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 $[\bar{u}(t)-\vartheta_{\rm LTD}][u-\vartheta_{\rm LTP}]$

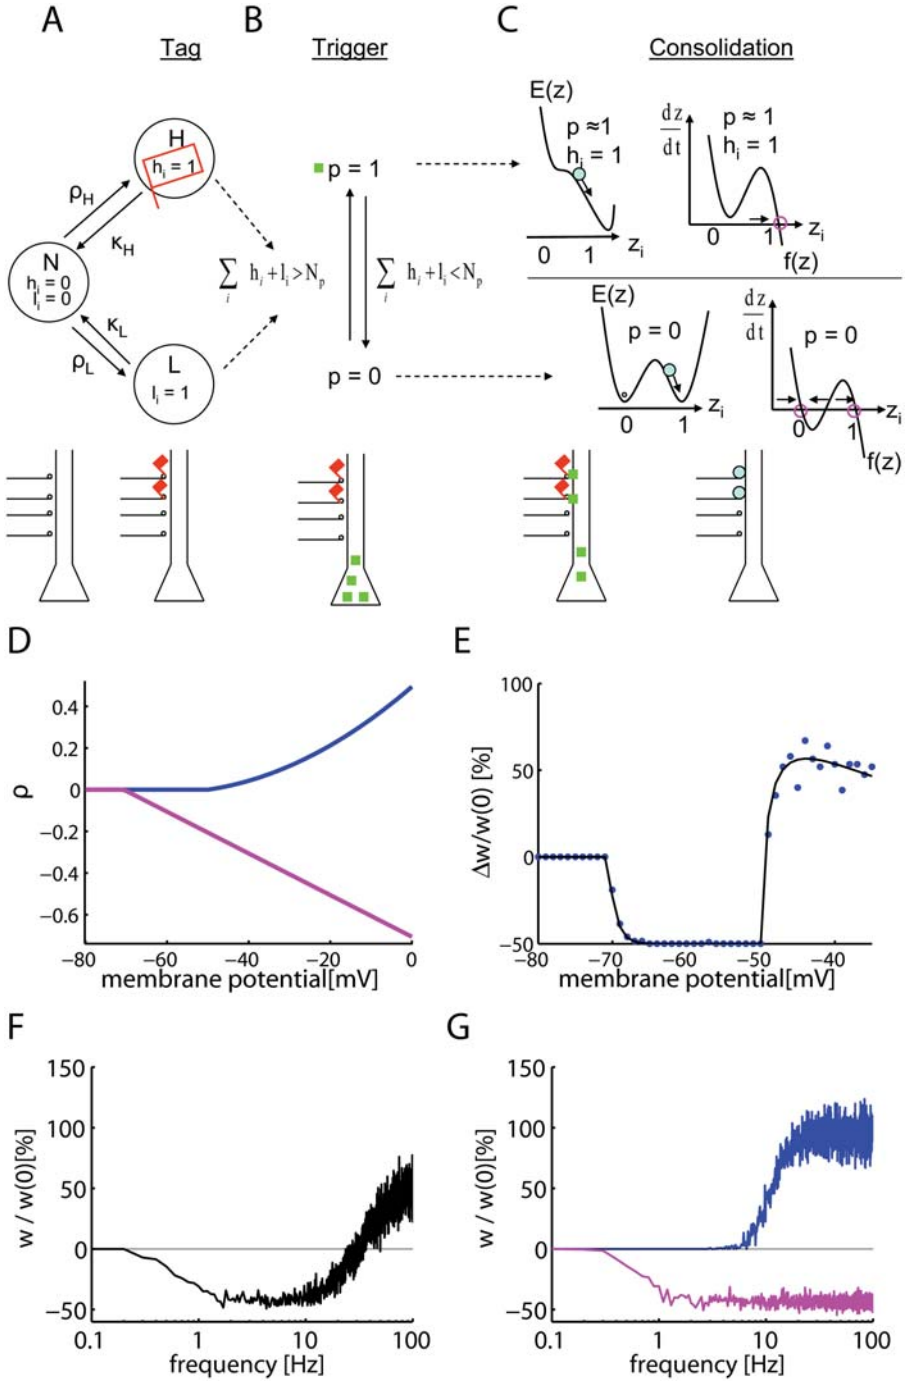

Figure 1. The three components of the Tag-Trigger-Consolidation (TagTriC) model. (A) A synapse can be in the non-tagged state N, the high state H or the low state L. A synapse i in H (or L) has a tag $h_i = 1$ (or $l_i = 1$, respectively). Transitions to a tagged state occur with rates ρ_H for potentiation 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.

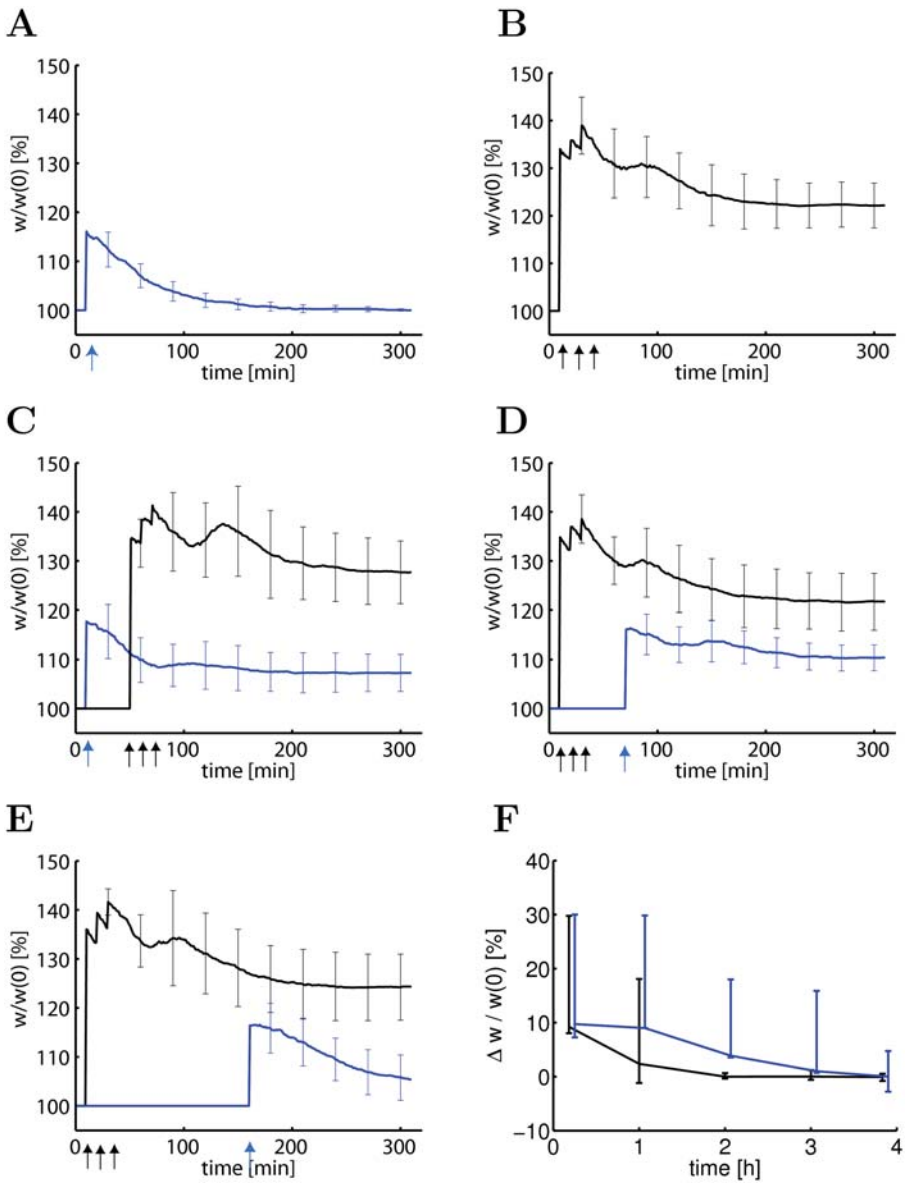

Figure 2. The model accounts for tagging paradigms. (A) A weak tetanus (21 pulses at 100 Hz) applied at a group of 100 synapses at $t = 10$ min (arrow) leads to an increased connection weight ($w/w(0)$, blue line) that decays back to baseline. (B) A strong tetanus (100 pulses at 100 Hz repeated three 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

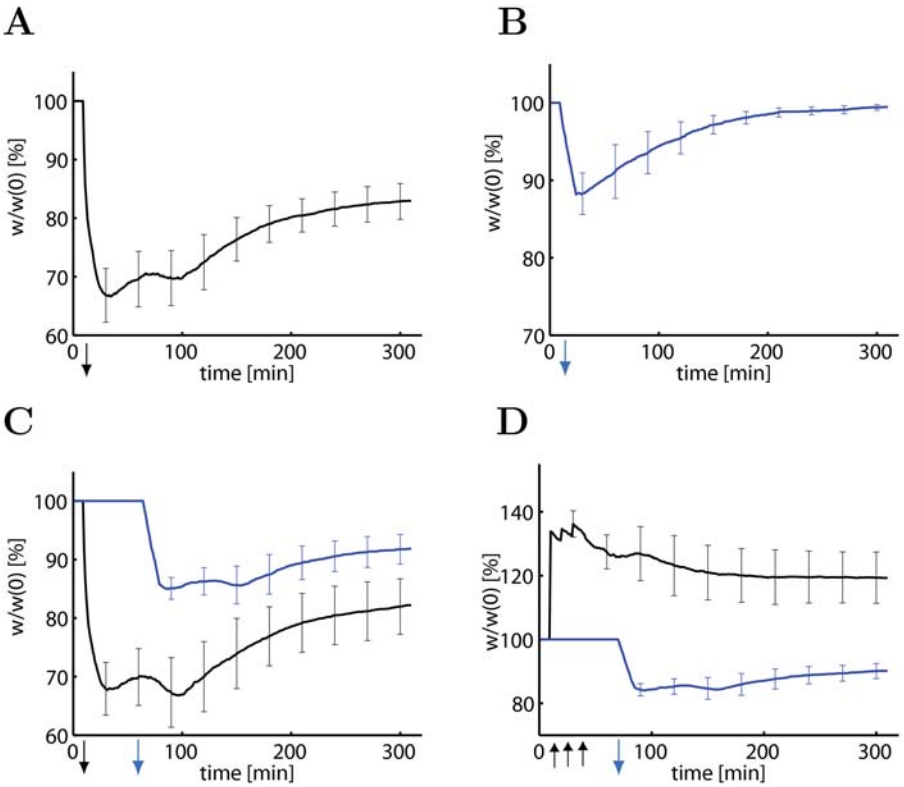

Figure 3. The model accounts for cross-tagging between LTP and LTD. (A) A strong low-frequency stimulus (3 pulses at 20 Hz, repeated 900 times every second) applied to a group of $N = 100$ synapses induces LTD with mean weights ($w/w(0)$) stabilized at $83 \pm 3\%$ of initial value after 5 hours (black 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

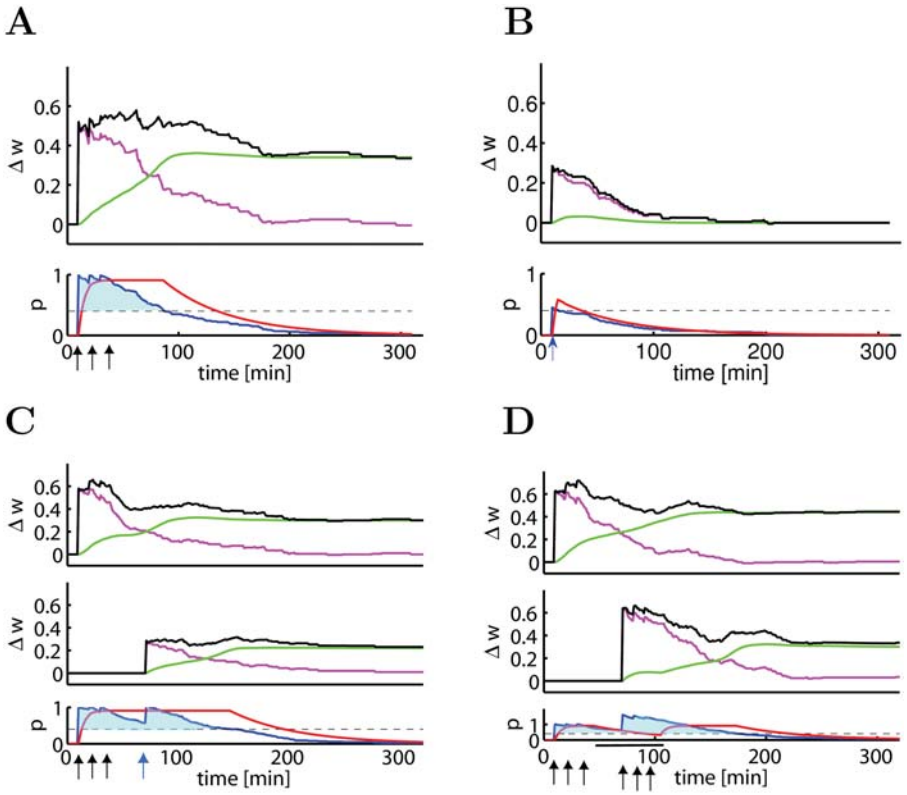

Figure 4. Dynamics of the TagTriC Model during different tagging protocols and protein synthesis blocking. The change of the total synaptic weight (top panels, black line $\Delta w = \sum_{i=1}^{N} [w_i(t) - w_i(0)/N]$) has contribution from early LTP (top panels, magenta line represents $\sum_{i=1}^{N} (h_i - \alpha l_i/N)$) 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

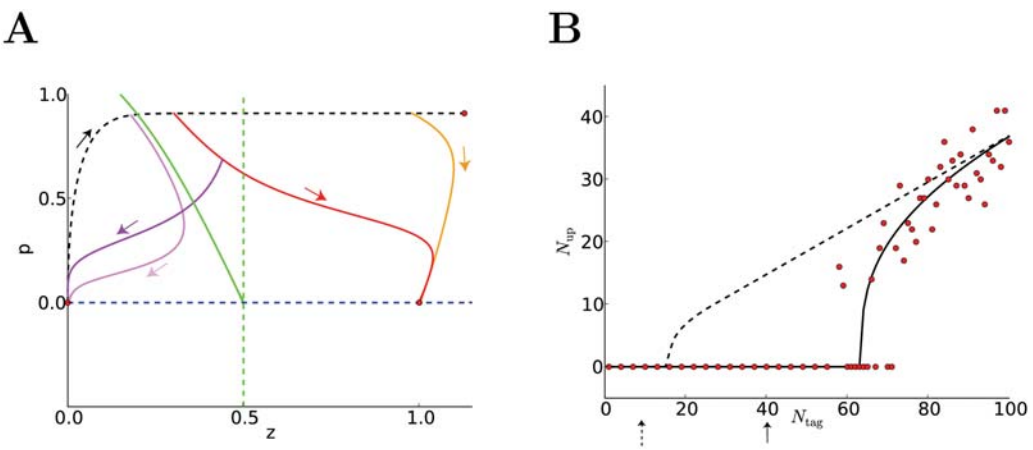

Figure 5. Theory and predictions. (A) Evolution of the variables p and z during tagging. If protein synthesis is 'ON' and the synapse tagged, p and z move along the black dashed line towards the stable fixed point on the upper right ($p \approx 1$, $z \approx 1$) (red filled circle). If protein synthesis 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 $N_p = n_0/(\mathrm{DA_{bg}}+c_0)$ 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 $0<\mathrm{DA_{bg}}<1$. 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 $\left(\sum_i^N h_i + l_i\right)\left(\mathrm{DA_{bg}} + c_0\right) > n_0$ 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 $\vartheta_{\rm LTD}$ we get a rate of LTD

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

where $A_{\rm LTD}>0$ 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 $P_{l=0\to l=1}=1-\exp(-A_{\rm LTD}[\bar{u}(t)-\vartheta_{\rm LTD}]^+\Delta)\approx A_{\rm LTD}[\bar{u}(t)-\vartheta_{\rm LTD}]^+\Delta$ 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

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

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 $\vartheta_{\rm LTP}$. 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 $A_{\rm LTP}$ is not the same as $A_{\rm LTD}$. 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 $\vartheta_{\rm LTP}$; 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 $0\le\mathrm{DA_{bg}}\le1$ 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 $t_1 \approx 28$ min and $t_2 \approx 60$ 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.

nature neuroscience

Connectivity reflects coding: a model of voltage-based STDP with homeostasis

Claudia Clopath[1], Lars Büsing[1,2], Eleni Vasilaki[1,2] & Wulfram Gerstner[1]

Electrophysiological connectivity patterns in cortex often have a few strong connections, which are sometimes bidirectional, among a lot of weak connections. To explain these connectivity patterns, we created a model of spike timing–dependent plasticity (STDP) in which synaptic changes depend on presynaptic spike arrival and the postsynaptic membrane potential, filtered with two different time constants. Our model describes several nonlinear effects that are observed in STDP experiments, as well as the voltage dependence of plasticity. We found that, in a simulated recurrent network of spiking neurons, our plasticity rule led not only to development of localized receptive fields but also to connectivity patterns that reflect the neural code. For temporal coding procedures with spatio-temporal input correlations, strong connections were predominantly unidirectional, whereas they were bidirectional under rate-coded input with spatial correlations only. Thus, variable connectivity patterns in the brain could reflect different coding principles across brain areas; moreover, our simulations suggested that plasticity is fast.

Experience-dependent changes in receptive fields[1,2] or in learned behavior relate to changes in synaptic strength. Electrophysiological measurements of functional connectivity patterns in slices of neural tissue[3,4] or anatomical connectivity measures can only present a snapshot of the momentary connectivity, which may change over time[5]. The question then arises of whether the connectivity patterns and their changes can be connected to basic forms of synaptic plasticity[6] such as long-term potentiation (LTP) and depression (LTD)[7]. LTP and LTD depend on the exact timing of pre- and postsynaptic action potentials[2–8] but also on postsynaptic voltage[9,10] and presynaptic stimulation frequency[11]. STDP has attracted particular interest in recent years, as temporal coding schemes in which information is contained in the exact timing of spikes rather than mean frequency can be learned by a neural system using STDP[12,13] (review in ref. 14). However, the question of whether STDP is more fundamental than frequency-dependent plasticity or voltage-dependent plasticity rules has not been resolved despite an intense debate[15]. Moreover, it is unclear how the interplay of coding and plasticity yields the functional connectivity patterns that are 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[3] or somatosensory cortex[4]. Recent experiments have shown that STDP is strongly influenced by postsynaptic voltage before action potential firing[16] but were unable to answer the question of 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[16], suggesting an interaction between timing- and frequency-dependent processes[16].

We found that a simple Hebbian plasticity rule that pairs presynaptic spike arrival with the postsynaptic membrane potential was sufficient to explain STDP and the dependence of plasticity on presynaptic stimulation frequency. Moreover, the intricate interplay of voltage and spike timing as well as the frequency dependence of STDP can be explained in our model from one single principle. In contrast with earlier attempts towards a unified description of synaptic plasticity[17,18], our model is a phenomenological one. It does not give an explicit interpretation in terms of biophysical quantities such a calcium concentration[17], CaMKII[18], glutamate binding, NMDA receptors, etc. Instead, 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[19–21], and possibly even large[22,23], networks. We found that the learning rule led to input specificity in small networks of up to ten neurons, which is necessary for receptive field development, similar to earlier models of STDP[12,19] or rate-based plasticity rules[24,25]. We explicitly addressed the question of whether functional connectivity patterns of cortical pyramidal neurons measured in recent electrophysiological studies[3,4] could be the result of plasticity during continued stimulation of neuronal model networks, particularly bidirectional connections[3] that are incompatible with standard STDP models[12,19]. The mathematical simplicity of our model enabled us to identify conditions under which it becomes equivalent to the well-known Bienenstock-Cooper-Munro (BCM) model[24] used in classical rate-based descriptions of developmental learning and, similar to some earlier models of STDP[26,27], and why our model is fundamentally different from classical STDP models[12,14,19], is widely used for temporal coding.

RESULTS

To study the means by which connectivity patterns in cortex can emerge from plasticity, we needed a plasticity rule that was consistent

[1]Laboratory of Computational Neuroscience, Brain-Mind Institute and School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland. [2]Present address: Institut für Grundlagen der Informationsverarbeitung, Graz University of Technology, Austria (L.B.), and Department of Computer Science, University of Sheffield, Sheffield, UK (E.V.). Correspondence should be addressed to C.C. (claudia.clopath@epfl.ch).

Received 19 October 2009; accepted 1 December 2009; published online 24 January 2010; doi:10.1038/nn.2479

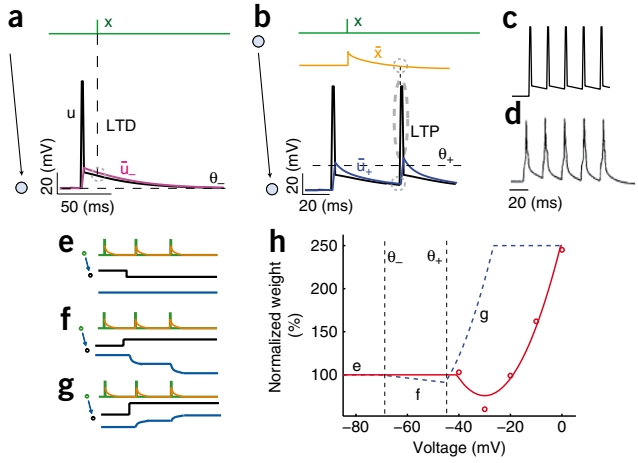

Figure 1 Illustration of the model. Synaptic weights react to presynaptic events (top) and postsynaptic membrane potential (bottom). (**a**) The synaptic weight was decreased if a presynaptic spike x (green) arrived when the low-pass-filtered value $\bar{u}_-$ (magenta) of the membrane potential was above θ_- (dashed horizontal line). (**b**) The synaptic weight was increased if the membrane potential u (black) was above a threshold θ_+ and the low-pass-filtered value of the membrane potential $\bar{u}_+$ (blue) was higher than a threshold θ_- and the presynaptic low-pass filter $\bar{x}$ (orange) was nonzero. (**c**) Step current injection made the postsynaptic neuron fire at 50 Hz in the absence of presynaptic stimulation (membrane potential u in black). No weight change was observed. Note the depolarizing spike afterpotential, consistent with experimental data. (**d**) Reproduced from ref. 16. (**e**–**h**) Voltage-clamp experiment. A neuron received weak presynaptic stimulation of 2 Hz during 50 s while the postsynaptic voltage was clamped to values between –60 mV and 0 mV. (**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 hyperpolarization (**e**), slight depolarization (**f**) and large depolarization (**g**). (**h**) The weight change as a function of clamped voltage using the standard set of parameters for visual cortex data (dashed blue line, voltage paired with 25 spikes at the synapse). With a different set of parameters, the model fit the experimental data (red circles) in hippocampal slices[10] (see Online Methods for details).

with a large body of experimental data. Because synaptic depression and potentiation occur via different pathways[28], our model used separate additive contributions to the plasticity rule, one for LTD and another one for LTP (see **Fig. 1** and Online Methods).

Fitting the plasticity model to experimental data

Consistent with voltage-clamp[10] and stationary-depolarization experiments[9], 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_- that is usually set to resting potential), whereas LTP occurs if depolarization is big (above a second threshold θ_+; **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 ($\bar{u}_-$ 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 $\bar{x}$ 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 $\bar{u}_+$ is above θ_- (this is fulfilled if there was 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; **Fig. 1b**). LTD occurs if the average voltage $\bar{u}_-$ is above θ_- at the moment of a presynaptic spike arrival (**Fig. 1a**). The amount of LTD in our model depended on a homeostatic process on a slower time scale[29]. Low-pass filtering of the voltage by the variable ($\bar{u}_-$or $\bar{u}_+$) refers to some unidentified intracellular processes triggered by depolarization, for example, increase in calcium concentration or second 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 used a STDP protocol in which presynaptic spikes arrive a few milliseconds before or after a postsynaptic spike (**Fig. 2** and **Supplementary Methods**). If a post-pre pairing with a timing difference of 10 ms was repeated at frequencies below 35 Hz, LTD occurred

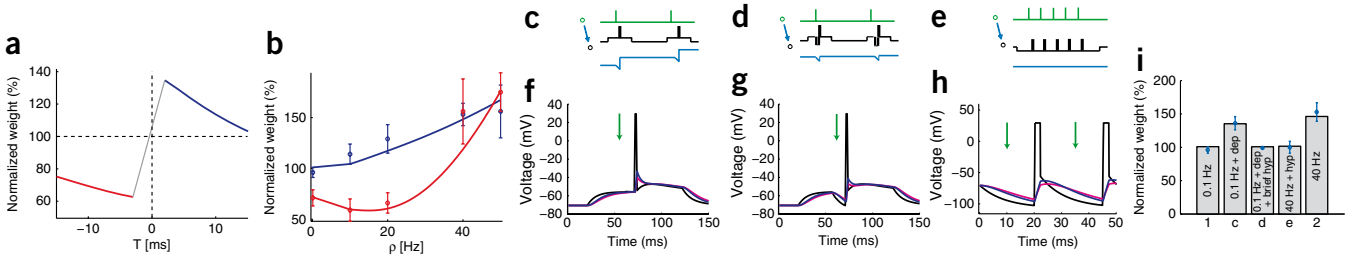

Figure 2 Fitting the model to experimental data. (**a**,**b**) Simulated STDP experiments. (**a**) Spike timing–dependent learning window: synaptic weight change for different time intervals T between pre- and postsynaptic firing using 60 pre-post-pairs at 20 Hz. (**b**) Weight change as a function of pairing repetition frequency ρ using pairings with a time delay of +10 ms (pre-post, blue) and –10 ms (post-pre, red). Dots represent data from ref. 16 and lines represent our plasticity model. (**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[16]. Low-frequency (0.1 Hz) pre-post spike pairs yielded LTP if a 100-ms-long depolarized current was injected around the pairing. (**d**) LTP failed if an additional brief hyperpolarized pulse was applied 14 ms before postsynaptic firing so that voltage is brought to rest. (**e**) Hyperpolarization preceding action potential prevents potentiation that normally occurred at 40 Hz[16]. (**f**) The simulated postsynaptic voltage u (black) is shown after using the protocol described in **c**, together with temporal averages $\bar{u}_-$ (magenta) and $\bar{u}_+$ (blue). The presynaptic spike time is indicated by the green arrow. Using the model (equation (3)) with this setting yielded potentiation. (**g**) Data are presented as in **f** but using the protocol described in **d**. No weight change was measured. (**h**) Data are presented as in **f** but using the protocol described in **e**. No weight change was measured. (**i**) Histogram summarizing the normalized synaptic weight of the simulation (bar) and the experimental data[16] (dot, blue bar indicates variance) for 0.1-Hz pairing (control 1), 0.1-Hz pairing with the depolarization (protocol **c**), 0.1-Hz pairing with the depolarization and brief hyperpolarization (protocol **d**), 40-Hz pairing (control 2), and 40-Hz pairing with the constant hyperpolarization (protocol **e**); parameters are described in **Table 1**.

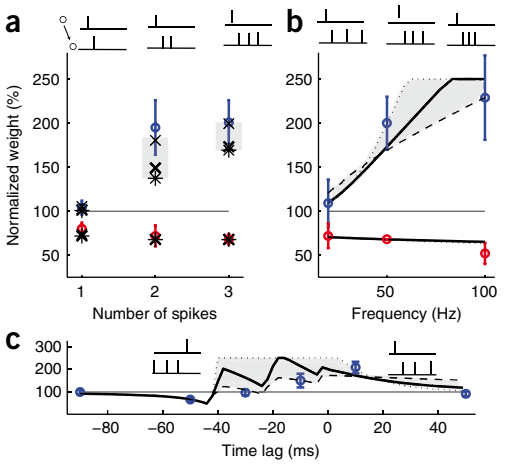

Figure 3 Burst timing–dependent plasticity. One presynaptic spike was paired with a burst of postsynaptic spikes. This pairing was repeated 60 times at 0.1 Hz. (**a**) Normalized weight as a function of the number of postsynaptic spikes (1, 2, 3) at 50 Hz (dots represent data from ref. 30, crosses represent simulation). The presynaptic spike was paired +10 ms before the first postsynaptic spike (blue) or −10 ms after (red). (**b**) Normalized weight as a function of the frequency between the three postsynaptic action potentials (dots indicate data, lines indicate simulation, blue indicates pre-post, red indicates post-pre). (**c**) Normalized weight as a function of the timing between the presynaptic spike and the first postsynaptic spike of a three-spike burst at 50 Hz (dot indicates data, black lines indicate simulation). A hard upper bound was set to 250% normalized weight. The dashed line and crosses and the dotted line and stars represent simulations with alternative sets of parameters, $A_{LTD} = 21 \times 10^{-5}$ mV^{-1}, $A_{LTP} = 50 \times 10^{-4}$ mV^{-2}, $\tau_x = 143$ ms, $\tau_- = 6$ ms, $\tau_+ = 5$ ms and $A_{LTD} = 21 \times 10^{-5}$ mV^{-1}, $A_{LTP} = 67 \times 10^{-4}$ mV^{-2}, $\tau_x = 5$ ms, $\tau_- = 8$ ms, $\tau_+ = 5$ ms, respectively. Shading indicates reachable data points generated by the model with different parameters.

in our model (**Fig. 2a**,**b**), consistent with experimental data[16]. Repeated pre-post pairings (with 10-ms timing difference) at frequencies above 10 Hz yielded LTP, but pairings at 0.1 Hz did not show any change in the model or in experiments[16]. In the model, these results can be explained by the fact that, at a 0.1-Hz repetition frequency, the low-pass-filtered voltage $\bar{u}_+$, which increases abruptly during postsynaptic spiking, decays back to zero before the next impulse arrives; thus, LTP cannot be triggered. However, as 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 a very low frequency. At repetition frequencies of 50 Hz, the post-pre procedure is nearly indistinguishable from a pre-post timing and LTP dominates.

If a pre-post protocol at 0.1 Hz that normally does not induce LTP was combined with a depolarizing current pulse, then potentiation was observed in experiments[16] and in our model (**Fig. 2c**,**f**,**i**). As a result of the injected current, the low-pass-filtered voltage variable $\bar{u}_+$ is depolarized before the pairing. Thus, at 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. 2e**,**h**,**i**).

To study nonlinear aspects of STDP, we simulated a protocol of burst timing–dependent plasticity in which presynaptic spikes are paired with 1–3 postsynaptic spikes[30] (see Online Methods). Although pairings at 0.1 Hz did not change the synaptic weight, repeated triplets pre-post-post generated potentiation in our model, as the first postsynaptic spike induced a depolarizing spike after potential so that $\bar{u}_+$ was depolarized. Adding a third postsynaptic spike to the protocol (that is, quadruplets pre-post-post-post) did not lead to stronger LTP (**Fig. 3a**). Our model also describes the dependence of LTP on the intra-burst frequency (**Fig. 3b**). At an intra-burst frequency of 20 Hz, no LTP occurred because the second spike in the burst came so late that the presynaptic trace $\bar{x}$ had 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**) that occurs when the timing of one presynaptic spike is changed with respect to a burst of three postsynaptic spikes is qualitatively similar to that found experimentally[30,31], but only four of the six experimental data points are quantitatively reproduced by the model with a given set of parameters. Notably, our model predicted that the curve of burst timing–dependent plasticity should show a change in the amount of potentiation whenever the presynaptic spike is shifted across one of the three postsynaptic spikes (**Fig. 3c**). Because dendritic spikes, which are relevant for burst timing–dependent STDP[31], are broader than somatic action potentials, the 'jumps' in the burst-STDP curves would be blurred.

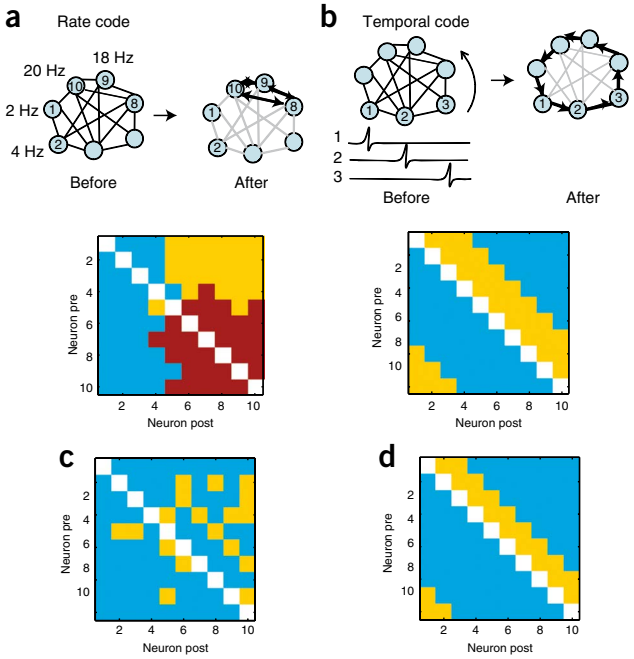

Figure 4 Weight evolution in an all-to-all connected network of ten neurons. (**a**) Rate code. Neurons fired at different frequencies, neuron 1 at 2 Hz, neuron 2 at 4 Hz, neuron 10 at 20 Hz. The weights (bottom) averaged over 100 s indicate that neurons with high firing rates developed strong bidirectional 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 (after). (**b**) Temporal code. Neurons fired successively every 20 ms (neuron 1, followed by neuron 2 20 ms later, followed by neuron 3 20 ms later, etc). Connections (bottom) were 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. (**c**,**d**) Data are presented as in **a** and **b**, but we used a standard STDP rule[12,14,19]. Bidirectional connections are impossible.

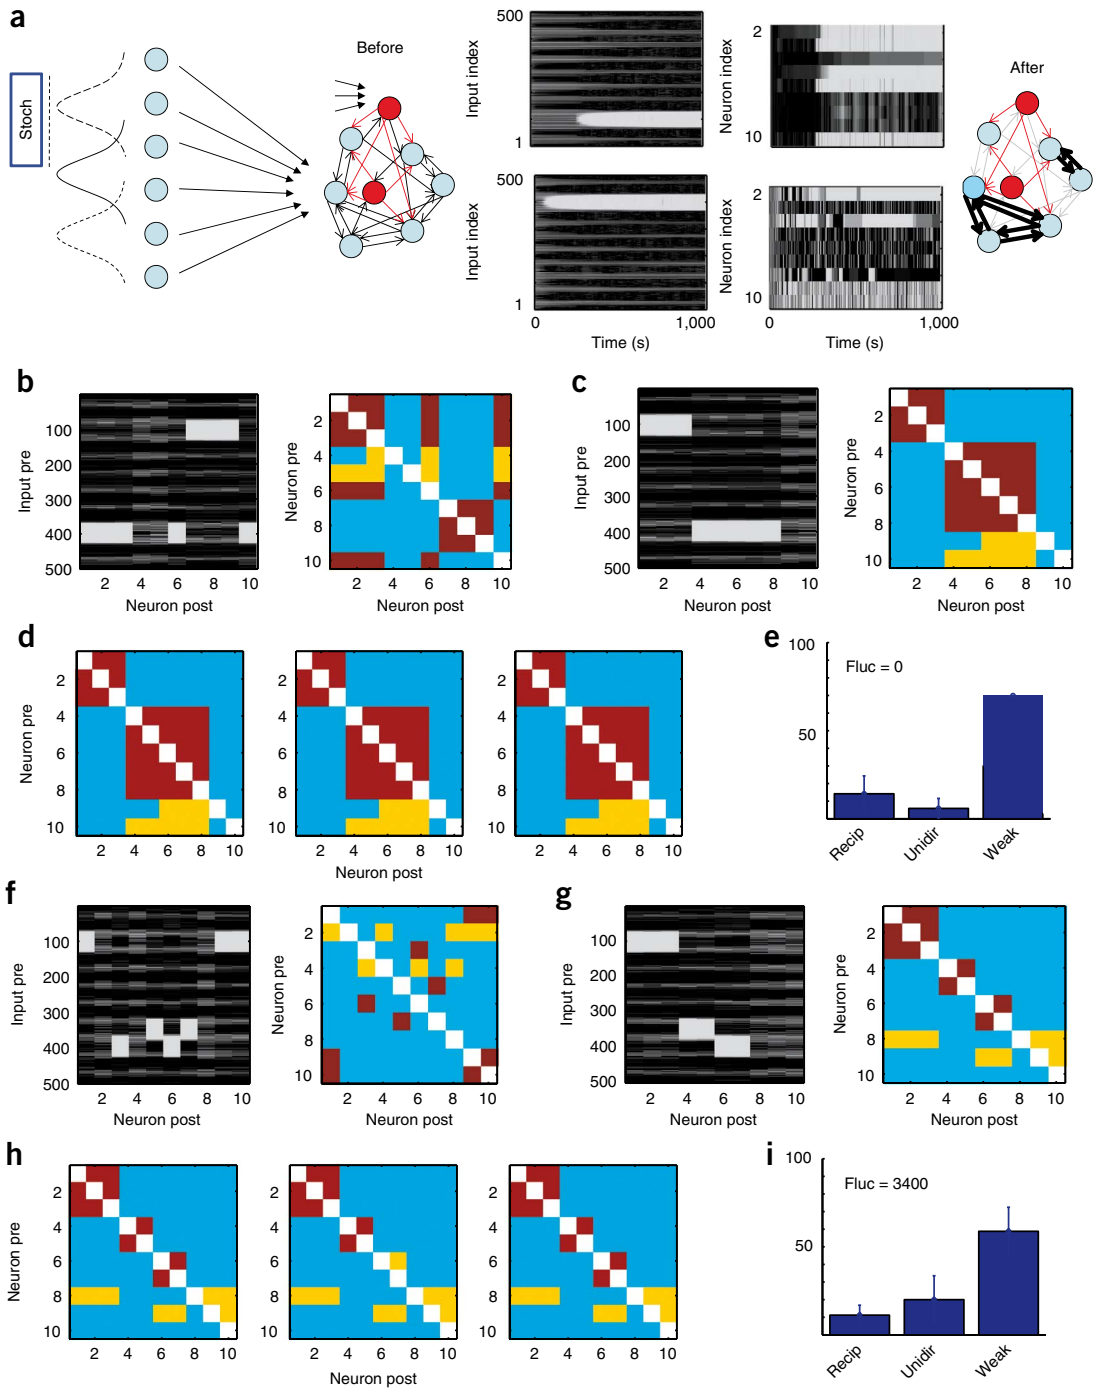

Figure 5 Plasticity during rate coding. (**a**) A network of ten excitatory (light blue) and three inhibitory neurons (red) received feedforward inputs from 500 Poisson spike trains with a Gaussian profile of firing rates. The center of the Gaussian was shifted randomly every 100 ms (schematic network before (left) and after the plasticity experiment (right)). The temporal evolution of the weights are shown (top, small amplitudes of plasticity; bottom, normal amplitudes of plasticity; left, feedforward connections onto neuron 1; right, recurrent connections onto neuron 1). (**b**–**e**) Learning with small amplitudes. We used the parameters detailed in **Table 1b** (visual cortex) except for the amplitudes A_{LTP} and A_{LTD}, which were reduced by a factor 100. (**b**) Mean feedforward weights (left) and recurrent excitatory weights (right) averaged over 100 s. The feedforward weights (left) indicate that the neurons developed localized receptive fields (light gray). The recurrent weights (right) were classified as weak (less than 2/3 of the maximal weight, light blue), strong unidirectional (more than 2/3 of the maximal weight, yellow) or strong reciprocal (brown) connections. The diagonal is white, as self-connections do not exist in the model. (**c**) Data are presented as in **b**, but the neuron index was reordered. (**d**) Three snapshots of the recurrent connections taken 5 s apart indicate that recurrent connections were stable. (**e**) Histogram of reciprocal, unidirectional and weak connections in the recurrent network averaged over 100 s, as shown in **b** (fluc, fluctuations). The total number of weight fluctuations during 100 s was zero. The histogram shows an average of ten repetitions (error bars represent s.d.). (**f**–**i**) Rate code during learning with normal amplitudes. We used the network described above but with a standard set of parameters (**Table 1b**, visual cortex). (**f**) Receptive fields were localized. (**g**) Reordering showed clusters of neurons with bidirectional coupling. These clusters were stable when averaged over 100 s. (**h**) Connections were able to change from one time step to the next. (**i**) The percentage of reciprocal connections was high, but because of fluctuations, more than 1,000 transitions between strong unidirectional to strong bidirectional or back occurred during 100 s.

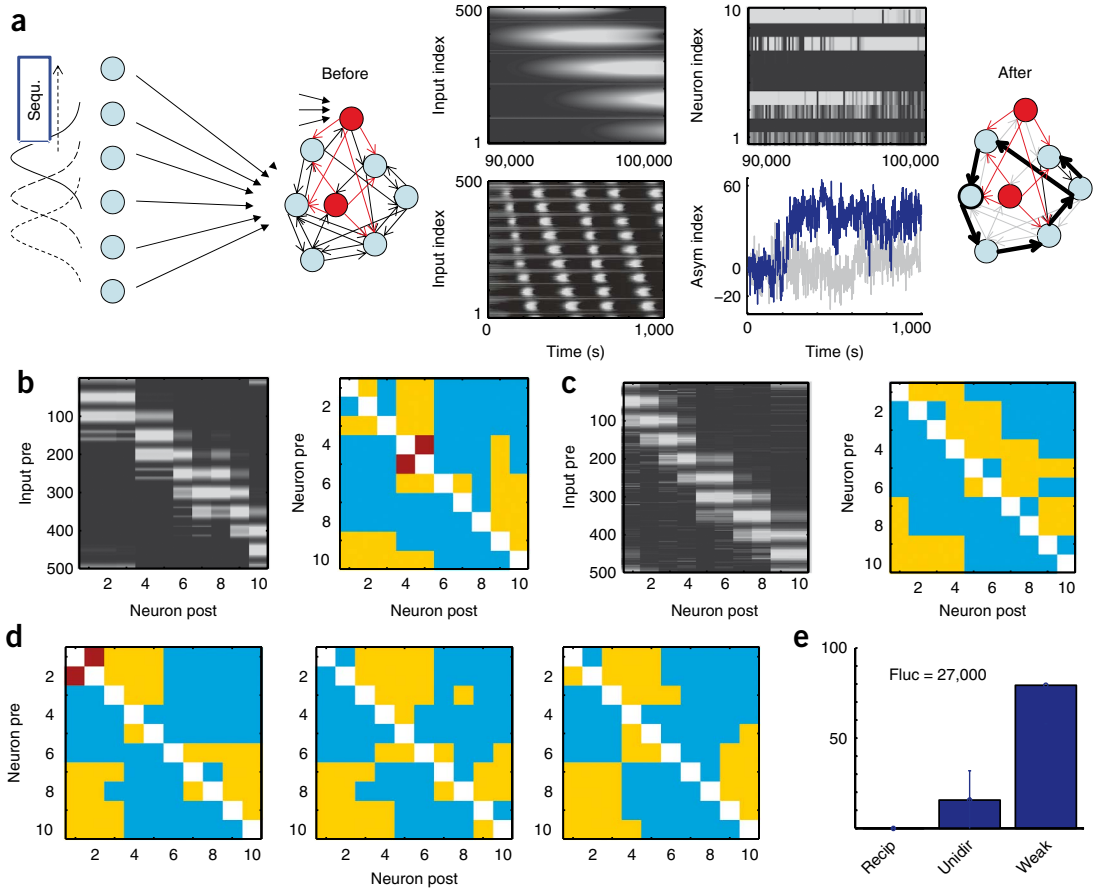

Figure 6 Temporal-coding procedure. The same parameters were used as in **Figure 5** (**Table 1b**, visual cortex), but input patterns were moved successively every 20 ms, corresponding to a step-wise motion of the Gaussian stimulus profile across the input neurons. (**a**) The schematic figure shows the network before and after the plasticity experiment. Shown are the temporal evolution of the weights (top panels: amplitude of synaptic plasticity for feedforward connections reduced by a factor of 100; left, feedforward weights onto neuron 6; right, lateral connections onto neuron 6; bottom panels: normal amplitude of plasticity; left, feedforward connections onto neuron 1; right, temporal evolution of asymmetry index of connection pattern (gray line indicates asymmetrical index for simulation; **Fig. 5**). Positive values indicate the weights from neurons n to $n + k$ are stronger than those from n to $n - k$ for $1 \leq k \leq 3$). (**b**) Receptive fields are localized (left). The recurrent network developed a ring-like structure with strong unidirectional connections from neuron 8 (vertical axis) to neurons 9 and 10 (horizontal axis), etc. (small amplitudes of plasticity). (**c**) Data are presented as in **b**, but normal plasticity values were used. (**d**) Some of the strong unilateral connections appeared or disappeared from one time step to the next, but the ring-like network structure persisted, as the lines just above the diagonal are much more populated than the line below the diagonal. (**e**) Reciprocal connections are absent, but unidirectional connections fluctuated several times between weak and strong during 100 s.

Functional implications

Connectivity patterns in a local cortical circuit have been shown to be nonrandom; that is, the majority of connections are weak and the rare strong ones have a high probability of being bidirectional[3]. However, standard models of STDP[14] do not exhibit stable bidirectional connections[19,32]. Intuitively, if cell A fires before 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. It is therefore impossible to strengthen both connections at the same time. Moreover, to assure long-term stability of firing rates, parameters in standard STDP rules are typically chosen such that inhibition slightly dominates excitation[14], which implies that random spike firing decreases connections. However, the nonlinear aspects of plasticity in our model changed such a simple picture. From the results (**Figs. 2b** and **3b**), we expect that our model could develop stable bidirectional connections at higher neuronal firing rates, in contrast with standard STDP rules.

We first simulated a small network of ten all-to-all connected neurons in which each neuron fired at a fixed frequency, but the frequency varied across neurons. We found that bidirectional connections were formed only between those neurons that both fired at a high rate (**Fig. 4a**). In a second simulation, the neurons in the same network were stimulated cyclically such that they fired in a distinct temporal order (1, 2 , 3,…). In this case, the weights form, after a period of synaptic plasticity, a loop in which strong connections from 1 to 2, 2 to 3… develop but bidirectional connections do not (**Fig. 4b**). These results contrast with those of simulation experiments using a standard STDP rule, where connections are always unidirectional, independently of the stimulation procedure (**Fig. 4c,d**). Theoretical arguments (**Supplementary Methods**) indicate that bidirectional connections cannot exist under the cyclic temporal stimulation procedure (for standard STDP or for our plasticity model). Bidirectional connections did develop in our nonlinear voltage-dependent plasticity model under the assumption of slowly varying rates, in contrast with standard STDP (**Fig. 4c,d**).

To move to a more realistic scenario, we simulated a network of ten excitatory neurons (with all-to-all connectivity) and three inhibitory neurons. Each inhibitory neuron received input from eight randomly selected excitatory neurons and randomly projected

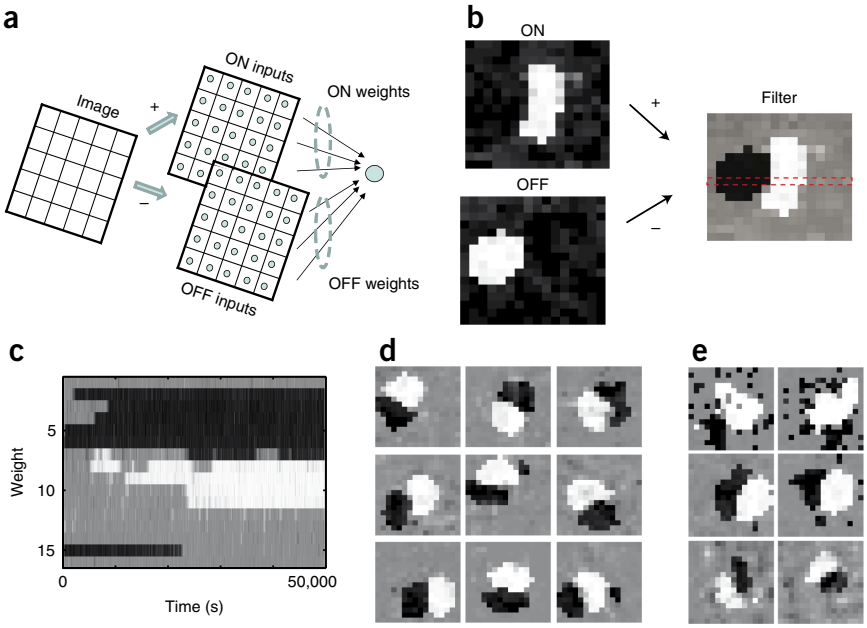

Figure 7 Receptive fields development. (**a**) A small patch of 16 × 16 pixels was chosen from the whitened natural images benchmark[35]. The patch was selected randomly and was presented as input to 512 neurons for 200 ms. The positive part of the image was used as the firing rate to generate Poisson 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 16 × 16 image. The filter was calculated by subtracting the OFF weights from the ON weights. The filter was localized and bimodal, corresponding to an oriented receptive field. (**c**) Temporal evolution of the weights shown in the red dashed box in **b**. (**d**) Nine different neurons. (**e**) Two different neurons receiving presynaptic input with varying firing rates from 0–25 Hz (top), 0–37.5 Hz (middle) and 0–75 Hz (bottom).

back to six excitatory neurons. In addition to the recurrent input, each excitatory and inhibitory neuron received feedforward spike input from 500 presynaptic neurons j that generated stochastic Poisson input at a rate ν_j. The rates of neighboring input neurons were correlated, mimicking the presence or absence of spatially extended objects. The location of the stimulus was switched every 100 ms to a new random position. In case of retinal input, this would correspond to a situation where the subject fixates every 100 ms on a new stationary stimulus. Depending on the retinal position of stimulus, a given postsynaptic neuron responds with a low, medium or high firing rate, which is stationary during the 100-ms stimulation period; the firing rates of the ten neurons in the network encode the current position of the stimulus (rate-coding procedure). In a temporal-coding procedure, the model input is shifted every 20 ms to a neighboring location, mimicking rapid movement of an object across an array of sensory receptors (for example, during whisking behavior)[33]. In this scenario, a given model neuron exhibits only short, transient bursts of a few spikes; thus, it is the temporal structure of the activity (as opposed to stationary firing rates) that encodes the position and movement of the stimulus. 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.

During the first 100–400 s of stimulation in the rate-coding procedure, the excitatory neurons developed localized receptive fields; that is, weights from neighboring inputs to the same postsynaptic neuron became either strong or weak together and stayed stable thereafter (**Fig. 5a**). Similarly, lateral connections onto the same postsynaptic neuron developed to strong or weak synapses, which remained, apart from fluctuations, stable thereafter (**Fig. 5a**), leading to a structured pattern of synaptic connections (**Fig. 5b**). After reordering from the neurons according to similarity of receptive fields, we found that three

Table 1 Parameters

a

Parameters	Value
C, membrane capacitance	281 pF
g_L, leak conductance	30 nS
E_L, resting potential	–70.6 mV
Δ_T, slope factor	2 mV
$V_{T_{rest}}$, threshold potential at rest	–50.4 mV
$\tau_{w_{ad}}$, adaptation time constant	144 ms
a, subthreshold adaptation	4 nS
b, spike triggered adaptation	0.805 pA
Isp, spike current after a spike	400 pA
τ_z, spike current time constant	40 ms
τ_{V_T}, threshold potential time constant	50 ms
$VTmax$, threshold potential after a spike	30.4 mV

b

Experiments	θ_- (mV)	θ_+ (mV)	A_{LTD} (mV^{-1})	A_{LTD} (mV^{-2})	τ_x (ms)	τ_- (ms)	τ_+ (ms)
Visual cortex*,[16]	–70.6	–45.3	14×10^{-5} **	8×10^{-5} **	15**	10**	7**
Somatosensory cortex[30]	–70.6	–45.3	21×10^{-5} **	30×10^{-5} **	30**	6**	5**
Hippocampal[10]	–41**	–38**	38×10^{-5} **	2×10^{-5} **	16		

(**a**) Parameters for the neuron model. (**b**) Plasticity rule parameters for the various experiments. * indicates the standard set of parameters. ** indicates the free parameters fitted to experimental data. Other parameters were set in advance on the basis of previous studies.

groups of neurons had formed, which were characterized by strong bidirectional connectivity in the group, and different receptive fields and no lateral connectivity between groups (**Fig. 5c**). If the overall amplitude of plastic changes was small (compared with that found in the experiments), the pattern of lateral connectivity was stable and had only a few strong bidirectional connections amidst a great deal of weak lateral connectivity. The reason for this is that two neurons with similar receptive fields are both active at high rate whenever the stimulus is in the center region of their receptive field, which gives rise to strong bidirectional lateral connections (**Fig. 4**). Unidirectional strong connections were nearly absent (**Fig. 5**). If the amplitude and rate of plasticity is more realistic and consistent with our data (**Fig. 2**), then the pattern of lateral connectivity changed between one snapshot in time and another one 5 s later, but the overall pattern was stable when averaged over 100 s (**Fig. 5f–h**). In each snapshot, about half of the strong connections were bidirectional (**Fig. 5h,i**).

This connectivity pattern contrasts with that shown under a temporal coding procedure (**Fig. 6**). Neurons developed receptive fields similar to those seen with the rate-coding procedure, but, as expected for temporal Hebbian learning[14], the receptive field shifted over time (**Fig. 6a**). With a small learning rate, this shift was slow, as in previous models[14], but with realistic learning parameters extracted from our experiments (**Fig. 2**), the shift of the receptive field was rapid (**Fig. 6a**). Notably, among the lateral connections, strong reciprocal links were nearly absent, whereas strong unidirectional connections from neuron n to neurons $n + 1$, $n + 2$ and $n + 3$ dominated (**Fig. 6b–e**). As the pattern of feedforward connections forming the receptive fields changed, the structure of lateral connections changed as well on the time scale of 10 min. Nevertheless, at each moment in time, the pattern of lateral connections was highly asymmetric, favoring connections from neuron n to $n + k$ (with $k = 1$, 2 and 3) over those from n to $n - k$, where n is the neuronal index after relabeling according to the receptive field position (**Fig. 6a**). This suggests that temporal coding procedures in which stimuli are nonstationary and exhibit systematic spatio-temporal correlations are reflected in the functional connectivity pattern by strong unidirectional connections, whereas rate coding (characterized by stationary input with spatial correlations only) leads to strong bidirectional connections. We confirmed this for a broad range of stimuli and in the presence of noise (**Supplementary Figs. 1–3**).

Development of localized receptive fields

The results for the feedforward connectivity in the previous subsection lead to the question of the behavior of our plasticity model under stimulation procedures previously used for rate models[24,34,35]. Both our spiking rule and the rate-based BCM model[24] require presynaptic activity to induce a change. Furthermore, for our rule, as well as for the simplest BCM rule (see ref. 24), the depression terms are linear and the potentiation terms are quadratic in the postsynaptic variables (that is, the postsynaptic potential or the postsynaptic firing rate). More quantitatively, for Poisson input, the total weight change Δw in our model is proportional to $v^{\mathrm{pre}} v^{\mathrm{post}} (v^{\mathrm{post}} - \vartheta)$, where v^{pre} and v^{post} denote the firing rates of pre- and postsynaptic neurons, respectively, and ϑ is a sliding threshold related to the ratio between the LTP- and LTD-inducing processes (equation (8)). The sliding threshold arises in our plasticity model because the amount of LTD A_{LTD} depends on the long-term average of the voltage on the slow time scale of homeostatic processes. Because of its similarities to BCM, we were not surprised that our spike-based learning rule with sliding threshold was able to support the development of localized receptive fields, a feature related to independent component analysis (ICA) and sparse coding[24,34].

In our experiments, the input consists of small patches of natural images using standard preprocessing[36]. After learning with our plasticity rule, the weights exhibit a stable spatial structure that can be interpreted as a receptive field (**Fig. 7**). In contrast with a principal component analysis of image patches (as, for example, implemented by Hebbian learning in linear neurons[37]), the receptive fields were localized (that is, the region with strong weights did not stretch across the whole image patch). Nine runs of the learning experiments gave receptive fields with different locations and orientations (**Fig. 7d**). Because of the homeostatic control of LTD in our plasticity model, the neuron compensated in experiments with increased input firing rates by developing smaller receptive fields that were even more localized (**Fig. 7e**). Development of localized receptive fields has been interpreted as a signature of ICA or sparse coding[35]. In contrast with most other ICA algorithms[36], our rule is biologically more plausible, as it is consistent with data from a large body of plasticity experiments.

DISCUSSION

Because traditional plasticity rules are rate models, the relation between coding and connectivity cannot be studied. Our plasticity rule is formulated on the level of postsynaptic voltage. Because action potentials are sharp voltage peaks, they act as singular events in the voltage so that, in the presence of a spike, our rule turns automatically into a spike timing–dependent rule. Indeed, for spike coding (and without subthreshold voltage manipulations), our plasticity rule behaves similar to a STDP rule in which triplets of spikes with pre-post-post or post-pre-post timing evoke LTP[26,27], whereas pairs with post-pre timing evoke LTD. In contrast with standard STDP rules (reviewed in ref. 14), pairing-frequency dependence[16] and burst-timing dependence[30] are qualitatively described. In addition, the rule is expected to reproduce the triplet and quadruplet experiments in hippocampal slices[38] (data not shown), as for all STDP protocols the plasticity rule that we used is similar to an earlier nonlinear STDP rule[27]. Deriving STDP rules from voltage dependence has been attempted before[16,39,40]. However, because these earlier models use the momentary voltage[40] or its derivative[39], rather than the combination of momentary and averaged voltage that we used 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. The voltage-based model[16] uses separate empirical functions for timing dependence, voltage dependence, frequency dependence and multiple spike summation with preference for LTP to capture the nonlinear effects of LTP. Our model is similar in that it also uses momentary voltage before the spike as one of the variables, but it requires neither an explicit frequency-dependent term nor an explicit timing-dependent term. Instead, frequency and timing dependence follow from the model dynamics. Our model has similarities with LTP induction in the TagTriC model[41], but the TagTriC model focuses on the long-term stability of synapses, rather than spike-timing dependence of the induction mechanism.

Even though our model does not require a biophysical interpretation of the variables, it is tempting to speculate about potential mechanisms. For the depression term in our model, a trace $\bar{u}\_$ left by previous activity of the postsynaptic neuron is combined with spike arrival x at the presynaptic terminal (**Fig. 1a**). In light of the results on LTD in layer V neocortical neurons[42], this trace could be related to endocannabinoids released from the postsynaptic site. Coincidence of this slow trace with the activation of presynaptic NMDA receptors (which rapidly respond to the glutamate released by presynaptic activity $x(t)$) could be the trigger signal for LTD[42]. Indeed, the duration of the LTD component in the STDP function increases if the endocannabinoid trace is artificially prolonged (see **Fig. 9** of

ref. 42). In other neuron types and brain areas, the same mathematical model (but with different parameters) could correspond to different biophysical mechanisms of LTD. For example, in hippocampal CA1 neurons, the trace $\bar{u}_-$ could reflect calcium entry through voltage-gated ion channels during depolarization, which, when combined with synaptic signals (caused by the presynaptic spike arrival x), would give rise to the calcium signals that are necessary to trigger LTD (reviewed in refs. 17,18,42). Potentiation is induced in our model by the combination of three factors: a momentary depolarization above spike threshold, a depolarization just before the spike, $\bar{u}_+$, above rest, and the presence of a trace $\bar{x}$ left by presynaptic spike arrival (**Fig. 1a**). The trace $\bar{x}$ could correspond to the amount of glutamate bound to the postsynaptic NMDA receptor, but this is controversial[42]. A high momentary voltage u can be induced by a backpropagating action potential; notably, backpropagation of action potentials is more likely to occur and will more reliably occur in the background of a weak depolarization of the dendrite[42], and such a weak depolarization potentially corresponds to the term $\bar{u}_+$ in our model. Because we have a depolarizing afterpotential after each spike in our model (**Fig. 1c,d**), the value of $\bar{u}$ just before the next spike increases with the repetition frequency of the STDP protocol, consistent with previous experiments (**Fig. 5d** in ref. 42). Our model is therefore consistent with previous results showing that LTP can be induced in distal synapses only if additional cooperative input or dendritic depolarization prevents failure of backpropagating action potentials[43]. In the context of the classical view of the NMDA receptor as a coincidence detector[42], it is quite natural to see why a sequence post-pre-post of two postsynaptic action potentials and one presynaptic spike are ideal for LTP. The spike afterpotential of the first postsynaptic action potential removes the calcium block and prepares the dendrite for successful backpropagation of a later action potential. If the backpropagating action potential caused by the second postsynaptic spike occurs just slightly after presynaptic spike arrival, this causes a sharply peaked and large calcium transient that would be sufficient to trigger the LTP induction chain.

Even though our model is formulated on the level of voltage, we do not imply that voltage itself is the essential biophysical mechanism. Rather, under physiological conditions, the voltage transient (or current or conductance transient) caused by synaptic input or action potential firing is the starting point of long biochemical signaling chains that lead to induction of plasticity. In our phenomenological model, the signature of the inputs (here, voltage transients) are directly linked (via mathematical variables or traces) to the induction of plasticity, jumping over the biophysical mechanisms of the signal transduction chain.

Our plasticity rule allows us to explain experiments from two different studies with a single principle. Both the 'potentiation is rescued by depolarization'[16] scenario (**Fig. 2f**) and that of burst-timing dependent LTP[30] (**Fig. 3**) indicate that LTP is induced at low frequency when the membrane is depolarized before the pre-post pairing. This depolarization can be the result of a previous spike during a postsynaptic burst[30] or to a depolarization current. A further unexpected result is that, with the set of parameters derived from visual cortex slice experiments, synapses fluctuated rapidly between strong and weak weights. This aspect is interesting in light of the synapse mobility that has been reported in imaging experiments[5].

Possible extensions of the model include a weight dependence of synaptic plasticity. We assumed that weights can grow to a hard upper bound, but the rule can easily be changed to soft bounds[14] by changing the prefactors A_{LTP} and A_{LTD} accordingly[41]. Second, short-term plasticity[44] could be added for a better description of the plasticity phenomena that occurs during high-frequency protocols. Third, additional mechanisms need to be implemented to describe the transition from early to late LTP/LTD[41,45]. Finally, we can generalize from point neurons to spatially extended neurons using a multicompartment neuron model (for example, distinct compartments for the soma and dendrites). We did not do this here because detailed spatial models introduce a considerable number of new parameters, making overfitting more likely to occur. Notably, our voltage-based formulation of plasticity, if applied locally in a compartmental model, would allow potentiation to occur in a dendritic branch whenever the three conditions—presynaptic activity, recent postsynaptic depolarization and momentary large depolarization—occur together, independent of the source of depolarization. Thus, dendritic spikes could lead to potentiation in the absence of somatic action potentials, consistent with data from experiments in hippocampal[46–48] and cortical slices[31].

Our plasticity model leads to several predictions that could be tested in slice experiments. First, the model predicts that in voltage-clamp experiments the weight change is dependent on the voltage and the number of presynaptic spikes but not on their exact timing (for example, low frequency, tetanus or burst). Second, in the scenario in which 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.

The influence of STDP on temporal coding has been previously studied with respect to changes in the feedforward connections (reviewed in ref. 14). The effect of STDP on lateral connectivity has been much less studied[20–23]. We found that, because of STDP, coding influences the network topology; that is, different stimulation procedures generate different patterns of lateral connectivity. Our results contrast with those of standard STDP rules, which always suppress short loops and, in particular, bidirectional connections[19,32]. Our more realistic plasticity model shows that under a rate-coding procedure (where the neuron is stimulated by different stationary patterns), bidirectional connectivity and highly connected clusters with multiple loops are not only possible but even dominant. It is only for temporal coding (characterized by stimulation with substantial spatiotemporal correlations) that our biologically plausible rule leads to dominant unilateral directions. 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 in which cells in a recurrent network are repeatedly stimulated in a fixed order would decrease the fraction of strong bidirectional connections, whereas a stimulation pattern in which clusters of neurons fire at a high rate during episodes of a few hundred milliseconds would increase this fraction. In this view, it is tempting to connect the low degree of bidirectional connectivity in barrel cortex[4] to the bigger importance of temporal structure in whisker input[33], compared with visual input[3].

METHODS

Methods and any associated references are available in the online version of the paper at http://www.nature.com/natureneuroscience/.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

This work was supported by the European project FACETS and the Swiss National Science Foundation.

AUTHOR CONTRIBUTIONS

C.C. developed the model and carried out the experiments. L.B. and E.V. participated in discussions. W.G. supervised the project and wrote most of the manuscript.

COMPETING INTERESTS STATEMENT
The authors declare no competing financial interests.

Published online at http://www.nature.com/natureneuroscience/.
Reprints and permissions information is available online at http://www.nature.com/reprintsandpermissions/.

1. Buonomano, D.V. & Merzenich, M.M. Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* **21**, 149–186 (1998).
2. Dan, Y. & Poo, M. Spike timing–dependent plasticity of neural circuits. *Neuron* **44**, 23–30 (2004).
3. Song, S., Sjöström, P.J., Reigl, M., Nelson, S. & Chklovskii, D.B. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* **3**, e350 (2005).
4. Lefort, S., Tomm, C., Sarria, J.C.F. & Petersen, C.C.H. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron* **61**, 301–316 (2009).
5. Yuste, R. & Bonhoeffer, T. Genesis of dendritic spines: insights from ultrastructural and imaging studies. *Nat. Rev. Neurosci.* **5**, 24–34 (2004).
6. Hebb, D.O. *The Organization of Behavior* (Wiley, New York, 1949).
7. Malenka, R.C. & Bear, M.F. LTP and LTD: an embarassment of riches. *Neuron* **44**, 5–21 (2004).
8. Markram, H., Lübke, J., Frotscher, M. & Sakmann, B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**, 213–215 (1997).
9. 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).
10. Ngezahayo, A., Schachner, M. & Artola, A. Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *J. Neurosci.* **20**, 2451–2458 (2000).
11. Dudek, S.M. & Bear, M.F. Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus. *J. Neurosci.* **13**, 2910–2918 (1993).
12. Gerstner, W., Kempter, R., Van Hemmen, L. & Wagner, H. A neuronal learning rule for sub-millisecond temporal coding. *Nature* **383**, 76–81 (1996).
13. Legenstein, R., Naeger, C. & Maass, W. What can a neuron learn with spike timing–dependent plasticity? *Neural Comput.* **17**, 2337–2382 (2005).
14. Gerstner, W. & Kistler, W.M. *Spiking Neuron Models* (Cambridge University Press, New York, 2002).
15. Lisman, J. & Spruston, N. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nat. Neurosci.* **8**, 839–841 (2005).
16. Sjöström, P.J., Turrigiano, G.G. & Nelson, S.B. Rate, timing and cooperativity jointly determine cortical synaptic plasticity. *Neuron* **32**, 1149–1164 (2001).
17. Shouval, H.Z., Bear, M.F. & Cooper, L.N. A unified model of NMDA receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA* **99**, 10831–10836 (2002).
18. Lisman, J.E. & Zhabotinsky, A.M. A model of synaptic memory: a CaMKII/PP1 switch that potentiates transmission by organizing an AMPA receptor anchoring assembly. *Neuron* **31**, 191–201 (2001).
19. Song, S. & Abbott, L.F. Cortical development and remapping through spike timing–dependent plasticity. *Neuron* **32**, 339–350 (2001).
20. Lubenov, E.V. & Siapas, A.G. Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron* **58**, 118–131 (2008).
21. Levy, N., Horn, D., Meilijson, I. & Ruppin, E. Distributed synchrony in a cell assembly of spiking neurons. *Neural Netw.* **14**, 815–824 (2001).
22. Morrison, A., Aertsen, A. & Diesmann, M. Spike timing–dependent plasticity in balanced random networks. *Neural Comput.* **19**, 1437–1467 (2007).
23. Izhikevich, E.M. & Edelman, G.M. Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. USA* **105**, 3593–3598 (2008).
24. Cooper, L.N., Intrator, N., Blais, B.S. & Shouval, H.Z. *Theory of Cortical Plasticity* (World Scientific, Singapore, 2004).
25. Miller, K.D. 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).
26. Senn, W., Tsodyks, M. & Markram, H. An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput.* **13**, 35–67 (2001).
27. Pfister, J.-P. & Gerstner, W. Triplets of spikes in a model of spike timing–dependent plasticity. *J. Neurosci.* **26**, 9673–9682 (2006).
28. O'Connor, D.H., Wittenberg, G.M. & Wang, S.S.H. Dissection of bidirectional synaptic plasticity into saturable unidirectional processes. *J. Neurophysiol.* **94**, 1565–1573 (2005).
29. Turrigiano, G.G. & Nelson, S.B. Homeostatic plasticity in the developing nervous system. *Nat. Rev. Neurosci.* **5**, 97–107 (2004).
30. Nevian, T. & Sakmann, B. Spine Ca^{2+} signaling in spike timing–dependent plasticity. *J. Neurosci.* **26**, 11001–11013 (2006).
31. Kampa, B.M., Letzkus, J.J. & Stuart, G.J. Requirement of dendritic calcium spikes for induction of spike timing–dependent synaptic plasticity. *J. Physiol. (Lond.)* **574**, 283–290 (2006).
32. Kozloski, J. & Cecchi, G.A. Topological effects of synaptic spike timing–dependent plasticity. Preprint at <http://arxiv.org/abs/0810.0029> (2008).
33. Jadhav, S.P., Wolfe, J. & Feldman, D.E. Sparse temporal coding of elementary tactile features during active whisker sensation. *Nat. Neurosci.* (2009).
34. Blais, B.S., Intrator, N., Shouval, H. & Cooper, L. Receptive field formation in natural scene environments. Comparison of single-cell learning rules. *Neural Comput.* **10**, 1797–1813 (1998).
35. Olshausen, B.A. & Field, D.J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (1996).
36. Hyvärinen, A., Karhunen, J. & Oja, E. *Independent Component Analysis* (Wiley, New York, 2001).
37. Oja, E. A simplified neuron as a principal component analyzer. *J. Math. Biol.* **15**, 267–273 (1982).
38. Wang, H.X., Gerkin, R.C., Nauen, D.W. & Bi, G.-Q. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.* **8**, 187–193 (2005).
39. Saudargiene, A., Porr, B. & Wörgötter, F. How the shape of pre- and postsynaptic signals can influence STDP: a biophysical model. *Neural Comput.* **16**, 595–626 (2004).
40. Brader, J.M., Senn, W. & Fusi, S. Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Comput.* **19**, 2881–2912 (2007).
41. Clopath, C., Ziegler, L., Vasilaki, E., Büsing, L. & Gerstner, W. Tag-trigger consolidation: a model of early and late long-term potentiation and depression. *PLOS Comput. Biol.* **4**, e1000248 (2008).
42. Sjöström, P.J., Turrigiano, G.G. & Nelson, S.B. Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. *Neuron* **39**, 641–654 (2003).
43. Sjöström, P.J. & Häusser, M. A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. *Neuron* **51**, 227–238 (2006).
44. Tsodyks, M.V. & Markram, H. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA* **94**, 719–723 (1997).
45. Frey, U. & Morris, R.G.M. Synaptic tagging and long-term potentiation. *Nature* **385**, 533–536 (1997).
46. Remy, S. & Spruston, N. Dendritic spikes induce single-burst long-term potentiation. *PNAS* **104**, 17192–17197 (2007).
47. Hardie, J. & Spruston, N. Synaptic depolarization is more effective than back-propagating action potentials during induction of associative long-term potentiation in hippocampal pyramidal neurons. *J. Neurosci.* **29**, 3233–3241 (2009).
48. Golding, N.L., Staff, N.P. & Spruston, N. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* **418**, 326–331 (2002).

ONLINE METHODS

Neuron model. In contrast with standard models of STDP, our plasticity model uses the postsynaptic membrane potential $u(t)$. As a model for neuronal voltage, we chose the adaptive exponential integrate-and-fire (AdEx) model[49] with an additional current describing the depolarizing spike after potential[50]. The voltage evolution is

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

where C is the membrane capacitance, g_L is the leak conductance, E_L is the resting potential and I is the stimulating current. The exponential term describes the activation of sodium current. The parameter Δ_T is the slope factor and V_T is the threshold potential. 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 and a is a parameter. On firing, the variable u is reset to the fixed value V_{reset}, whereas w_{ad} is increased by the amount b. The main difference between this and a previously described model[23] is that the voltage is exponential rather than quadratic, allowing for a better fit to data[50]. The spike afterpotential of the cells used in typical STDP experiments[16] have a long depolarizing spike afterpotential. We therefore added 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 was modeled with the adaptive threshold V_T, which starts at $V_{T_{\max}}$ after a spike and decays to $V_{T_{\text{rest}}}$ with a time constant τ_{V_T}[50], that is,

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

Parameters for the neuron model are taken from ref. 49 for the AdEx model, τ_z was set to 40 ms, consistent with ref. 16 (see also ref. 50), and kept fixed throughout all simulations (see **Table 1a**).

Plasticity model. Our model exhibits separate additive contributions to the plasticity rule, one for LTD and another one for LTP[28]. For the LTD part, we assumed that presynaptic spike arrival at synapse i induces depression of the synaptic weight w_i by $-A_{\text{LTD}}(\bar{u}_-(t)-\theta_-)_+$. $(\,)_+$ indicate rectification, that is, any value $\bar{u}_- < \theta_-$ does not lead to a change[9] (see **Fig. 1h**). The quantity $\bar{u}_-(t)$ is an exponential low-pass-filtered version of the postsynaptic membrane potential $u(t)$ with time constant τ_-

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

The variable $\bar{u}_-$ is an abstract variable that could, for example, reflect the level of calcium concentration[17] or the release of endocannabinoids[42], although such an interpretation is not necessary for our rule. Because 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 represented by

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

where $A_{\text{LTD}}(\bar{\bar{u}})$ is an amplitude parameter that is under the control of homeostatic process[29]. For slice experiments, the parameter has a fixed value that was determined experimentally. For the network simulations shown in **Figures 5–7**, the parameter depends on the mean depolarization $\bar{u}_-$ of the postsynaptic neuron, averaged over a time scale of 1 s. Equation (1) is a simple method for implementing homeostasis; other methods, such as weight rescaling, would also be possible[29]. The time scale of 1 s is not critical (100 s or more would be more realistic for homeostasis) but is convenient for the numerical implementation.

For the LTP component, we assumed 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[12,27]

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

where $X_i(t)$ is the spike train defined above. The quantity $\bar{x}_i(t)$ could, for example, represent the amount of glutamate bound to postsynaptic receptors[27] or the number of NMDA receptors in an activated state[26]. Potentiation is given by

$$\frac{d}{dt}w_i^+ = A_{\text{LTP}}\bar{x}_i(u-\theta_+)_+(\bar{u}_+ - \theta_-)_+ \quad \text{if } w_i < w_{\max} \qquad (2)$$

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

The final rule used in the simulation was

$$\frac{d}{dt}w_i = -A_{\text{LTD}}(\bar{\bar{u}})X_i(\bar{u}_- - \theta_-)_+ + A_{\text{LTP}}\bar{x}_i(u-\theta_+)_+(\bar{u}_+ - \theta_-)_+ \qquad (3)$$

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

It is unlikely that the model can be simplified further. First, voltage is necessary as a variable whenever voltage is manipulated in experiments. Second, dependence on voltage must be nonlinear[9–11]. Phenomenological models have some freedom in the choice of the mathematical form of the nonlinearities (for example, exponential, polynomial Hill functions or piecewise linear) and we chose a suitable combination of piecewise linear functions with thresholds θ_+ and θ_-. Third, STDP experiments indicate that the temporal relation between stimulation events is important. All timing relations have been implemented as (first order) linear filtering. For the case of classical STDP experiments, where all spikes are triggered by the experimenter, our phenomenological model can be simplified and becomes identical or closely related to existing nonlinear STDP models[26,27], but regarding the interaction between voltage and spike timing, such a further simplification is not possible. Finally, the fact that the curve of burst timing–dependent plasticity (**Fig. 3c**) is not perfectly reproduced indicates that our plasticity model does not have an unnecessarily large number of free parameters.

Analysis of plasticity model. We established a quantitative link between our plasticity model (equation (3)) and BCM theory[24] under the assumption of a linear Poisson neuron model with spikes. In a linear Poisson neuron, input spike trains $X_j(t) = \sum_i \delta(t-t_i^f)$ are low-pass filtered and weighted to give a subthreshold potential $u^s(t) = \sum_j \int_0^\infty \varepsilon(s)X_j(t-s)ds$, where $\varepsilon(s)$ is the time course of an excitatory postsynaptic potential and u^s is measured with respect to the resting potential θ_-. The linear Poisson neuron generates spikes stochastically with stochastic firing intensity ν^{post} proportional (with parameter $1/\alpha$) to u^s, hence the probability of firing in a short time between t and $t+\Delta$ is $P_F(t;t+\Delta) = \nu^{\text{post}}(t)\Delta = u^s(t)\frac{\Delta}{\alpha}$. If the linear Poisson neuron spikes at time t_f^{post}, we add a short voltage pulse $\beta\delta(t-t_f^{\text{post}})$. The total membrane potential is therefore

$$u(t) = u^s(t) + \beta Y(t) + \theta_- \qquad (4)$$

where $Y(t) = \sum_f \delta(t-t_f^{\text{post}})$ is the spike train of the postsynaptic neuron and β is the integral weight of spikes. To illustrate the importance of β, suppose that in a hypothetical experiment of 100-ms duration we found a single triangular action potential with amplitude 120 mV and a 1-ms duration at half-maximum, and that the voltage was otherwise constant at a value of 2 mV above rest. The mean voltage averaged over this 100-ms period would therefore be

doi:10.1038/nn.2479

$\int_0^{100} \frac{u(t)dt}{100} = 2$ mV + 1.2 mV + θ_- so that the weight parameter β in equation (4) should have a value of 1.2 mV.

By construction, the expected number of spikes of the linear Poisson neuron is equal to its instantaneous rate $< Y > (t) = \nu^{\text{post}}(t) = \frac{u^s(t)}{\alpha}$. In the following derivation, the time dependence of the variables is not explicitly denoted for the sake of simplicity (except for a few special cases); for example, $u(t)$ is abbreviated as u.

We assumed that the neuron has N excitatory synapses stimulated by N presynaptic Poisson spike trains of rates $\nu^{\text{pre}} = (\nu^{\text{pre}}, ..., \nu_{N1}^{\text{pre}})$. Furthermore, we assumed that the presynaptic rates ν^{pre} are slowly varying quantities compared with the intrinsic time scales τ_+ and τ_- of our plasticity model or those of our neuron model (for example, excitatory postsynaptic potential duration), which were all below 50 ms. This assumption explicitly resulted in the following simplifications: $\bar{\nu}^{\text{pre}} \approx \nu^{\text{pre}}$, $\bar{\nu}_-^{\text{post}} \approx \nu^{\text{post}}$ and $\bar{\nu}_+^{\text{post}} \approx \nu^{\text{post}}$. For a variable q, $\bar{q}$ denotes low-pass filtering with the time constant τ_q, and $\bar{q}_+$ and $\bar{q}_-$ correspond to the time constants τ_+ and τ_-, respectively.

Using the linear Poisson model defined above in the plasticity rule (equation (3)) yields (if we suppress for the moment the dependence on the homeostatic variable $\bar{\bar{u}}$)

$$\frac{d}{dt} w_i = -A_{\text{LTD}} X_i(t)(\bar{u}_-^s + \beta \bar{Y}_-) + A_{\text{LTP}} \bar{x}_i(t) \beta Y(\bar{u}_+^s + \beta \bar{Y}_+) \quad (5)$$

with θ_- being equal to the resting potential, all voltages being above resting potential, as only excitatory inputs are considered, and only Y being above the firing threshold θ_+, as u^s was the subthreshold voltage. Taking the average $< . >_{\text{post}}$ over the postsynaptic spikes given the postsynaptic rate ν^{post} yields

$$< \frac{d}{dt} w_i >_{\text{post}} = -(\alpha + \beta) A_{\text{LTD}} X_i(t) \bar{\nu}_-^{\text{post}} + (\alpha + \beta) A_{\text{LTP}} \bar{x}_i(t) \beta \nu^{\text{post}} \bar{\nu}_+^{\text{post}} \quad (6)$$

Here, we used $< Y(t) \bar{Y}_+(t) >_{\text{post}} = \nu^{\text{post}}(t) \bar{\nu}_+^{\text{post}}(t)$, which holds because $\bar{Y}_+(t)$ is not influenced by a possible spike at time t (just by spikes at times s with $s < t$), and it is therefore uncorrelated with $Y(t)$ given ν^{post}. For slowly varying input rates

$$< \frac{d}{dt} w_i >_{\text{post}} = -(\alpha + \beta) A_{\text{LTD}} X_i(t) \nu^{\text{post}} + (\alpha + \beta) A_{\text{LTP}} \bar{x}_i(t) \beta \nu^{\text{post}} \nu^{\text{post}} \quad (7)$$

Taking the average $< . >_{\text{post}}$ over the presynaptic spikes given the presynaptic firing rates ν^{pre} and neglecting spike-spike correlations (that is, correlations between X_i and ν^{post} beyond rate correlations between ν^{pre} and ν^{post}) gives

$$\begin{aligned} < \frac{d}{dt} w_i >_{\text{post}} &= -(\alpha + \beta) A_{\text{LTD}} \nu_i^{\text{pre}} \nu^{\text{post}} + (\alpha + \beta) A_{\text{LTP}} \nu_i^{\text{pre}} \beta \nu^{\text{post}} \nu^{\text{post}} \\ &= (\alpha + \beta) \beta A_{\text{LTP}} \nu_i^{\text{pre}} \nu^{\text{post}} (\nu^{\text{post}} - \frac{A_{\text{LTD}}}{\beta A_{\text{LTP}}}) \end{aligned} \quad (8)$$

Here, $<< . >_{\text{post}} >_{\text{pre}}$ was abbreviated as $< . >$. The factor $(\alpha + \beta)\beta$ can be interpreted as the learning rate in a rate-based plasticity model and $\frac{A_{\text{LTD}}}{\beta A_{\text{LTP}}} = \vartheta$ as the threshold for the transition of LTD to LTP in the quadratic BCM model[24]. Because A_{LTD} depends on the slow time scale of homeostatic processes on the long-term averaged potential $\bar{\bar{u}}$, the threshold ϑ is a sliding one. Just as in the BCM model[24], our plasticity model responds to persistent periods of high activity with an increase in the threshold ϑ.

Parameters and data fitting. For the plasticity slice experiments, we took $\bar{\bar{u}} = u_{\text{ref}}$ as fixed and fit the parameter A_{LTD}. The total number of parameters of the plasticity model is therefore seven. For all data sets, except the one taken from ref. 10, the threshold θ_- was set to the resting potential and θ_+ to the firing threshold of the AdEx model, that is, $\theta_- = -70.6$ mV and $\theta_+ = -45.3$ mV. The remaining five parameters, τ_x, τ_-, τ_+, A_{LTD} and A_{LTP}, were fitted to each data set individually by the following procedure. We calculated the theoretically predicted weight change $\Delta w_i^{\text{th},j}$ by integrating (analytically or numerically) equation (3), for a given experimental protocol j, as a function of the free parameters. We then estimated the free parameters by minimizing the mean-square error E between the theoretical calculations and the experimental data $\Delta w_i^{\text{exp},j}$

$$E = \sum_j (\Delta w_i^{\text{th},j} - \Delta w_i^{\text{exp},j})^2$$

For the data set in hippocampus[10], we also fit the two parameters θ_- and θ_+, as completely different preparations and cell type were used. Moreover, for this data set, the time constant τ_x was taken from physiological measurements given in ref. 2 and fixed to the value of 16 ms. The parameters for the various experiments are summarized in **Table 1b**.

Voltage-clamp experiment. The postsynaptic membrane potential was switched in the simulations to a constant value, u_{clamp}, chosen from –80 to 0 mV while synapses were stimulated with either 25 (blue line) or 100 pulses (red line) at 50 Hz. As a result of voltage clamping, the actual value of the voltage u itself and the low-pass-filtered versions $\bar{\bar{u}}$ are constant and equal to u_{clamp}. Thus, the synaptic plasticity rule becomes

$$\frac{d}{dt} w_i = -A_{\text{LTD}} X_i(t)(u_{\text{clamp}} - \theta_-)_+ + A_{\text{LTP}} \bar{x}_i(t)(u_{\text{clamp}} - \theta_+)_+ (u_{\text{clamp}} - \theta_-)_+.$$

STDP experiment and frequency dependence. Presynaptic spikes in the simulation were paired with postsynaptic spikes that were either advanced by +10 ms or delayed by –10 ms with respect to the presynaptic spike. Postsynaptic spikes were triggered by brief, strong current pulses into the postsynaptic neuron. The pairing was repeated five times with different frequencies ranging from 0.1 to 50 Hz. These five pairings were repeated 15 times at 0.1 Hz. However, the five pairings at 0.1 Hz were repeated only ten times to mimic the experimental protocol[16].

Burst timing–dependent plasticity For **Figure 3a**, the presynaptic spike was paired $\Delta t = +10$ ms before (or $\Delta t = -10$ ms after) 1, 2 or 3 postsynaptic spikes. The frequency of the burst was 50 Hz. The neuron received 60 pairings at a frequency of 0.1 Hz. For **Figure 3b**, the presynaptic spike was paired with a burst of three action potentials ($\Delta t = +10$ ms and –10 ms), whereas the burst frequency varies from 20 to 100 Hz. For **Figure 3c**, a presynaptic spike is paired with a burst of three postsynaptic action potentials with burst frequency of 50 Hz. 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 ref. 30.

Poisson input for functional scenarios. Poisson inputs were used in all of the following experiments. They were generated by a stochastic process where the spike was elicited with a stochastic intensity ν.

Relation between connectivity and coding: toy model. Weights of ten all-to-all connected neurons were initialized at 1, bounded between 0 and 3. Weights evolved with the voltage-based rule (equation (3)) for 100 s. The model was compared with a canonical pair-based STDP model written as $\frac{d}{dt} w_i = -A_{\text{LTD}}^{\text{pair}} X_i \bar{y} + A_{\text{LTP}}^{\text{pair}} \bar{x}_i Y$, where Y is the postsynaptic spike train defined in the same manner as the presynaptic spike train X_i with a filter of the postsynaptic spikes $\bar{y}$ similar to $\bar{x}_i$.

We chose the parameters $A_{\text{LTD}}^{\text{pair}} = A_{\text{LTP}}^{\text{pair}} = 1 \times 10^{-5}$ 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 $\bar{y}$. Neuron 1 fired at 2 Hz, neuron 2 at 4 Hz. neuron 10 at 20 Hz following Poisson statistics; that is, short current pulses were injected to make the neuron fire with Poisson statistics at this frequency. Neurons fired successively every 20 ms, with neuron 1 firing, followed 20 ms later by neuron 2,… followed by neuron 10 and then back to neuron 1, etc. in a loop.

Rate coding in network simulation. 500 presynaptic Poisson neurons with firing rates $\nu_i^{\text{pre}} (1 \le i \le 500)$ are connected to 10 postsynaptic excitatory neurons. The input rates ν_i^{pre} follow a Gaussian profile, that is, $\nu_i^{\text{pre}} = Ae^{\frac{-(i-\mu)^2}{2\sigma^2}}$, with variance $\sigma = 10$ and amplitude $A = 30$ Hz. The center μ of the Gaussian shifts randomly every 100 ms between ten different, equally distributed positions.

Circular boundary conditions are assumed, that is, neuron $i = 500$ is considered to be a neighbor of $i = 1$. Synaptic weights of the feedforward connections to the excitatory neurons are initialized randomly (uniformly in [0.5, 2]) and hard bound are set to 0 and 3. The ten 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, three inhibitory neurons are driven by eight excitatory neurons and the feedforward inputs; they project onto six excitatory neurons, connectivity chosen randomly. Those random recurrent connections are fixed and have a weight equal to 1. The feedforward connections onto the inhibitory neurons are also fixed and chosen randomly between 0 and 0.5. The reference value is set to $u^2_{\mathrm{ref}} = 60\ \mathrm{mV}^2$ and the simulation time to 1,000 s. Parameters were normally chosen as in **Table 1b** (visual cortex data), except for **Fig. 5b–e**, where A_{LTP} and A_{LTD} were reduced by a factor 100.

Temporal coding in network simulation. Settings were determined as described above, but patterns were presented for 20 ms successively (from center position 50 to 100 to 150, etc. in a circular manner). The reference value was set to $u^2_{\mathrm{ref}} = 80\ \mathrm{mV}^2$. We used an asymmetry index calculated by relabeling the neurons according to the current position of their receptive field so that with the cyclic stimulation they get activated one after the other: $n \to n+1 \ldots \to n+k \to n-1 \to n$. We then compared the connection from n to $n+k$ with that from n to $n-k$ and computed $AS = \sum_k w_{n,n+k} - w_{n,n-k}$, $k = 1$–3. Connectivity patterns were also analyzed in model networks where neurons received (in addition to feedforward and lateral input) unspecific stochastic background activity that made them fire spontaneously (**Supplementary Fig. 1**).

ICA-like computation: orientation selectivity with natural images. Ten natural images have been taken from a previously determined benchmark[35]. A small patch of 16×16 pixels from any of the images is randomly chosen every 200 ms, which is on the order of the fixation time between saccades. Half of the time the image matrix is transposed, flipped around the vertical axis or the horizontal axis to remove any statistical orientation bias. 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 gray value reflecting the ensemble mean) with maximum frequency of 50 Hz. The $2 \times 16 \times 16$ 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 (equation (3)), where the reference value is set to $u^2_{\mathrm{ref}} = 50\ \mathrm{mV}^2$. Parameters were chosen as in **Table 1b** (visual cortex data), but A_{LTP} and A_{LTD} were reduced by a factor 10. Every 20 s, an extra normalization was applied to equalize the norm of the ON weights to one of the OFF weights[25].

49. Brette, R. & Gerstner, W. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* **94**, 3637–3642 (2005).
50. Badel, L. *et al.* Dynamic I-V curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *J. Neurophysiol.* **99**, 656–666 (2008).

doi:10.1038/nn.2479

Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks

T. P. Vogels, *et al.*

Science **334**, 1569 (2011);
DOI: 10.1126/science.1211095

pronounced than in *drbp* nulls (*5*, *14*). Functionally, *drbp* and *bruchpilot* phenotypes appear similar: Both demonstrate decreased and desynchronized evoked SV release with atypical short-term facilitation. However, the deficits in evoked SV release are much more severe in *drbp* nulls than in *bruchpilot* nulls [i.e., release occurs at 5% versus 30% (*5*) of the respective wild-type level]. DRBP levels were clearly reduced in *bruchpilot* mutants (fig. S7), whereas gross Bruchpilot levels were not altered in *drbp* mutants (Fig. 2B). Given that even a partial loss of DRBP causes marked reduction in SV release (Fig. 3A), deficits in *bruchpilot* mutants might be explained, at least in part, by a concomitant loss of DRBP, and DRBP probably serves functions beyond the structural and Ca^{2+} channel–clustering roles of Bruchpilot.

Taken together, we identified DRBP as a central part of the AZ cytomatrix. How, in detail, DRBP functionally integrates into this protein network is subject to future analyses. Notably, the short-term plasticity phenotype of *drbp* mutants is reminiscent of mammalian *munc13-1* KO and *caps-1* and *caps-2* DKO mutants (*25*, *26*), which implicates functional links between priming factors and DRBP. Consistent with the functional importance of the DRBP protein family suggested by our study, human genetics recently identified two *rbp* loci associated with autism with high confidence (*27*, *28*).

References and Notes

1. Y. Jin, C. C. Garner, *Annu. Rev. Cell Dev. Biol.* **24**, 237 (2008).
2. S. J. Sigrist, D. Schmitz, *Curr. Opin. Neurobiol.* **21**, 144 (2011).
3. S. Schoch, E. D. Gundelfinger, *Cell Tissue Res.* **326**, 379 (2006).
4. L. Siksou, A. Triller, S. Marty, *Curr. Opin. Neurobiol.* **21**, 261 (2011).
5. R. J. Kittel *et al.*, *Science* **312**, 1051 (2006).
6. Y. Wang, S. Sugita, T. C. Sudhof, *J. Biol. Chem.* **275**, 20033 (2000).
7. H. Hibino *et al.*, *Neuron* **34**, 411 (2002).
8. S. A. Spangler, C. C. Hoogenraad, *Biochem. Soc. Trans.* **35**, 1278 (2007).
9. T. Mittelstaedt, S. Schoch, *Gene* **403**, 70 (2007).
10. S. W. Hell, *Science* **316**, 1153 (2007).
11. J. Bückers, D. Wildanger, G. Vicidomini, L. Kastrup, S. W. Hell, *Opt. Express* **19**, 3130 (2011).
12. J. Hou, T. Tamura, Y. Kidokoro, *J. Neurophysiol.* **100**, 2833 (2008).
13. F. Kawasaki, R. Felling, R. W. Ordway, *J. Neurosci.* **20**, 4885 (2000).
14. W. Fouquet *et al.*, *J. Cell Biol.* **186**, 129 (2009).
15. F. Kawasaki, S. C. Collins, R. W. Ordway, *J. Neurosci.* **22**, 5856 (2002).
16. K. J. Venken, Y. He, R. A. Hoskins, H. J. Bellen, *Science* **314**, 1747 (2006).
17. L. Siksou *et al.*, *J. Neurosci.* **27**, 6868 (2007).
18. P. Rostaing, R. M. Weimer, E. M. Jorgensen, A. Triller, J. L. Bessereau, *J. Histochem. Cytochem.* **52**, 1 (2004).
19. S. Hallermann *et al.*, *J. Neurosci.* **30**, 14340 (2010).
20. E. O. Gracheva, E. B. Maryon, M. Berthelot-Grosjean, J. E. Richmond, *Front. Synaptic Neurosci.* **2**, 141 (2010).
21. E. Neher, T. Sakaba, *Neuron* **59**, 861 (2008).
22. F. Kawasaki, B. Zou, X. Xu, R. W. Ordway, *J. Neurosci.* **24**, 282 (2004).
23. P. S. Kaeser *et al.*, *Cell* **144**, 282 (2011).
24. Y. Han, P. S. Kaeser, T. C. Südhof, R. Schneggenburger, *Neuron* **69**, 304 (2011).
25. C. Rosenmund *et al.*, *Neuron* **33**, 411 (2002).
26. W. J. Jockusch *et al.*, *Cell* **131**, 796 (2007).
27. M. Bucan *et al.*, *PLoS Genet.* **5**, e1000536 (2009).
28. D. Pinto *et al.*, *Nature* **466**, 368 (2010).

Acknowledgments: This work was supported by Deutsche Forschungsgemeinschaft (DFG) grants (SFB 665, SFB 958, and EXC 257) to S.J.S and D.S., as well as Bundesministerium für Bildung und Forschung (The German Federal Agency of Education and Research) funding for Deutsche Zentrum für Neurodegenerative Erkrankungen (DZNE) to D.S. Also, M.S. was supported by a Ph.D. fellowship from the Max Delbrück Center for Molecular Medicine and a Boehringer Ingelheim Fonds Ph.D. fellowship. E.K. and S.W. were supported by Ph.D. fellowships from the graduate school GRK 1123 funded by the DFG. M.M. was supported by a fellowship of the Swiss National Science Foundation (PBSKP3-123456/1).

Supporting Online Material

www.sciencemag.org/cgi/content/full/334/6062/1565/DC1
Materials and Methods
Figs. S1 to S7
References (*29*–*39*)

22 August 2011; accepted 2 November 2011
10.1126/science.1212991

Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks

T. P. Vogels,[1]*† H. Sprekeler,[1]* F. Zenke,[1] C. Clopath,[1,2] W. Gerstner[1]

Cortical neurons receive balanced excitatory and inhibitory synaptic currents. Such a balance could be established and maintained in an experience-dependent manner by synaptic plasticity at inhibitory synapses. We show that this mechanism provides an explanation for the sparse firing patterns observed in response to natural stimuli and fits well with a recently observed interaction of excitatory and inhibitory receptive field plasticity. The introduction of inhibitory plasticity in suitable recurrent networks provides a homeostatic mechanism that leads to asynchronous irregular network states. Further, it can accommodate synaptic memories with activity patterns that become indiscernible from the background state but can be reactivated by external stimuli. Our results suggest an essential role of inhibitory plasticity in the formation and maintenance of functional cortical circuitry.

The balance of excitatory and inhibitory membrane currents that a neuron experiences during stimulated and ongoing activity has been the topic of many studies (*1*–*11*). This balance, first defined as equal average amounts of de- and hyperpolarizing membrane currents (from here on referred to as "global balance"), is essential for maintaining stability of cortical networks (*1*, *2*). Balanced networks display asynchronous irregular (AI) dynamics that mimic activity patterns observed in cortical neurons. Such asynchronous network states facilitate rapid responses to small changes in the input (*2*, *3*, *12*), providing an ideal substrate for cortical signal processing (*4*, *13*, *14*).

Moreover, input currents to cortical neurons are not merely globally balanced but also coupled in time (*5*, *6*, *15*) and cotuned for different stimulus features (*7*, *8*). The tight coupling of excitation and inhibition suggests a more precise, detailed balance, in which each excitatory input arrives at the cell together with an inhibitory counterpart (Fig. 1A), permitting sensory inputs to be transiently (*9*) or persistently turned on by targeted disruptions of the balance (*10*, *11*).

Although the excitatory-inhibitory balance plays an important role for stability and information processing in cortical networks, it is not understood by which mechanisms this balance is established and maintained during ongoing sensory experiences. Inspired by recent experimental results (*7*), we investigated the hypothesis that synaptic plasticity at inhibitory synapses plays a central role in balancing the excitatory and inhibitory inputs a cell receives.

We simulated a single postsynaptic integrate-and-fire neuron receiving correlated excitatory and inhibitory input signals. The cell received input through 1000 synapses (Fig. 1B), which were divided into eight independent groups of 100 excitatory and 25 inhibitory synapses. All excitatory and inhibitory synapses within each group followed the same temporally modulated rate signal (time constant $\tau \sim 50$ ms) to mimic ongoing sensory activity (*13*, *16*). Spikes were generated from independent Poisson processes, leading to 125 different spike trains per signal. This architecture allowed each signal to reach the cell simultaneously through both excitatory and inhibitory synapses (Fig. 1B). To mimic glutamatergic and γ-aminobutyric acid (GABAergic) transmission, the synapses were conductance-based

[1]School of Computer and Communication Sciences and Brain-Mind Institute, École Polytechnique Fédérale de Lausanne, 1015 Lausanne EPFL, Switzerland. [2]CNRS, UMR 8119, Université Paris Descartes, 45 Rue des Saints Pères, 75270 Paris Cedex 06, France.

*These authors contributed equally to this work.
†To whom correspondence should be addressed. E-mail: tim.vogels@epfl.ch

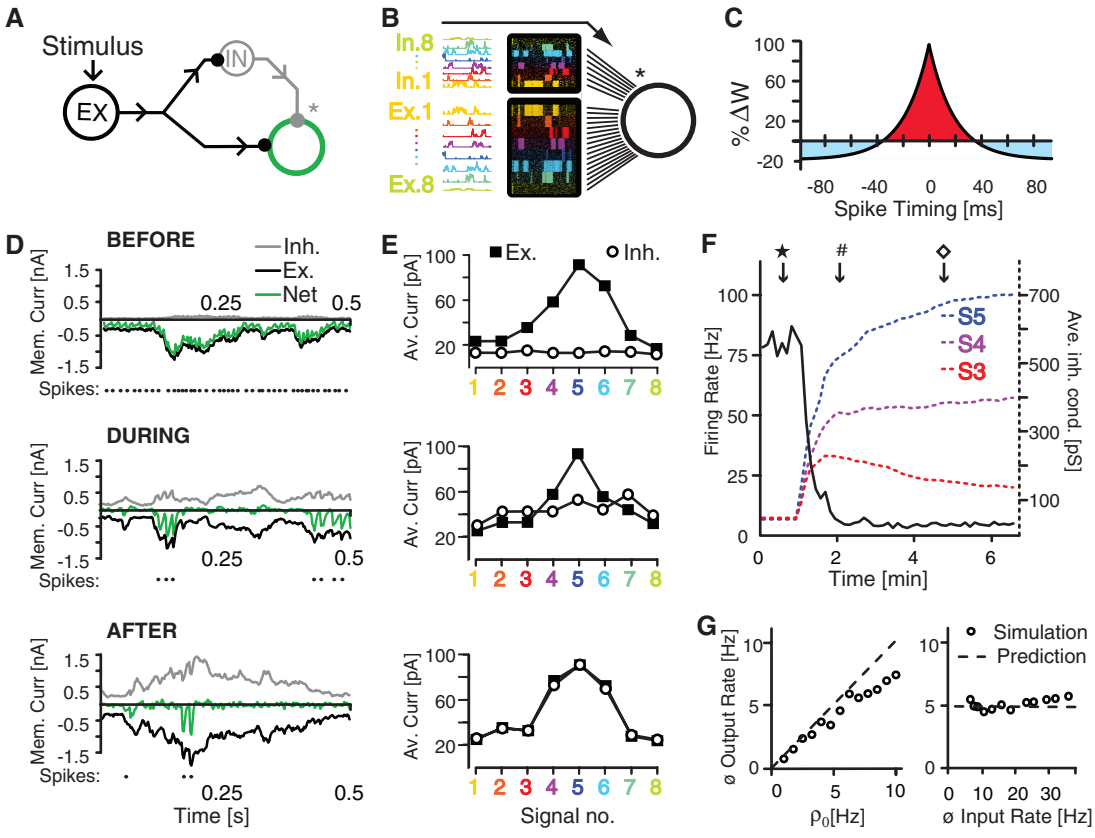

Fig. 1. Inhibitory synaptic plasticity balances excitation and inhibition. (**A**) Feedforward inhibition: Excitatory input reaches a target region through both direct excitation and indirect disynaptic inhibition. (**B**) Feedforward inhibition for a single postsynaptic cell: Eight groups of 100 excitatory and 25 inhibitory synapses each deliver spikes to a single postsynaptic cell. Spiking probabilities are homogeneous within the groups but vary in time, simulating eight separate (color-coded) signal channels that reach the cell simultaneously through excitatory and inhibitory synapses. (**C**) Spike-timing–dependent learning rule: Near-coincident pre- and postsynaptic spikes potentiate inhibitory synapses [marked with * in (A) and (B)], whereas every presynaptic spike causes synaptic depression. (**D**) Total excitatory (black), inhibitory (gray), and net (green) membrane currents before, during, and after inhibitory synaptic plasticity. The resulting spikes are indicated as dots underneath each current plot. (**E**) Excitatory and inhibitory membrane currents (black and white symbols, respectively) evoked by each signal channel, averaged over 4 s, before, during, and after inhibitory synaptic plasticity (top, middle, and bottom, respectively). (**F**) Temporal evolution of the postsynaptic firing rate (solid line) and the average synaptic weights of the inhibitory synapses associated with three representative signals (dotted lines). ★, #, and ◊ indicate the times at which the top, middle, and bottom graphs of (D) and (E) were recorded. (**G**) Average firing rate of the postsynaptic neuron after learning, plotted for different values of target firing rate ρ_0 (left) and different input rates (right). The dashed lines in both graphs show the analytical predictions.

with reversal potentials $V^E = 0$ mV and $V^I = -80$ mV and time constants $\tau^E = 5$ ms, and $\tau^I = 10$ ms for excitation and inhibition, respectively [see supporting online material (SOM)]. The strength of the inhibitory synapses was initially weak but could change according to a spike-timing–dependent plasticity rule, in which near-coincident pre- and postsynaptic spikes induce potentiation of the synapse (*17*–*19*). Additionally, every presynaptic spike leads to synaptic depression (*17*, *18*) (Fig. 1C). This learning rule can be summarized as

$$\Delta w = \eta(pre \times post - \rho_0 \times pre) \qquad (1)$$

where Δw denotes the change in synaptic efficacy, *pre* and *post* are the pre- and postsynaptic activity, η is the learning rate, and ρ_0 is a constant that acts as a target rate for the postsynaptic neuron (see SOM Sec. 2 for a mathematical analysis).

Whereas inhibitory synapses were plastic, the efficacies of the excitatory model synapses were fixed at the beginning of a simulation and left unchanged unless otherwise noted. Analogous to frequency- or orientation-tuned sensory neurons, excitatory synapses were tuned to have a preferred signal (Fig. 1E). Because all excitatory synapses were set to nonzero strengths, the postsynaptic neuron fired at high rates when the inhibitory synapses were weak at the beginning of a simulation (Fig. 1, D and E, top, and F). The resulting high number of pairs of pre- and postsynaptic spikes led to relatively indiscriminate strengthening of all inhibitory synapses (Fig. 1, D and E, middle) until excitatory and inhibitory membrane currents became approximately balanced and the postsynaptic firing rate was dramatically reduced (Fig. 1F). In this globally balanced state, only unbalanced excitatory signals led to coincident pairs of pre- and postsynaptic spikes, consequently strengthening underpowered inhibitory synapses. Those inhibitory synapses that were stronger than their excitatory counterparts kept the postsynaptic side unresponsive and were thus weakened (because of sole presynaptic firing) until they allowed postsynaptic spiking again. Over time, this led to a precise, detailed balance of excitatory and inhibitory synaptic weights for each channel (Fig. 1, D and E, bottom). In agreement with the mathematical analysis, the postsynaptic firing rate was determined mainly by the depression factor, ρ_0, but not by the average input firing rate to the postsynaptic neuron (Fig. 1G). The mechanism was robust to plausible delays of several milliseconds. However, because detailed balance requires a correlation between excitatory and inhibitory synaptic inputs, the balance deteriorated when the delay between excitation and inhibition increased to values larger than the autocorrelation time of the input signals and the coincidence time of the Hebbian learning rule, but global balance still persisted (fig. S2).

To investigate how the state of the balance affects the neuron's response properties, we presented a fixed stimulus sequence to the neuron (Fig. 2A) and compared the spiking response over 50 trials to the input rates of each signal. In the globally balanced state (Fig. 2B, top) in which inhibitory synapses were distributed so that excitation and inhibition were balanced only on average across all channels, the peristimulus time histogram (PSTH) faithfully reproduced the firing rates of the preferred signals. The other, non-preferred input signals evoked more inhibition than excitation and thus had no impact on the cell's firing behavior. An additional steplike input rate protocol, in which 100-ms-long pulses of various step sizes (Fig. 2C) were presented to one channel at a time, revealed that spiking responses are largely insensitive to stimulus intensity and indeed narrowly tuned to the preferred stimulus, giving rise to an all-or-none response (Fig. 2, D and E).

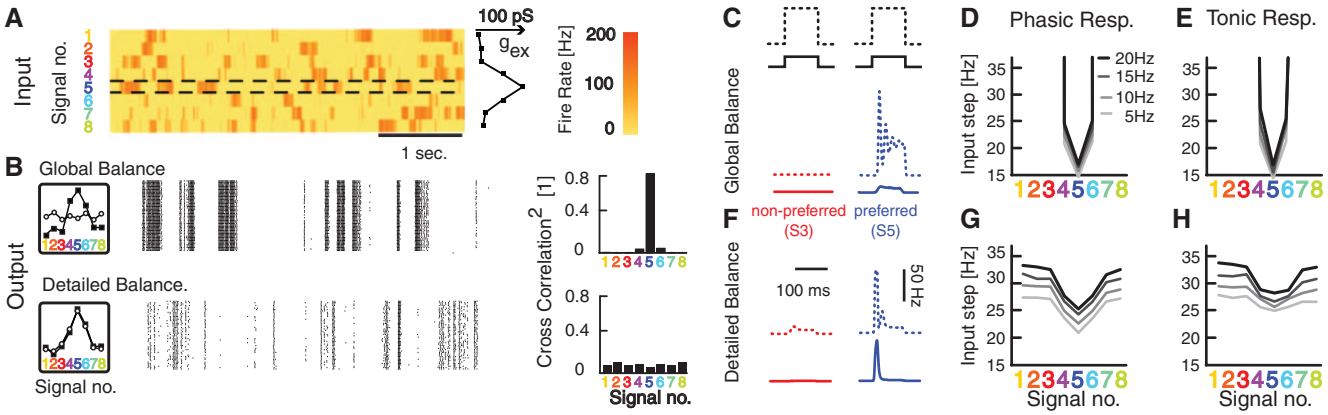

Fig. 2. Inhibitory synaptic plasticity sparsifies and democratizes receptive fields. (**A**) A fixed sequence of eight stimuli of varying firing rates is fed repetitively into a postsynaptic cell. Excitatory synapses are strength-tuned by signal group (see conductance graph on the right) so that signal five (marked also by dashed lines) is the preferred signal. (**B**) Postsynaptic spikes over 50 trials with globally or detailed balanced inhibitory synapses (top and bottom graphs, respectively) as indicated by the schematics on the left (compare with Fig. 1E). The normalized and squared cross-correlation coefficients between each input signal and the PSTH are also shown (right). (**C**) Schematic of a step stimulus delivered with large and small step sizes (solid and dotted black lines respectively); Sample PSTHs for nonpreferred (red) and preferred (blue) stimuli to both step sizes are shown for a globally balanced cell. (**D** and **E**) Iso-response contour lines of the postsynaptic cell in the globally balanced regime during the onset (phasic) (0 to 50 ms) (D) and tonic (50 to 100 ms) (E) parts of the response. (**F**) Sample responses for nonpreferred (red) and preferred (blue) stimuli to both step stimuli [as in (C)]. (**G** and **H**) Iso-response contour lines [as in (D) and (E)] for a detailed balanced cell.

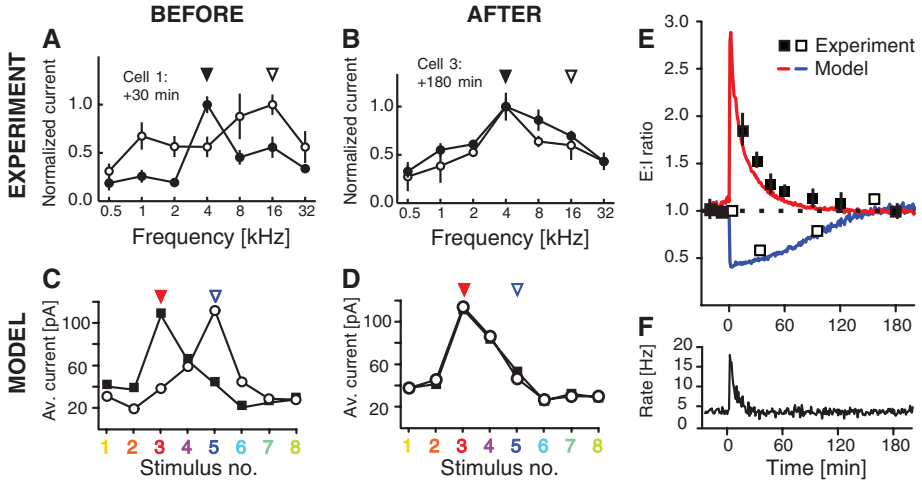

Fig. 3. Temporal dynamics of inhibitory plasticity, experiment, and model. Frequency-tuned excitatory and inhibitory membrane currents (black and white symbols, respectively) as recorded from pyramidal cells in the primary auditory cortex of adult rat (*7*) (**A**) 30 min and (**B**) 180 min after a stimulus protocol shifted the preferred frequency of the excitatory membrane currents from 16 to 4 kHz. Similarly stimulus-tuned input currents in a simulation (**C**) 30 min and (**D**) 180 min after (manually) changing the excitatory tuning curve. Solid and open arrowheads indicate the previous and the new preferred stimuli in all panels. (**E**) Summary plot of the ratios of excitatory and inhibitory current amplitudes of previously preferred stimuli and new preferred stimuli, as indicated in (A) to (D), in the experiment (open and solid symbols, respectively) and simulations (blue and red lines, respectively). (**F**) Firing rate of the simulated neuron over the time of the simulation in (E). Error bars indicate SEM. [(A), (B), and (E) adapted from (*7*) with permission]

In the detailed balanced state, the response of the cell was sparse (Fig. 2B, bottom) and reminiscent of experimental observations (*16*, *20*–*22*) across many sensory systems. Spikes were caused primarily by transients in the input signals, during which the faster dynamics of the excitatory synapses momentarily overcame inhibition. Sustained episodes of presynaptic firing, on the other hand, caused steady membrane currents that canceled each other and thus failed to evoke a reliable postsynaptic response. Seemingly indifferent to the tuning of the excitatory synapses, each signal contributed an equal part to the PSTH of the output signal, but the effect of the excitatory synaptic weights was uncovered by the steplike input protocol (Fig. 2F). The broad, graded responses (as opposed to all-or-none) to preferred and nonpreferred stimuli (Fig. 2, G and H) were in accord with experimental results (*5*, *7*, *8*, *23*, *24*) and confirm earlier theoretical studies arguing that sharp tuning is not a necessary feature for a sparse sensory representation (*25*, *26*). The sparsity of the response to each signal was a direct consequence of the detailed balance of correlated excitatory and inhibitory synapses as described above, not of the specificity of the tuning curve.

The self-organizing dynamics of inhibitory plasticity imply that the excitatory-inhibitory balance is maintained, even in the presence of ongoing excitatory plasticity (Fig. 3). Experiments (*7*) in which a stimulus alters the frequency tuning of excitatory input currents to pyramidal neurons in rat primary auditory cortex point in a similar direction: The disrupted cotuning of excitatory and inhibitory input currents (Fig. 3A) prompts a compensatory response that subsequently changes the amplitude of the inhibitory input currents. After 180 min, the cell returns to a cotuned state, albeit with a different preferred frequency (Fig. 3B). When we disturbed the cotuning of a simulated neuron in a similar way (Fig. 3C), inhibitory plasticity rebalanced the excitatory input currents (Fig. 3, D and E) and stabilized the output firing rates of the postsynaptic neurons (Fig. 3F). Quantitative agreement with the rebalancing dynamics observed in the experiment (for both synaptic depression and potentiation) was achieved by adjusting η, ρ_0, and the average firing rate of the inhibitory input neurons.

The learning rule for inhibitory synapses does not rely on a feedforward structure to achieve low firing rates. It simply matches excitatory and inhibitory synapses that show correlated activity. We therefore tested whether inhibitory plasticity was able to stabilize the dynamics of recurrent networks. In simulations of such networks (*13*)

with plastic inhibitory synapses that were initially weak (Fig. 4A), the resulting high firing rates and subsequent increase in inhibitory synaptic strengths caused by the plasticity rule indeed produced globally balanced input currents that led to a self-organized AI network state (Fig. 4B) with firing rates between 3 and 15 Hz.

We wondered whether it was possible to introduce associative memories to the stabilized network by strengthening specific excitatory connections within dedicated groups of neurons. First proposed by Hebb (*27*), such "cell assemblies" aim to provide a physiologically plausible explanation of how groups of neurons form a memory. Groups of highly connected neurons have since been successfully embedded into large spiking networks (*28*) and shown to self-sustain their activity without disrupting the global dynamics of the host network (*13*, *29*, *30*), but the parameter space that guarantees stable performance is narrow and tuning is arduous. The question has been raised how useful such memory attractors can be for long-term memory systems if only one of all stored memories can be active at a time, and potentially remains active for long periods of time, broadcasting the stored information into the network (*29*).

Inhibitory plasticity can solve some of these problems. After two arbitrarily chosen groups of excitatory neurons were turned into Hebbian assemblies by strengthening the excitatory connections within the groups fivefold, the assemblies temporarily fired at high rates and raised the background firing rate across the network (Fig. 4C). The resulting increase of coincident spike pairs caused inhibitory plasticity to increase the inhibitory synapses onto neurons in both assemblies until the global AI state was reestablished (Fig. 4D). After the excitatory and inhibitory inputs onto these neurons had been rebalanced, the firing rates of neurons in the cell assemblies became indistinguishable from the rest of the network, despite the imprinted memory traces in the excitatory synapses. Electrophysiological recordings of neuronal activity would thus not reveal the presence of a synaptic memory trace in this state.

Retrieval of previously quiescent memory items could be achieved by momentarily disrupting the balance within a cell assembly, for example, through additional excitatory input. It was sufficient to drive a small fraction of the cells of one assembly to reactivate all cells of that assembly. Notably, the recall was asynchronous and irregular, as indicated by low correlations between neurons and large variability of the interspike intervals (Fig. 4E). Although we embedded two overlapping assemblies into the network, only one was activated. The rest of the network remained nearly unperturbed in the AI state. Unlike traditional attractor networks, both assemblies could also be activated in unison by driving cells of both memories simultaneously (figs. S4 and S5), and their activity decayed to the background state after the stimulus was turned off.

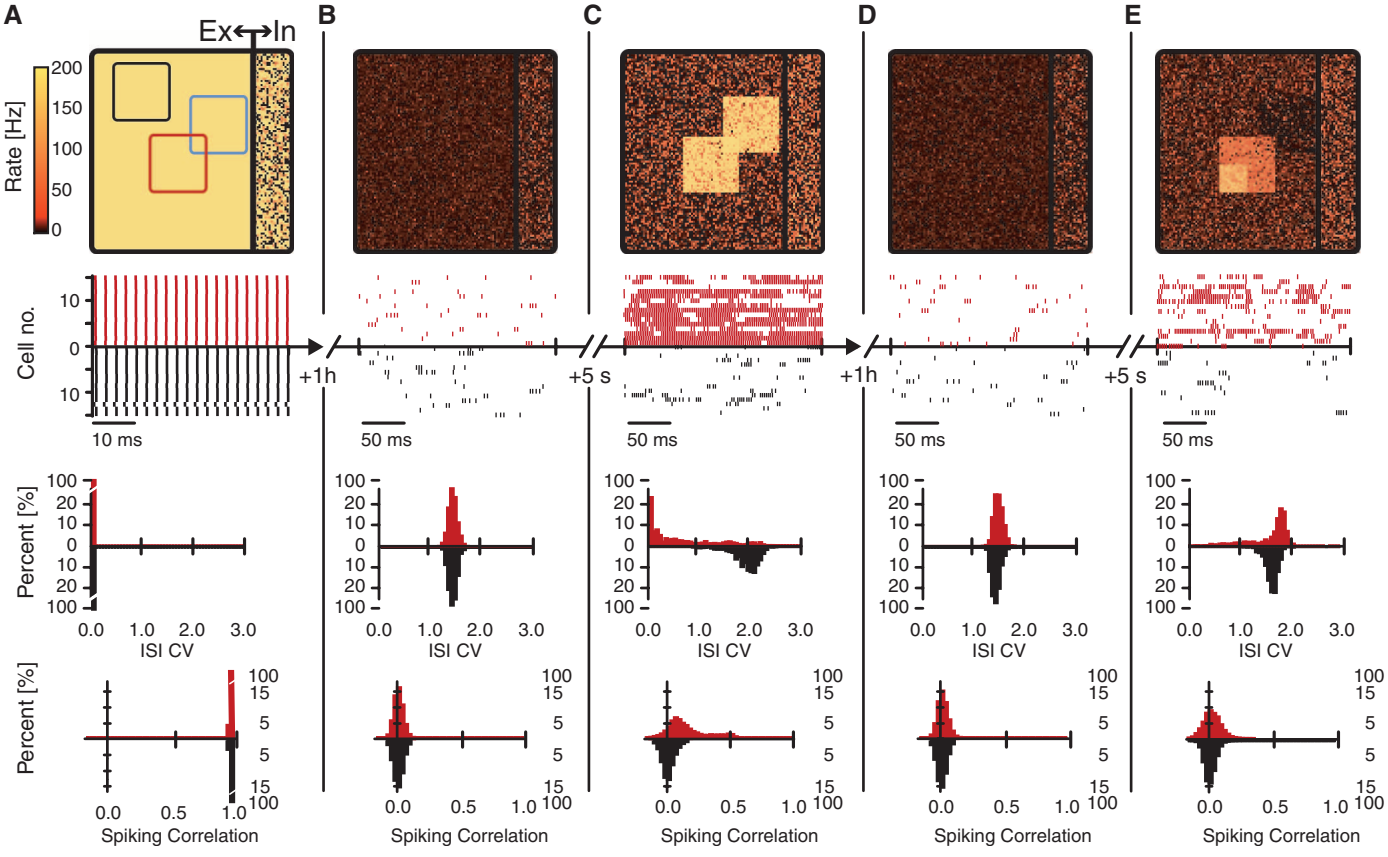

Fig. 4. Inhibitory plasticity in recurrent networks. Five consecutive snapshots of the momentary activity of a network of 10,000 integrate-and-fire cells with inhibitory plasticity. (**A**) Synchronous regular network dynamics with high firing rates at the beginning of the simulation with weak inhibitory synapses. (**B**) Establishment of the AI (steady) state with low firing rates through up-regulation of inhibitory synaptic weights by the synaptic plasticity rule. (**C**) The introduction of two synaptic memory patterns (cell assemblies) by fivefold increased excitatory synaptic weights between neurons outlined in red and blue in (A) leads to high firing rates. (**D**) Recovery of the AI state at low firing rates. (**E**) Memory retrieval through externally driving the lower left quarter of the red cell assembly with an additional excitatory stimulus. Each snapshot (A) to (E) shows (from top to bottom) the following: (i) The momentary (1-s) average firing rate of all neurons on a grid of 100^2 cells and separated into excitatory and inhibitory cells [left and right of the vertical line in (A), respectively]. Three groups of neurons play the role of either a cell assembly (red and blue outlines) or a control group (black outline). (ii) A raster plot of 30 randomly drawn neurons from one (red) cell assembly and the control group, indicated by a red and a black square in the plot above. (iii) The distributions of coefficients of variation of interspike intervals (ISI CVs) recorded from the neurons in the red and black groups. (iv) The distributions of spiking correlations between spike trains from neurons in the same designated groups. For methods and additional statistics, please see SOM.

Our results offer an explanation for how long-term memories can be stably embedded into networks as quiescent and overlapping Hebbian assemblies. Unlike previous studies, our network does not exhibit the behavior of an attractor network, in which activated cell assemblies will compete with each other and the winning pattern often exhibits persistent elevated activity. Instead, the network remains quiet unless the balance of one or more assemblies is modulated in favor of the excitation and returns to the background state when the modulation is turned off. We have shown this effect here by driving a subset of cells with an external stimulus, but there are several conceivable methods to modulate the balance of excitation and inhibition (SOM). The possibility to activate several patterns simultaneously allows the analog combination of patterns into larger composite memories. The capacity of storable and retrievable patterns is likely to depend on complex interactions between dynamics, size, and connectivity of the assemblies and the host network, as well as several other parameters.

We show that a simple, Hebbian plasticity rule on inhibitory synapses leads to robust and self-organized balance of excitation and inhibition that requires virtually no fine-tuning (figs. S6 to S9) and captures an unexpected number of recent experimental findings. The precision of the learned balance depends on the degree of correlation between the excitatory and the inhibitory inputs to the cell, ranging from a global balance in the absence of correlated inputs to a detailed balance for strong correlations. The phenomenon is robust to the shape of the learning rule, as long as it obeys two fundamental requirements: Postsynaptic activity must potentiate activated inhibitory synapses, whereas in the absence of postsynaptic firing inhibitory synapses must decay. Because the balance is self-organized, inhibitory plasticity will most likely maintain balance also in the presence of excitatory plasticity, as long as excitation changes more slowly than inhibition or when excitatory plasticity events are rare.

The mammalian brain hosts a wide variety of inhibitory cell types with different synaptic time scales, response patterns, and morphological target regions. Presumably, these cell types serve different functions, and consequently their synapses may obey several different plasticity rules (*31*). In our simplified model, the dynamics of inhibitory plasticity powerfully contributes to the functional state of cortical architectures and may have a strong impact on cortical coding schemes.

References and Notes

1. N. Brunel, *J. Comput. Neurosci.* **8**, 183 (2000).
2. C. van Vreeswijk, H. Sompolinsky, *Science* **274**, 1724 (1996).
3. M. Tsodyks, T. Sejnowski, *Network Comput. Neural Syst.* **6**, 111 (1995).
4. A. Renart *et al.*, *Science* **327**, 587 (2010).
5. M. Wehr, A. M. Zador, *Nature* **426**, 442 (2003).
6. M. Okun, I. Lampl, *Nat. Neurosci.* **11**, 535 (2008).
7. R. C. Froemke, M. M. Merzenich, C. E. Schreiner, *Nature* **450**, 425 (2007).
8. J. de la Rocha, C. Marchetti, M. Schiff, A. D. Reyes, *J. Neurosci.* **28**, 9151 (2008).
9. B. K. Murphy, K. D. Miller, *Neuron* **61**, 635 (2009).
10. Y. Shu, A. Hasenstaub, D. A. McCormick, *Nature* **423**, 288 (2003).
11. T. P. Vogels, L. F. Abbott, *Nat. Neurosci.* **12**, 483 (2009).
12. W. Gerstner, *Neural Comput.* **12**, 43 (2000).
13. T. P. Vogels, L. F. Abbott, *J. Neurosci.* **25**, 10786 (2005).
14. A. Kumar, S. Rotter, A. Aertsen, *Nat. Rev. Neurosci.* **11**, 615 (2010).
15. J. Cafaro, F. Rieke, *Nature* **468**, 964 (2010).
16. T. Hromádka, M. R. Deweese, A. M. Zador, *PLoS Biol.* **6**, e16 (2008).
17. M. A. Woodin, K. Ganguly, M. M. Poo, *Neuron* **39**, 807 (2003).
18. V. Kilman, M. C. W. van Rossum, G. G. Turrigiano, *J. Neurosci.* **22**, 1328 (2002).
19. K. Hartmann, C. Bruehl, T. Golovko, A. Draguhn, *PLoS One* **3**, e2979 (2008).
20. M. R. DeWeese, M. Wehr, A. M. Zador, *J. Neurosci.* **23**, 7940 (2003).
21. H. Yao, L. Shi, F. Han, H. Gao, Y. Dan, *Nat. Neurosci.* **10**, 772 (2007).
22. S. Crochet, J. F. Poulet, Y. Kremer, C. C. Petersen, *Neuron* **69**, 1160 (2011).
23. L. M. Aitkin, D. J. Anderson, J. F. Brugge, *J. Neurophysiol.* **33**, 421 (1970).
24. I. O. Volkov, A. V. Galazjuk, *Neuroscience* **43**, 307 (1991).
25. P. Seriès, P. E. Latham, A. Pouget, *Nat. Neurosci.* **7**, 1129 (2004).
26. J. Beck, V. R. Bejjanki, A. Pouget, *Neural Comput.* **23**, 1484 (2011).
27. D. Hebb, *The Organization of Behavior; a Neuropsychological Theory* (Wiley-Interscience, New York, 1949).
28. W. Gerstner, R. Ritz, J. L. van Hemmen, *Biol. Cybern.* **69**, 503 (1993).
29. D. J. Amit, N. Brunel, *Cereb. Cortex* **7**, 237 (1997).
30. A. Renart, R. Moreno-Bote, X.-J. Wang, N. Parga, *Neural Comput.* **19**, 1 (2007).
31. M. A. Woodin, A. Maffei, *Inhibitory Synaptic Plasticity* (Springer, New York, 2010).

Acknowledgments: Research was supported by Swiss National Science Foundation grant no. 200020 13287 (Coding Characteristics) and CRSIK0 122697 (Sinergia). Additionally, T.P.V. was supported by the European Community's Seventh Framework Marie Curie International Reintegration grant no. 268436, and H.S. and F.Z. by the European Community's Seventh Framework Program under grant agreement no. 243914 (BRAIN-I-NETS) and 237955 (FACETS-ITN), respectively. C.C. received additional support from a French National Science grant ANR-08-SYSC-005. Thanks to G. Hennequin and A. Woodruff for helpful discussions.

Supporting Online Material

www.sciencemag.org/cgi/content/full/science.1211095/DC1
Materials and Methods
SOM Text
Figs. S1 to S10
Tables S1 and S2
References (*31*–*49*)

13 July 2011; accepted 20 October 2011
Published online 10 November 2011;
10.1126/science.1211095

Autophagy-Dependent Anticancer Immune Responses Induced by Chemotherapeutic Agents in Mice

Mickaël Michaud,[1,2,3]* Isabelle Martins,[1,2,3]* Abdul Qader Sukkurwala,[1,2,3] Sandy Adjemian,[1,2,3] Yuting Ma,[2,3,4,5] Patrizia Pellegatti,[6] Shensi Shen,[1,2,3] Oliver Kepp,[1,2,3] Marie Scoazec,[2,7] Grégoire Mignot,[8,9] Santiago Rello-Varona,[1,2,3] Maximilien Tailler,[1,2,3] Laurie Menger,[1,2,3] Erika Vacchelli,[1,2,3] Lorenzo Galluzzi,[1,2,3] François Ghiringhelli,[8,9] Francesco di Virgilio,[6] Laurence Zitvogel,[2,3,4,5]† Guido Kroemer[1,2,10,11,12]†

Antineoplastic chemotherapies are particularly efficient when they elicit immunogenic cell death, thus provoking an anticancer immune response. Here we demonstrate that autophagy, which is often disabled in cancer, is dispensable for chemotherapy-induced cell death but required for its immunogenicity. In response to chemotherapy, autophagy-competent, but not autophagy-deficient, cancers attracted dendritic cells and T lymphocytes into the tumor bed. Suppression of autophagy inhibited the release of adenosine triphosphate (ATP) from dying tumor cells. Conversely, inhibition of extracellular ATP-degrading enzymes increased pericellular ATP in autophagy-deficient tumors, reestablished the recruitment of immune cells, and restored chemotherapeutic responses but only in immunocompetent hosts. Thus, autophagy is essential for the immunogenic release of ATP from dying cells, and increased extracellular ATP concentrations improve the efficacy of antineoplastic chemotherapies when autophagy is disabled.

Transplantable or primary murine cancers respond to chemotherapy with anthracyclines or oxaliplatin much more efficiently when they grow in syngenic immunocompetent mice than in immunodeficient hosts (*1*, *2*). Similarly, clinical studies indicate that severe lymphopenia negatively affects the chemotherapeutic response of solid cancers (*3*), and immune defects are negative predictors of the response to chemotherapy with anthracyclines or oxaliplatin (*2*, *4*, *5*). Apparently, some successful chemotherapeutics can induce a type of tumor cell stress and death that is immunogenic (*6*–*8*), implying that the patient's dying cancer cells serve as a therapeutic vaccine that stimulates an antitumor immune response, which in turn can control residual cancer cells (*9*, *10*). Immunogenic cell death is characterized by the preapoptotic exposure of calreticulin (CRT) on the cell surface (*11*), postapoptotic release of the chromatin-binding protein high mobility group B1 (HMGB1) (*2*), and secretion of adenosine triphosphate (ATP) (*4*).