(t)\ t \geq 0,$$

where ν_o is the equilibrium firing rate and $h(t)$ is the impulse response of the firing rate with respect to a δ-shaped perturbation of the mean input. So Eq. 30 becomes

$$\varepsilon_\mathrm{eff} = \tfrac{\nu_\mathrm{i}}{\nu_\mathrm{o}}w\tau\int_0^\infty h(t)\,\mathrm{d}t \tag{35}$$

where $\int_0^\infty h(t')\,\mathrm{d}t'$ is the step-response of the firing rate for $t \to \infty$ with respect to a unit-step change of the mean input. Up to linear order in w

$$\int_0^\infty h(t')\,\mathrm{d}t' = \frac{\partial \nu_o}{\partial \mu} + O(w^2)$$
$$= \sqrt{\pi}\nu_o^2\tau\frac{1}{\sigma}\left(e^{y_{th}^2}(1+\mathrm{erf}(y_{th})) - e^{y_r^2}(1+\mathrm{erf}(y_r))\right) + O(w^2)$$
$$=: \Omega(\mu,\sigma) + O(w^2). \tag{36}$$

In the second step we employ Eq. 34 and in the third summarize terms by the quantity $\Omega(\mu, \sigma)$ called "dc-susceptibility" in linear response theory. Using Eq. 36 in Eq. 35 results in the desired explicit expression. A related approach can be found in De la Rocha et al. (2007).

Correlated poisson input

The input–output correlation coefficient can also be calculated in the case, where the neuron receives spikes from N excitatory inputs of weight w generated by a multiple-interaction-process (mip) Kuhn et al. (2003) with mother spike rate ν_{mother} and copy probability c. In addition, the neuron receives uncorrelated excitatory and inhibitory Poisson input which we treat as a Gaussian white noise background input with mean μ_{bg} and standard deviation σ_{bg}. As in "Poission Input" we are interested in the correlation coefficient in Eq. 30 of a given input $i \in [1,\ldots,N]$ and the output spike train. A spike at time t_i at input i, indicates that there is a mother-spike that was successfully copied to input i. Hence, with probability c each of the remaining $N-1$ inputs also has a spike at time t_i. Therefore the expectation value of the weight of the composed event is $\langle w\rangle_{mip} = (1+(N-1)c)w$. For small $\langle w\rangle_{mip} \ll V_{th} - \mu$ the we can apply the framework of linear response theory as before and calculate Eq. 30 using Eqs. 36 and 35 as

$$\varepsilon_{eff} = \Omega(\mu,\sigma)\frac{\nu_i}{\nu_o}\tau\langle w\rangle_{mip}, \tag{37}$$

where

$$\mu = \mu_{bg} + \underbrace{wNc\nu_{mother}}_{\mu_{mip}}\tau$$
$$\sigma^2 = \sigma_{bg}^2 + \underbrace{w^2\nu_{mother}\tau(Nc(1-c)+N^2c^2)}_{\sigma_{mip}^2}.$$

The last two equations using μ_{mip} and σ_{mip}^2 result from Eq. 33.

ALGORITHMIC IMPLEMENTATION OF SYNAPSE MATURATION AND SYNAPSE DEATH

A synapse connecting an axon of neuron i to the dendrite of neuron j is stored in a list associated with neuron i. Without loss of generality we restrict the discussion to the case of a single synapse between i and j. In a distributed simulation, each process stores only the part of the list referring to the resident target neurons (see Morrison et al., 2005). The parameters and dynamic variables stored for each synapse j are: the synaptic weight w_j, the synaptic delay $d_j = d_j^D + d_j^A$, composed of the dendritic delay d_j^D and the axonal delay d_j^A, the number x_j of active CaMKII molecules, and the boolean variable mature_j indicating whether the synapse is mature (True) or premature (False). Axonal delays and the NMDA receptor rise time can be taken into account as long as $d_j^A + \tau_{rise,nmda} \le d_j^D$ (compare Morrison et al., 2005). However, in our simulations, we assume $d_j^A + \tau_{rise,nmda} = 0$ for sake of simplicity. We also assume the time constant τ_{nmda} to be identical for all synapses. In this case the fraction of glutamate bound NMDA receptors $n(t)$ as defined in Eq. 3 is the same for all axonal synapses of neuron i. In addition, each neuron stores the time t_{old} of its last spike. The spike distribution algorithm invokes the method `spike(t)` (see Algorithm 1) for each spike t of the presynaptic neuron i. It propagates the dynamics of each synapse in the local target list. Here we use the function $T(x, n')$ (see Algorithm 2) as an abbreviation for the transitions of x due to the number $n' = n(t_{post})$ of glutamate bound NMDA receptors as defined in the section "A Counter for Correlated Events". To evolve the synaptic dynamics, we need the spike history of the postsynaptic neuron. Here we

Algorithm 1 spike(t)

Require: t_{old} last presynaptic spike processed, $n = n(t_{old})$
Ensure: $t_{old} = t$, $n = n(t)$ on exit
  for all postsynaptic neurons j **do**
    if mature_j **then**
      send spike(w_j, d_j) to neuron j
    end if
    history $\leftarrow j$.`get_history`$(t_{old} - d_j^D + d_j^A + \tau_{rise,nmda}, t - d_j^D + d_j^A + \tau_{rise,nmda})$
    for all spikes $t_j \in$ history **do**
      $t_{BP} \leftarrow t_j + d_j^D$
      $n' \leftarrow n\exp\left(-\frac{t_{BP}-t_{old}}{\tau_{nmda}}\right)$
      $x_j \leftarrow \mathrm{T}(x_j, n')$
      $\mathrm{mature}_j \leftarrow \mathrm{mature}_j \vee (x_j > X_m)$ {synapse maturates, if $x_j > X_m$}
      if $\mathrm{mature}_j \wedge (x_j < X_d)$ **then**
        j.`stop_recording`$(t - d_j^D + d_j^A + \tau_{rise,nmda})$
        delete synapse j {mature synapse dies, if $x_j < X_d$}
      end if
    end for
  end for
  $n \leftarrow n\exp\left(\frac{t-t_{old}}{\tau_{nmda}}\right) + 1$
  $t_{old} = t$

Algorithm 2 T(x, n)

Require: `Binomial`$(N, p) \in [0, N]$ binomially distributed random number
  if $n > \Theta_h$ **then**
    $x \leftarrow x +$ `Binomial`$(N - x, p)$
  else if $n > \Theta_b \wedge n <= \Theta_l$ **then**
    $x \leftarrow x -$ `Binomial`(x, q)
  end if
  return x

Algorithm 3 start_recording(t_{first})

  start at beginning of spike register
  while $t_{SP} <= t_{first}$ **do**
    $\mathrm{counter}_{SP} \leftarrow \mathrm{counter}_{SP} + 1$
    move to next element
  end while

Algorithm 4 stop_recording(t_{last})

start at beginning of spike register
while $t_{\text{SP}} <= t_{\text{last}}$ **do**
  $\text{counter}_{\text{SP}} \leftarrow \text{counter}_{\text{SP}} - 1$
  move to next element
end while

closely follow Morrison et al. (2007) and use the same definition of `update_register`(t) and `get_history`(t_1, t_2). An extension, however, is necessary, because in the presence of structural plasticity the number of incoming connections of a neuron may change over time: The access counter $\text{counter}_{\text{SP}}$ of each entry in the spike register has to be adapted whenever a synapse is created or dies. We define the method `start_recording`(t_{first}) (see Algorithm 3) called on the postsynaptic neuron when a new synapse is created. The argument is the earliest $t_{\text{first}} = t_{\text{old}} - d_j^D + d_j^A + \tau_{\text{rise,nmda}}$ the new synapse will start accessing the spike history (excluding t_{first}). Analogously, `stop_recording`(t_{last}) is called when a synapse dies (see Algorithm 4). Its argument $t_{\text{last}} = t - d_j^D + d_j^A + \tau_{\text{rise,nmda}}$ is the latest point in the history (including t_{last}) considered by the synapse before dawn (see Algorithm 1).

ACKNOWLEDGMENTS

Many thanks to Birgit Kriener for very helpful discussions and to Tom Tetzlaff for comments on an earlier version of the manuscript. Ursula Körner and Andreas Knoblauch contributed essential inspiring ideas in the initial phase of the project. The implementation of the synapse model is the result of constructive discussions with the NEST initiative (www.nest-initiative.org). Partially funded by DIP F1.2, BMBF Grant 01GQ0420 to the Bernstein Center for Computational Neuroscience Freiburg, and EU Grant 15879 (FACETS), and "The Next-Generation Integrated Simulation of Living Matter" project, part of the Development and Use of the Next-Generation Supercomputer Project of the Ministry of Education, Culture, Sports, Science and Technology (MEXT) of Japan.

REFERENCES

Abramowitz, M., and Stegun, I. A. (1974). *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables.* New York, Dover Publications.

Atwood, H. L., and Wojtowicz, M. J. (2004). Silent synapses in neural plasticity: current evidence. *Learn. Mem.* 6, 542–571.

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

Bonhoeffer, T., and Yuste, R. (2002). Spine motility. Phenomenology, mechanisms, and function. *Neuron* 35, 1019–1027.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* 8, 183–208.

Brunel, N., and Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.* 11, 1621–1671.

Butz, M., Teuchert-Noodt, G., Grafen, K., and van Ooyen, A. (2008). Inverse relationship between adult hippocampal cell proliferation and synaptic rewiring in the dentate gyrus. *Hippocampus* 18, 879–898.

Cai, Y., Gavornik, J. P., Cooper, L. N., Yeung, L. C., and Shouval, H. Z. (2007). Effect of stochastic synaptic and dendritic dynamics on synaptic plasticity in visual cortex and hippocampus. *J. Neurophysiol.* 97, 375–386.

Cavazzini, M., Bliss, T., and Emptage, N. (2005). Ca^{2+} and synaptic plasticity. *Cell Calcium* 38, 355–367.

Chen, X., Vinade, L., Leapman, R. D., Nakagawa, T., Phillips, T. M., Sheng, M., and Reese, T. S. (2005). Mass of the postsynaptic density and enumeration of three key molecules. *Proc. Natl. Acad. Sci. U.S.A.* 102, 11551–11556.

Chklovskii, D. B., Mel, B. W., and Svoboda, K. (2004). Cortical rewiring and information storage. *Nature* 431, 782–788.

Cohen-Cory, S. (2002). The developing synapse: construction and modulation of synaptic structures and circuits. *Science* 298, 770–776.

Cormier, R. J., Greenwood, A. C., and Connor, J. A. (2001). Bidirectional synaptic plasticity correlated with the magnitude of dendritic calcium transients above a threshold. *J. Neurophysiol.* 85, 399–406.

Dammasch, I. E., Wagner, G. P., and Wolff, J. R. (1986). Self-stabilization of neuronal networks. *Biol. Cybern.* 54, 211–222.

De la Rocha, J., Doiron, B., Shea-Brown, E., Kresimir, J., and Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature* 448, 802–807.

Gerstner, W., and Kistler, W. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity.* Cambridge, Cambridge University Press.

Gewaltig, M.-O., and Diesmann, M. (2007). NEST (Neural Simulation Tool). Scholarpedia 2, 1430.

Gilbert, E. N., and Pollak, H. O. (1960). Amplitude distribution of shot noise. *Bell Syst. Tech. J.* 39, 333–350.

Graupner, M., and Brunel, N. (2007). STDP in a bistable synapse model based on CaMKII and associated signaling pathways. *PLoS Comput. Biol.* 3, e221.

Helias, M., Rotter, S., Gewaltig, M.-O., and Diesmann, M. (2007). A model for correlation detection based on Ca^{2+} concentration in spines. In *Sixteenth Annual Computational Neuroscience Meeting CNS 2007.* W. R. Holmes (ed.), Vol. 8, BMC *Neuroscience*, pp. 192.

Jahr, C. E., and Stevens, C. F. (1990a). A quantitative description of NMDA-receptor-channel kinetic behavior. *J. Neurosci.* 10, 1830–1837.

Jahr, C. E., and Stevens, C. F. (1990b). Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics. *J. Neurosci.* 10, 3178–3182.

Jones, E., Oliphant, T., Peterson, P., et al. (2001). *SciPy: Open Source Scientific Tools for Python.* Available at: http://www.scipy.org/.

Kalisman, N., Silberberg, G., and Markram, H. (2005). The neocortical circuit as a tabula rasa. *Proc. Natl. Acad. Sci. U.S.A.* 102, 880–885.

Kampa, B., Clements, J., Jonas, P., and Stuart, G. (2004). Kinetics of Mg^{2+} unblock of NMDA receptors: implications for spike-timing dependent synaptic plasticity. *J. Physiol.* 556, 337–345.

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

Kempter, R., Gerstner, W., and van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Comput.* 12, 2709–2742.

Knoblauch, A. (2006). *On Compressing the Memory Structures of Binary Neuronal Associative Networks.* Technical Report HRI-EU 06-02, Honda Research Institute Europe GmbH, Offenbach.

Knoblauch, A., Sommer, F., Gewaltig, M.-O., Kupper, R., Körner, U., and Körner, E. (2007). A model for structural plasticity in neocortical associative networks trained by the hippocampus. *BMC Neurosci.* 8, S14.

Kuhn, A., Aertsen, A., and Rotter, S. (2003). Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Comput.* 1, 67–101.

Le Be, J.-V., and Markram, H. (2006). Spontaneous and evoked synaptic rewiring in the neonatal neocortex. *Proc. Natl. Acad. Sci. U.S.A.* 103, 13214–13219.

Lisman, J., Schulman, H., and Cline, H. (2002). The molecular basis of CaMKII function in synaptic and behavioral memory. *Nat. Rev. Neurosci.* 2, 175–190.

Lüscher, C., Noicoll, R., Malenka, R., and Muller, D. (2000). Synaptic plasticity and dynamic modulation of the postsynaptic membrane. *Nat. Neurosci.* 3, 545–550.

MacCluer, C. R. (2000). The many proofs and applications of Perron's theorem. *SIAM Rev.* 42, 487–498.

Mainen, Z., Malinow, R., and Svoboda, K. (1999). Synaptic calcium transients in single spines indicate that NMDA receptors are not saturated. *Nature* 399, 151–155.

Miller, P., Zhabotinsky, A. M., Lisman, J. E., and Wang, X.-J. (2005). The stability of a stochastic CaMKII switch: dependence on the number of enzyme molecules and protein turnover. *PLoS Biol.* 3, e107.

Morrison, A., Aertsen, A., and Diesmann, M. (2007). Spike-timing dependent plasticity in balanced random networks. *Neural Comput.* 19, 1437–1467.

Morrison, A., Diesmann, M., and Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike-timing. *Biol. Cybern.* 98, 459–478.

Morrison, A., Mehring, C., Geisel, T., Aertsen, A., and Diesmann, M. (2005). Advancing the boundaries of high connectivity network simulation with distributed computing. *Neural Comput.* 17, 1776–1801.

Nevian, T., and Sakmann, B. (2004). Single spine Ca^{2+} signals evoked by coincident EPSPs and backpropagating action potentials in spiny stellate cells of layer 4 in the juvenile rat somatosensory barrel cortex. *J. Neurosci.* 24, 1689–1699.

Nevian, T., and Sakmann, B. (2006). Spine ca signaling in spike-timing-dependent plasticity. *J. Neurosci.* 26, 11001–11013.

Ricciardi, L. M. (1977). *Diffusion Processes and Related Topics on Biology.* Berlin, Springer-Verlag.

Ricciardi, L. M., Di Crescenzo, A., Giorno, V., and Nobile, A. G. (1999). An outline of theoretical and algorithmic approaches to first passage time problems with applications to biological modeling. *Math. Jpn.* 50, 247–322.

Risken, H. (1996). *The Fokker–Planck Equation*. Berlin, Heidelberg, Springer-Verlag.

Rubin, D. D. B. D., and Fusi, S. (2007). Long memory lifetimes require complex synapses and limited sparseness. *Front. Comput. Neurosci.* 1, 1–14.

Shi, S.-H., Hayashi, Y., Petralia, R. S., Zaman, S. H., Wenthold, R. J., Svoboda, K., and Malinow, R. (1999). Rapid spine delivery and redistribution of AMPA receptors after synaptic NMDA receptor activation. *Science* 284, 1811–1816.

Shouval, H. Z., Bear, M. F., and Cooper, L. N. (2002). A unified model of NMDA receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. U.S.A.* 99, 10831–10836.

Shouval, H. Z., and Kalantzis, G. (2005). Stochastic properties of synaptic transmission affect the shape of spike time dependent plasticity curves. *J. Neurophysiol.* 93, 1069–1073.

Stepanyants, A., Hirsch, J., Martinez, L. M., Kisvarday, Z. F., Ferecsko, A. S., and Chklovskii, D. B. (2007). Local potential connectivity in cat primary visual cortex. *Cereb. Cortex* 18, 13–28.

Stepanyants, A., Hof, P. R., and Chklovskii, D. B. (2002). Geometry and structural plasticity of synapse connectivity. *Neuron* 34, 275–288.

Stettler, D. D., Yamahachi, H., Li, W., Denk, W., and Gilbert, C. D. (2006). Axons and synaptic boutons are highly dynamic in adult visual cortex. *Neuron* 49, 877–887.

van Ooyen, A., van Pelt, J., and Corner, M. A. (1995). Implications of activity dependent neurite outgrowth for neuronal morphology and network development. *J. Theor. Biol.* 172, 63–82.

Waters, J., Larkum, M., Sakmann, B., and Helmchen, F. (2003). Supralinear Ca^{2+} influx into dendritic tufts of layer 2/3 neocortical pyramidal neurons *in vitro* and *in vivo*. *J. Neurosci.* 23, 8558–8567.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 12 August 2008; paper pending published: 13 October 2008; accepted: 02 December 2008; published online: 24 December 2008.

*Citation: Helias M, Rotter S, Gewaltig M-O and Diesmann M (2008) Structural plasticity controlled by calcium based correlation detection. Front. Comput. Neurosci. (2008) **2**:7. doi: 10.3389/neuro.10.007.2008*