
THEORY OF NON-LINEAR SPIKE-TIME-DEPENDENT PLASTICITY

By
Jean-Pascal Pfister

SUBMITTED IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
AT THE
SWISS FEDERAL INSTITUTE OF TECHNOLOGY LAUSANNE
LAUSANNE, SWITZERLAND
JULY 2006

Ralf Schneggenburger (president of the jury)
Wulfram Gerstner (research supervisor)
Carl Petersen (referee)
Walter Senn (referee)
Haim Sompolinsky (referee)

© Copyright by Jean-Pascal Pfister, 2006

A Celui qui donne le souffle.

Abstract

A fascinating property of the brain is its ability to continuously evolve and adapt to a constantly changing environment. This ability to change over time, called plasticity, is mainly implemented at the level of the connections between neurons (i.e. the synapses). So if we want to understand the ability of the brain to evolve and to store new memories, it is necessary to study the rules that govern synaptic plasticity.

Among the large variety of factors which influence synaptic plasticity, we focus our study on the dependence upon the precise timing of the pre- and postsynaptic spikes. This form of plasticity, called Spike-Timing-Dependent Plasticity (STDP), works as follows: if a presynaptic spike is elicited before a postsynaptic one, the synapse is up-regulated (or potentiated) whereas if the opposite occurs, the synapse is down-regulated (or depressed).

In this thesis, we propose several models of STDP which address the two following questions: (1) what is the functional role of a synapse which elicits STDP and (2) what is the most compact and accurate description of STDP? In the first two papers contained in this thesis, we show that in a supervised scenario, the best learning rule which enhances the precision of the postsynaptic spikes is consistent with STDP. In the three following papers, we show that the information transmission between the input and output spike trains is maximized if synaptic plasticity is governed by a rule similar to STDP. Moreover, we show that this infomax principle added to an homeostatic constraint leads to the well-known Bienenstock-Cooper-Munro (BCM) learning rule. Finally, in the last two papers, we propose a phenomenological model of STDP which considers not only pairs of pre- and postsynaptic spikes, but also triplets of spikes (e.g. 1 pre and 2 post or 1 post and 2 pre). This model can reproduce a lot of experimental results and can be mapped to the BCM learning rule.

Keywords: Spike-Timing-Dependent Plasticity, Hebbian Learning, Long-Term Potentiation, Information Theory, BCM Learning Rule, Triplets of Spike.

Une propriété fascinante du cerveau est sa capacité d'évoluer et de s'adapter continuellement à un environnement qui change constamment. Cette capacité de se modifier au cours du temps, appelée plasticité est principalement implémentée au niveau des connections entre les neurones (les synapses). Donc, si nous désirons comprendre la capacité du cerveau à changer et mémoriser de nouveaux éléments, il est nécessaire d'étudier les règles qui gouvernent la plasticité synaptique.

Parmi la grande variété de facteurs qui influencent la plasticité synaptique, nous focalisons notre étude sur la dépendance du temps de tir des potentiels d'action pre- et postsynaptiques. Cette forme de plasticité, appelée plasticité synaptique à modulation temporelle relative (STDP ou Spike-Timing-Dependent Plasticity) fonctionne de la façon suivante: si un potentiel d'action presynaptique a lieu avant un potentiel d'action postsynaptique, la synapse est renforcée (ou potentialisée); si au contraire, l'ordre des potentiels d'action est inversé, la synapse est affaiblie (ou déprimée).

Dans cette thèse, nous proposons plusieurs modèles de la STDP qui adrent les deux questions suivantes: (1) quel est le rôle fonctionnel d'une synapse suivant une règle STDP et (2) quelle est la description la plus compacte et la plus précise de la STDP? Dans les deux premiers papiers contenus dans cette thèse, nous montrons que dans un contexte supervisé, la meilleure règle d'apprentissage qui augmente la précision des potentiels d'action postsynaptiques est cohérente avec la STDP. Dans les trois papiers qui suivent, nous montrons que la transmission d'information entre les trains de potentiels d'action d'entrée et de sortie est maximisée si la plasticité synaptique est gouvernée par une règle similaire à la STDP. De plus, nous montrons que ce principe infomax ajouté à une contrainte homéostatique conduit à la règle d'apprentissage appelée Bienenstock-Cooper-Munro (BCM). Finalement, dans les deux derniers papiers, nous proposons un modèle phénoménologique de la STDP qui considère non seulement des paires de potentiels d'action pre- et postsynaptiques, mais aussi des

triplets de potentiels d'action (1 pre et 2 post ou 1 post et 2 pre). Ce modèle peut reproduire beaucoup de résultats expérimentaux et peut être relié à la règle d'apprentissage BCM.

Mot-clés: Plasticité synaptique à modulation temporelle relative (STDP), Apprentissage Hebbien, Potentialisation à long-terme, Théorie de l'information, Règle d'apprentissage BCM, Triplets de potentiels d'action.

Contents

Abstract	v
Version résumée	vii
Acknowledgments	xiii
Foreword	xv
1 Introduction	1
1.1 Experiments on Synaptic Plasticity	2
1.2 Models of STDP	5
1.2.1 Phenomenological Models	7
1.2.2 Biophysical Models	11
1.2.3 Optimality Models	16
1.3 Aims and achievements	19
2 Publications	25
2.1 Paper I	25
2.2 Paper II	35
2.3 Paper III	66
2.4 Paper IV	76
2.5 Paper V	100
2.6 Paper VI	133
2.7 Paper VII	143
3 Discussion	173
3.1 Summary of the Results	173

3.2	Limitations	177
3.3	Propositions for Future Developments	178
A	Volterra Expansion	183
A.1	All-to-All Interaction Scheme	183
A.2	Nearest-Past-Spike Interaction Scheme	184
B	Comparison with Other Models	187
B.1	Notes on the Tempotron	187
B.2	Triplets and the SMT Learning Rule	189
	Bibliography	194
	Curriculum Vitæ	203

Acknowledgments

I am very thankful to my family, friends and colleagues. In particular, I would like to thank the following persons:

My wife Natacha for her love and her continuous and priceless support, my son Corentin for his uncountable smiles and laughs.

My parents Jean-François and Maria for their interest in my work and for the great opportunity they gave me to study at the EPFL, my sister Marielle and her husband Michael, and my brother Pierre-Daniel for being much more than brother and sister.

Everybody in the lab. Wulfram Gerstner for introducing me to the fascinating field of synaptic plasticity, for the great time he spent in the supervision and the freedom he gave me at the end of the PhD. Brigitte Ramuz and lately Nicole Rauser for their devotion and their great help for logistic questions. Pierre-Edouard Sottas and Silvio Borer for their supervision during the diploma thesis. Angelo Arleo for advising me to do a Phd, Alix Hermann for developing noise models I used in my thesis, Edgar Oropesa Farras and especially José del R. Millán for introducing me to the EEG world, Arnaud Tonnelier for the discussions about rock climbing and ice climbing, Renaud Jolivet and Julien Mayor for transforming standard breaks into lively and enjoyable events, Taro Toyozumi for the nice collaboration and the challenging discussions about STDP, Magnus Richardson for having co-founded the LCN jogging club and for his precious academic and scientific advices, Thomas Strösslin for having discovered the best jogging path, Yuval Aviel and recently André Longtin for having shared the office, Ricardo Chavarriaga for participating to a special event of the LCN jogging club: the 10 km of Lausanne, Denis Sheynikhovich for his precious help when computers became capricious, Daniela Helbig for gliding discussions, François Fleuret for the endless discus-

sions about faith and evolution, Brice Bathellier for joining the LCN parents club, Claudia Cloppath for her careful reading of the manuscript, Laurence Meylan for her courage to have been a woman at the LCN, Romain Brette and Naoki Masuda for visiting the LCN, Laurent Badel and Guedi Lusky for scientific discussions.

Last, but not least, I would like to thank my friends Fabien Bourgeois and Stephan Gamper for their weekly prayers and advices for my PhD.

This work was supported by the Swiss National Science Foundation (grant numbers 200021-100215 (1.4.03-31.4.05) and 200020-108097 (1.4.05-31.3.06) and by the École Polytechnique Fédérale de Lausanne (EPFL).

Foreword

The following thesis was done at the Laboratory of Computational Neuroscience (Brain and Mind Institute and School of Computer and Communication Science, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland) between August 2002 and May 2006.

All the simulations were performed on a Sun Workstation and a Sun Server (*Sun Microsystems, Inc.*, Santa Clara, CA, USA). Original programs were written in MATLAB (*The MathWorks, Inc.*, Natick, MA, USA).

Most of the figures were drawn with MATLAB and some with Xmgrace (*Free Software Foundation, Inc.*, Boston, MA, USA). Schematic drawings were done with Xfig. Multi-panel figures were assembled directly with \LaTeX .

This thesis is a compilation of seven articles (see chapter 2) that are labeled with roman letters (I, II, . . . , VII). This thesis were prepared with \LaTeX and compiled with PDF \LaTeX .

CHAPTER 1

Introduction

One of the most remarkable features of the nervous system is its ability to change over time. This ability to evolve, named plasticity, is especially clear during development. This early developmental plasticity is mainly implemented by new physical connections between neurons. Even if plasticity is mainly present during the early stage of development, it is clear that at the adult stage, the nervous system can still acquire new behaviors and store new memories. It is now widely acknowledged that memories depend on a persistent change of synaptic transmission, but it has not been always the case. Indeed, in his own words Lashley (1924) stated that

Among the many unsubstantiated beliefs concerning the physiology of the learning process, none is more widely prevalent than the doctrine that the passage of the nerve impulse through the synapse somehow reduces synaptic resistance and leads to the fixation of a new habit.

Many years later, Lashley's student, D.O. Hebb took the risk to formulate as a postulate what his professor and others heavily criticized. His point of view is that if we assume a structural change at the level of the synapse, i.e. "synaptic activity makes the synapse more readily traversed", lasting memory is possible. More precisely, in his monograph, "The Organization of Behavior", which was highly influential, D.O. Hebb (1949) postulated that

When an axon of cell A is near enough to excite cell B and repeatedly or per-

sistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.

In other words, according to this postulate, correlated activity of neuron A and B increases the strength of the synapse (or weight) between A and B. Even if the notion of correlation-based learning has been formulated before Hebb (James 1890)¹, it is often called Hebbian learning. An important aspect of Hebb's postulate is its intrinsic notion of "locality". Only locally available information, namely the pre- and postsynaptic activities can trigger a weight change. Without such an unsupervised mechanism which adapts continuously the weights, it would be impossible to tune individually each of the 10^{14} synapses in the human brain.

Clearly, if synapses have only a mechanism for potentiation and none for depression, they will quickly saturate to their maximal value. Therefore, a mechanism which decreases the synaptic strength is necessary. This is why Stent (1973) proposed to add to Hebb's original postulate that "when the presynaptic axon of cell A repeatedly and persistently fails to excite the postsynaptic cell B while cell B is firing under the influence of other presynaptic axons, metabolic changes take place in one or both cells such that A's efficiency, as one of the cells firing B, is decreased".

1.1 Experiments on Synaptic Plasticity

At this stage, there was no physiological evidence for such a mechanism. It is only in 1973 (Bliss and Gardner-Medwin 1973; Bliss and Lomo 1973) that has been published the detailed description of long-term potentiation (LTP)². In LTP, the strength of a synapse increases for a prolonged period if the synapse is activated by a brief but intense stimulation. Initially, LTP was elicited in intact anesthetized rabbit (Bliss and Gardner-Medwin 1973) and later in awake rabbit (Bliss and Lomo 1973). This major discovery nicely confirmed the intuition of Hebb and inspired an enormous number of researchers. In 2006, more than 14'000 papers have

¹James (1890) stated that "When two elementary brain processes have been active together or in immediate succession, one of them, on re-occurring, tends to propagate its excitement into the other". Note that the temporal order of the processes are not mentioned here.

²The very first results on LTP have been already published in 1966 by Lomo (1966) in an abstract.

been written in relation to LTP. This phenomenon continues to fascinate the neuroscience community. As Bliss and Lynch (1988) said

No matter how often one has witnessed the phenomenon, it is impossible not to retain a sense of amazement that such modest stimulation can produce so immediate, so profound, and so persistent an effect.

Why such a fascination for this phenomenon? For our brains to store memories, it is necessary that short stimulations induce structural changes that persist for a long time. This is precisely what LTP does. The correspondence between the necessary condition and the observed phenomenon fascinates. See the introductory remarks of Bliss, Collingridge, and Morris (2003) as well as the entire issue on long-term potentiation.

It is only 20 years after the discovery of LTP that the opposite phenomenon, i.e. long-term depression, has been elicited for the first time. Although it was known that potentiation could be reversed by a phenomenon termed “depotentialization” (Barrionuevo et al. 1980), the key advance has been to show that synapses could be depressed by low-frequency stimulation (Dudek and Bear 1992): this phenomenon is now termed long-term depression (LTD).

Hebbian learning does not contain only the notion of locality in space. The concept of locality in time or simultaneity and even the idea of causality is also present in Hebb’s postulate (“... takes part in firing it ...”), but the time scale of simultaneity is not mentioned explicitly.

The first experimental work on this notion of simultaneity has been done by Levy and Stewart (1983). They stimulated extracellularly a burst of 8 spikes at 400 Hz and after a given delay stimulated the postsynaptic neuron with a similar burst of spikes. They observed that the size of the excitatory postsynaptic potential (EPSP) increased if the presynaptic burst occurred before the postsynaptic one and decreased if the timing is reversed. One important outcome of this study is that potentiation or depression occurs in a limited time window of the order of ± 100 ms. Although this type of experiment is different from the ones described below, named Spike-Timing-Dependent Plasticity (STDP), it shares most of the features and gives an experimental confirmation of Hebb’s postulate.

In order to achieve the behaviorally required temporal precision in the barn owl auditory system, Gerstner, Kempter, van Hemmen, and Wagner (1996) assumed a specific form of simultaneity between the pre- and postsynaptic spikes within which potentiation or depression

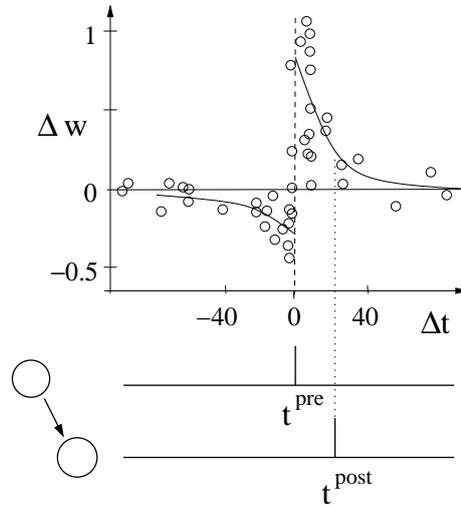


Figure 1.1: Synaptic plasticity depends on the relative timing $\Delta t = t^{post} - t^{pre}$ between pre- and postsynaptic spikes. If the presynaptic spike at time t^{pre} occurs before the postsynaptic spike at time t^{post} , the synaptic strength w increases whereas depression occurs if the timing is reversed. Data redrawn from Bi and Poo (1998).

can occur. More precisely, this learning window has been defined as follows. If a single presynaptic spike is before a postsynaptic one, maximal potentiation occurs whereas depression is elicited if the timing is reversed. With this mechanism, synapses from presynaptic neurons that “take part in firing” the postsynaptic one are potentiated whereas synapses that elicit a postsynaptic potential too late (i.e. after the postsynaptic spike) are depressed.

Shortly after this learning window had been introduced on theoretical grounds, a lot of in vitro experiments (Bell et al. 1997; Markram et al. 1997; Bi and Poo 1998; Zhang et al. 1998) showed that synaptic strength depends on the precise timing between the pre- and postsynaptic spikes. This phenomenon is now termed Spike-Timing-Dependent Plasticity (STDP). See Fig. 1.1 for the learning window obtained by Bi and Poo (1998). Potentiation occurs if pre- and postsynaptic spikes arrive in a causal order, i.e. pre-before-post, whereas depression is elicited for acausal pre- and postsynaptic spikes, i.e. pre-after-post.

Although STDP is now a well established phenomenon for in vitro experiments, it is however less clear whether and how STDP operates in vivo. Extracellular recordings in the visual cortex of anesthetized cats show that sequential visual stimulation induces a shift of

orientation tuning of the cells (Frégnac et al. 1992; Yao and Dan 2001; Yao and Dan 2004) and a shift of the cell's receptive field (Fu et al. 2002) in a way which is consistent with STDP. The problem with extracellular recordings is that the precise location where plasticity occurs is unclear. Recently Meliza and Dan (2006) used whole-cell recording techniques *in vivo* to show that the response of a single neuron to visual stimulus changes consistently with STDP. It should be noted that the LTP or LTD reported in (Yao and Dan 2001) lasts only of the order of 10 minutes and then vanishes. Therefore, it still remains to be elucidated if the observed *in vivo* phenomenon is identical to the one observed *in vitro* which lasts for hours.

1.2 Models of STDP

Since this thesis is about modeling STDP, it is necessary to be aware of the different types of models one can elaborate on a given phenomenon, STDP in our case; see Fig. 1.2. Even if there are various ways to classify the different types of models, we propose here to classify the existing models of STDP into three different groups: the phenomenological models, the biophysical models and the optimal models (see the preface of Dayan and Abbott (2001)).

What, how and why are the three questions addressed by those three types of model. The phenomenological models (or descriptive models) are intended to describe the phenomenon in a precise, but compact way, i.e. the number of required parameters should be as small as possible while describing the phenomenon as accurately as possible. Biophysical models (or mechanistic models) belong to a bottom-up approach and aim to explain how the mechanism operates. The goal of those models is to create a link between biophysical quantities such as the calcium concentration, the number of released vesicles, the ion current through NMDA channels, ... and the observed phenomenon. Finally optimal models (or interpretive models) have an opposite perspective: the role of such top-down models is to explain the functional role of the phenomenon. Practically, this is often done by showing that the phenomenon maximizes a given quantity which corresponds to a given function. In the following subsections, we will discuss the strengths and weaknesses of those different approaches.

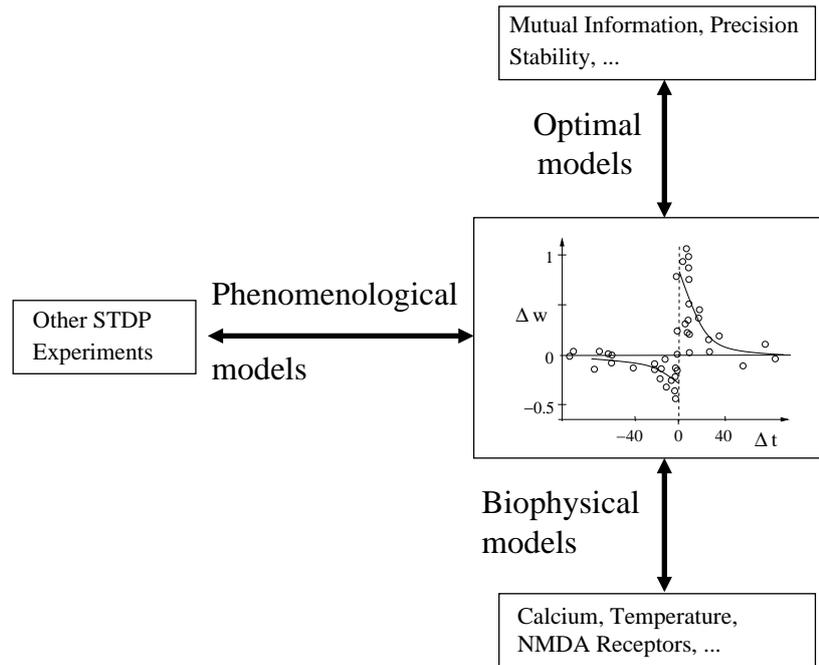


Figure 1.2: Classification of the different types of model for Spike-Timing Dependent Plasticity. The phenomenological models summarize the experimental data, giving a minimal description of the observed phenomena. Biophysical models create a link between physical quantities and the STDP data. Optimal models (or interpretive models) of STDP propose a functional role of the synapse. This is often done by showing that STDP maximizes a given quantity, which corresponds to the function of the synapse.

1.2.1 Phenomenological Models

A first observation about the results of STDP experiments shows that data are extremely noisy. This can be observed in Fig. 1.1 as well as in all papers dealing with STDP. This can have several implications on models of STDP. Indeed, when we compare plasticity experiments with other electrophysiological experiments such as input-output characterization of a neuron (Hodgkin and Huxley 1952), we realize that the high variability in the data does not allow us to develop a canonical model such as the Hodgkin and Huxley model. As a consequence, the models should be as simple as possible and therefore contain as few parameters as possible. This is exactly the approach of the phenomenological models which are able to create a link between similar types of experiments controlling the timing of the pre- and postsynaptic spikes (Sjöström et al. 2001; Froemke and Dan 2002; Wang et al. 2005; Froemke et al. 2006)

Volterra expansion

Let us present now a possible conceptual framework for phenomenological models. Since those models should be as simple as possible, let us consider a framework in which we can consider the simplest models of STDP and then incrementally increase the level of complexity. This can be done by assuming that the instantaneous weight change \dot{w} is given by the Volterra expansion (Volterra 1930; Gerstner and Kistler 2002) of an unknown functional $H[X, Y]$ which depends on the history of the presynaptic spike train $X(t) = \sum_{t^{pre}} \delta(t - t^{pre})$ and the postsynaptic spike train $Y(t) = \sum_{t^{post}} \delta(t - t^{post})$ where $\delta(s)$ is the Dirac function and t^{pre} and t^{post} are respectively the pre- and postsynaptic spike times:

$$\dot{w}(t) = H[X, Y] \tag{1.2.1}$$

Clearly, synaptic plasticity depends also on other neuronal quantities such as calcium concentration (Malenka et al. 1988; Lisman 1989; Shouval et al. 2002; Karmarkar and Buonomano 2002), the depolarization (Artola et al. 1990; Sjöström et al. 2001) and of course the current value of the weight w (Bi and Poo 1998; van Rossum et al. 2000; Kistler and van Hemmen 2000; Gütiig et al. 2003). All those variables could be introduced in the functional H . For the sake of simplicity, let us omit them from now on. If we further assume that the weight changes are instantaneous and occur whenever a pre- or a postsynaptic spike

is emitted, we can write

$$\dot{w}(t) = H_0 + X(t)F[X, Y] + Y(t)G[X, Y] \quad (1.2.2)$$

where F and G are unknown functional of the pre- and postsynaptic spike trains and H_0 is a constant. Let us perform a Volterra expansion of those functionals. This gives

$$\begin{aligned} F([X, Y]) &= F_1^x + \int_0^\infty F_2^{xx}(s)X(t-s)ds + \int_0^\infty F_2^{xy}(s)Y(t-s)ds \\ &+ \int_0^\infty \int_0^\infty F_3^{xxx}(s, s')X(t-s)X(t-s')ds'ds \\ &+ \int_0^\infty \int_0^\infty F_3^{xxy}(s, s')X(t-s)Y(t-s')ds'ds \\ &+ \int_0^\infty \int_0^\infty F_3^{xyy}(s, s')Y(t-s)Y(t-s')ds'ds + \dots \end{aligned} \quad (1.2.3)$$

Similarly, the expansion of G yields

$$\begin{aligned} G([X, Y]) &= G_1^y + \int_0^\infty G_2^{xy}(s)X(t-s)ds + \int_0^\infty G_2^{yy}(s)Y(t-s)ds \\ &+ \int_0^\infty \int_0^\infty G_3^{xxy}(s, s')X(t-s)X(t-s')ds'ds \\ &+ \int_0^\infty \int_0^\infty G_3^{xyy}(s, s')X(t-s)Y(t-s')ds'ds \\ &+ \int_0^\infty \int_0^\infty G_3^{yyy}(s, s')Y(t-s)Y(t-s')ds'ds + \dots \end{aligned} \quad (1.2.4)$$

Note that the upper index in functions represents the type of interaction. For example, G_3^{xyy} refers to a triplet interaction consisting of 1 pre- and 2 postsynaptic spikes. We remark that the G_3^{xxy} term could correspond to a *pre-post-post* sequence as well as a *post-pre-post* sequence. See appendix A for the compact description of the Volterra expansion.

The interest of such a framework is that it is possible to classify a lot of existing models within this framework. For example, the first model of STDP (Gerstner et al. 1996) considers the pair terms, i.e. F_2^{xy} and G_2^{xy} as well as a constant term H_0 . Kempter et al. (1999) and Kistler and van Hemmen (2000) modeled STDP with the pair terms F_2^{xy} , G_2^{xy} and added a pure presynaptic term F_1^x as well as a pure postsynaptic term G_1^y . Most of the STDP models, like the ones of Song et al. (2000) and Roberts (1999) only consider the pair terms F_2^{xy} , G_2^{xy} .

Since not all STDP models can be fitted in this framework, it is possible to broaden it in order to include new classes of STDP models. As given by the Eqs. (1.2.2), (1.2.3) and (1.2.4), whenever there is a pre- or postsynaptic spike, it interacts with all previous pre- or postsynaptic spikes. We will call this the *All-to-All* interaction scheme. There is no *a priori* reason to think that say each presynaptic spike should interact in the same way with all previous postsynaptic spikes and vice versa. This is why other schemes of interaction have been proposed (van Rossum et al. 2000; Bi 2002; Sjöström et al. 2001; Izhikevich 2003; Burkitt et al. 2004; Meffin et al. 2006); see Fig. 1.3.

For the sake of simplicity for the following discussion, let us restrict to the case of pair terms since the same logic applies for the higher order terms. In a *Nearest-Past-Spike* interaction scheme (see papers VI and VII)³, each presynaptic spike interacts only with the previous postsynaptic spike and each postsynaptic spike with the previous postsynaptic one. The *Nearest-Future-Spike* interaction scheme (van Rossum et al. 2000) (called “nearest neighbor” by (Bi 2002)) considers the opposite case. Each presynaptic spike interacts with the next postsynaptic spike and vice versa. In the *Presynaptic-Centric* case (which corresponds to the “output restricted” case of Burkitt, Meffin, and Grayden (2004)), each presynaptic spike interacts only with the last and the next postsynaptic spikes (Izhikevich 2003). Similarly, in the *Postsynaptic-Centric* case (which corresponds to the “input restricted” case of Burkitt, Meffin, and Grayden (2004)), each postsynaptic spike interacts with the previous and the next presynaptic spikes (Sjöström et al. 2001; Meffin et al. 2006).

In order to include those different schemes in a Volterra-like framework, we have to change the domain of integration of the F_k and G_k kernels. Instead of integrating from 0 to ∞ as done in Eqs. (1.2.3) and (1.2.4), it is possible to restrict the integration to the last pre- or postsynaptic spike. For a general expression of the *Nearest-Past-Spike* interaction scheme, see appendix A.

In summary, the high level of noise observed in the plasticity data lead us to phenomenological models which intend to be as simple as possible. Hence the above framework makes sense only if the terms above second order or third order can be neglected. Even if we can neglect terms that are above third order, we are left with a huge number of possible models. Indeed, those models combine some or all terms up to the third order with the different

³Note that in those articles the *Nearest-Past-Spike* interaction scheme is called *Nearest-Spike*.

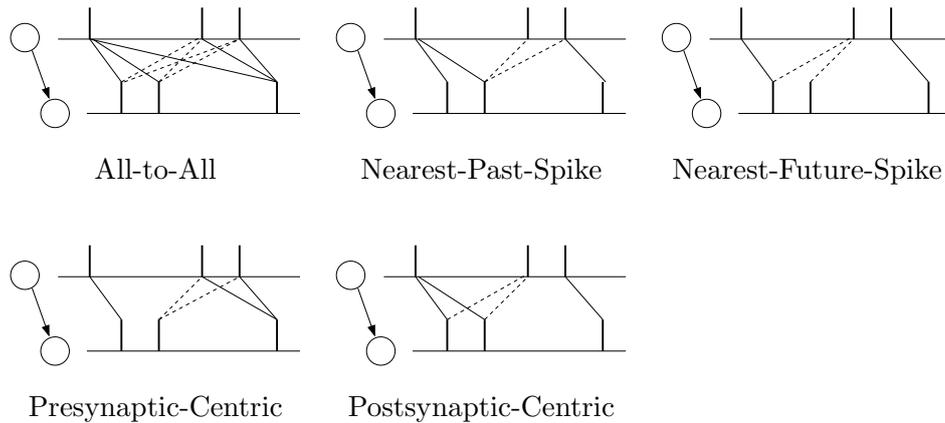


Figure 1.3: Different types of interaction schemes for pairs of pre- and postsynaptic spikes. Solid lines denote pre-post interactions and dashed lines post-pre interactions. **All-to-All**: each presynaptic spike interacts with *all* postsynaptic spikes and vice versa. **Nearest-Past-Spike**: each presynaptic (postsynaptic) spike interacts with the *last* postsynaptic (presynaptic) spike. **Nearest-Future-Spike**: each presynaptic spike interacts with the *next* postsynaptic spike and vice versa. **Presynaptic-Centric**: each presynaptic spike interacts with the *last* and the *next* postsynaptic spike. **Postsynaptic-Centric**: each postsynaptic spike interacts with the *last* and the *next* postsynaptic spike.

interaction schemes mentioned above. At this point, there are two strategies: either a brute-force strategy which tries to test all different possibilities or a different strategy which uses some insight from other approaches, e.g. from biophysical models, in order to choose which combination of terms and interaction scheme to test. This brings us to the next type of models.

1.2.2 Biophysical Models

As soon as we consider STDP experiments, we quickly realize the enormous number of possible protocols. This problem becomes even more serious when we consider the required time to get a single experimental data point ⁴. Even though great efforts have been done to cover a large number of configurations of spike (Froemke et al. 2006), it is simply impossible to cover all the possibilities for all different types of synapses. A different strategy would be to try to understand the mechanisms which lead to the up-regulation or down-regulation of the synapses. This is the approach of the biophysical models.

A large portion of those biophysical models (Karmarkar and Buonomano 2002; Shouval et al. 2002) are based on the dynamics of calcium concentration $[Ca^{2+}]$ which is known to play a crucial role in the induction of LTP or LTD (Malenka et al. 1988; Lisman 1989). Those calcium models follow what has been termed the “calcium control hypothesis” (Shouval et al. 2002). This hypothesis states that the weight changes such that it reaches asymptotically a function $\Omega([Ca^{2+}])$ with a time constant $\tau([Ca^{2+}])$ (see Fig. 1.4) which can be formulated as

$$\dot{w} = \frac{\Omega([Ca^{2+}]) - w}{\tau([Ca^{2+}])} \quad (1.2.5)$$

In those models, the *N*-methyl-*D*-aspartate (NMDA) receptors play a key role. Indeed, the NMDA channels let calcium ions enter the postsynaptic cell if two conditions are met. First, there should be a presynaptic release of glutamate which binds the NMDA receptor and second, the magnesium ion which blocks the NMDA channel should be removed by a depolarization of the postsynaptic cell. This can occur when a postsynaptic spike is back-propagated to the synapse. Therefore, the NMDA receptor acts as a coincidence detector

⁴According to a discussion I had with R. Froemke at Salt-Lake City (Cosyne 2006), when the whole set-up is ready, it takes at least three months full-time to get a reasonable STDP learning window similar to the one depicted in Fig. 1.1.

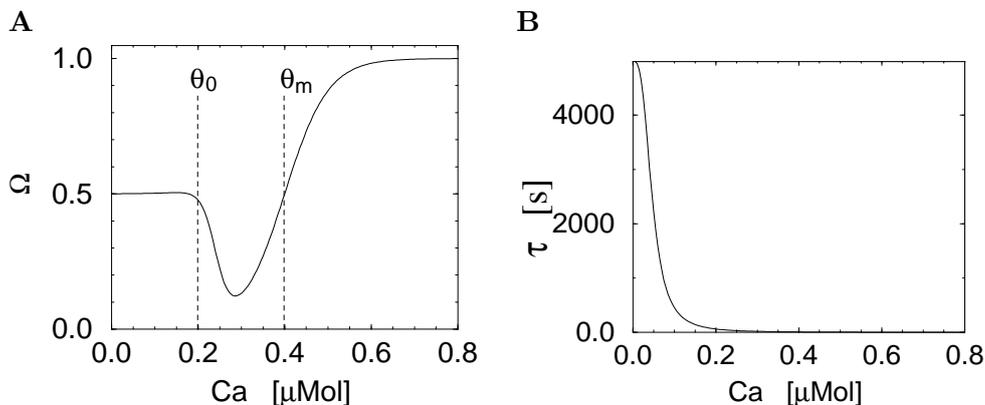


Figure 1.4: Calcium control hypothesis. **A.** Asymptotic value of the weight as a function of calcium concentration. **B.** Time constant of weight change as a function of calcium concentration.

between a pre- and postsynaptic event. This is in direct agreement with the notion of correlation-based learning formulated by Hebb.

SMT Model

Instead of giving some more details about calcium models, let us report another biophysical model, i.e. the model of Senn, Markram, and Tsodyks (2001) (SMT) on Spike-Timing-Dependent Plasticity since it reproduces a lot of experimental data and can be related to this present work, especially paper VI and VII.

This model is based on the dynamics of the NMDA receptors. Those receptors are assumed to be in three different states: the *recovered* state, the *up* state which correspond to the saturation of glutamate and the *down* state corresponding to the state altered by intracellular calcium (Mayer et al. 1987). Let N^u and N^d be the fraction of NMDA receptors in the *up* and *down* state respectively. $N^r = 1 - N^u - N^d$ denotes the fraction of receptors in the *recovered* state; see Fig. 1.5.

Let us further denote the pre- and postsynaptic spike trains by $X(t) = \sum_{t^{pre}} \delta(t - t^{pre})$ and $Y(t) = \sum_{t^{post}} \delta(t - t^{post})$, $\forall t \geq 0$. For the sake of convenience, we assume that $X(t) = Y(t) = 0$, $\forall t < 0$. The dynamics of those receptors becomes:

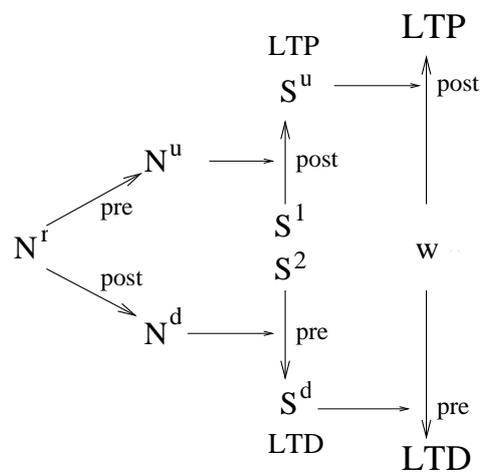


Figure 1.5: Schematic description of the SMT model. In order to elicit LTP, the model requires a first presynaptic spike (pre) that activates a fraction of NMDA receptors, a subsequent postsynaptic spike (post) in order to set some secondary messengers in an up-state. The second postsynaptic spike further induces LTP. Similarly, LTD is induced by a post-pre as well as a post-pre-pre sequence of spikes. See appendix B.2 for more details.

$$\dot{N}^u(t) = -\frac{N^u(t)}{\tau_{N^u}} + r^u N^r(t)X(t) \quad (1.2.6)$$

$$\dot{N}^d(t) = -\frac{N^d(t)}{\tau_{N^d(t)}} + r^d N^r(t)Y(t) \quad (1.2.7)$$

where r^u denotes the fraction of *recovered* NMDA receptors that enter in an *up* state for each presynaptic spike. Without presynaptic spikes, the number of NMDA receptors in an *up* state decreases with a time constant of τ_{N^u} . Similarly, for each postsynaptic spike the fraction of receptors that follow the transition $N^r \rightarrow N^d$ is given by r^d and the decay time constant of the NMDA receptors in a *down* state in the absence of postsynaptic spikes is τ_{N^d} .

The SMT model further assumes the existence of so-called secondary messengers governed by the following kinetic:

$$\dot{S}^u(t) = -\frac{S^u(t)}{\tau_{S^u}} + r^S N^u(t)(1 - S^u(t - \epsilon))Y(t) \quad (1.2.8)$$

$$\dot{S}^d(t) = -\frac{S^d(t)}{\tau_{S^d}} + r^S N^d(t)(1 - S^d(t - \epsilon))X(t) \quad (1.2.9)$$

where S^u and S^d are the secondary messengers in the *up* and *down* state and ϵ is an arbitrary small positive constant. The number of secondary messengers in the *up* state increases if (1) there are some NMDA receptors in the *up* state and if (2) a postsynaptic spike occurs. If at least one of those two conditions are not matched, S^u decreases with a time constant τ_{S^u} . Similarly S^d are activated if presynaptic spikes are elicited when there are NMDA receptors in a *down* state. If this is not the case, S^d decays with a time constant of τ_{S^d} . Finally, the target weight w^∞ dynamic is given by

$$\dot{w}^\infty(t) = r^P(1 - w^\infty(t - \epsilon)) [S^u(t + \epsilon) - \theta^u]^+ Y(t) - r^D w^\infty(t - \epsilon) [S^d(t + \epsilon) - \theta^d]^+ X(t) \quad (1.2.10)$$

Here, the argument $(t + \epsilon)$, $\epsilon > 0$ of the variable S^u and S^d ensures that the value of S^u (S^d) is taken just after its update due to a postsynaptic (presynaptic) spike. $[s]^+$ denotes a piecewise linear function, i.e. $[s]^+ = s$ if $s > 0$ and $[s]^+ = 0$ elsewhere. If the number of secondary messengers are above a given threshold θ^u , potentiation is triggered, with a

rate $r^P(1 - w^\infty)$ whenever there is a pre-post or a pre-post-post sequence (see Fig. 1.5 and appendix B.2 for more details). Analogously, if the number of secondary messengers in a *down* state are above a given threshold θ^d , long-term depression is elicited, with a rate $r^D w^\infty$, when there is a post-pre or a post-pre-pre sequence (see appendix B.2 for more details).

In order to take into account the delay between the time of induction of LTP or LTD and the time of expression, this model low-pass filters the target weight w^∞ with a time constant τ_M and gives the effective weight dynamics:

$$\dot{w}(t) = \frac{w^\infty(t) - w(t)}{\tau_M} \quad (1.2.11)$$

It should be noted that this biophysical model of synaptic plasticity is somehow at the boundary between biophysical and phenomenological models. Indeed, all the parameters of this model are chosen to fit the experimental data and not taken from known underlying biophysical dynamics. Furthermore the nature of the *secondary messengers* and not well identified.

This shows well the dilemma faced by all biophysical models. Either they aim to be rather simple, like the SMT model, but they fail to identify the variables of the model with biophysical quantities, or they take seriously the complex mechanisms involved in the up- or down-regulation of synaptic strength (see for example (Rubin et al. 2005)). In this latter case, it is necessary to consider that synaptic plasticity is cell-specific (Tzounopoulos et al. 2004), depends on the dendritic location (Froemke et al. 2005), the structure of dendritic spines (Kasai et al. 2003), the activation of neighboring synapses (heterosynaptic plasticity) (Bi 2002), the level of calcium concentration (Malenka et al. 1988), the level depolarization (Sjöström et al. 2001). It is also clear that a highly complex genetic machinery (Kandel 2001) is involved in the regulation of synaptic strength. It is even possible that astrocytes may play an active role since they are regulating synaptic transmission (Volterra and Meldolesi 2005).

When facing this gigantic level of complexity, it seems reasonable to approach the problem of synaptic plasticity from a different point of view. Instead of getting lost in a mass of (possibly irrelevant) details, we need criteria or guiding principles that allow us to focus on relevant mechanisms of synaptic plasticity. In order to do so, we need to have a picture of the functional role of the synapse. This is the approach of the optimality models developed in the next section.

1.2.3 Optimality Models

A wing would be a most mystifying structure if one did not know that birds flew.

Barlow (1961)

It is with those words that H. B. Barlow starts his influential 1961 paper. His point of view is radically different from the biophysical approach. He is more interested in the functional role of the synapse than in the mechanism that lead to the up- or down regulation of synaptic strength. According to him this approach is necessary “if one does not will get lost in a mass of irrelevant detail and fail to make the crucial observations” (Barlow 1961). Among the three hypothesis he develops in his article in order to explain the functional role of synapses, namely the *password hypothesis*, the *controlled pass-characteristic hypothesis* and the *redundancy-reducing hypothesis*, he focuses his interest on the latter. The idea is that “sensory relays” or synapses at a high level recode sensory messages so that the amount of information transmission is maximal or the redundancy, formulated as a function of the input entropy (Shannon 1948), is reduced.

At that time, this idea of redundancy reduction was not new. Indeed, few years earlier, Attneave (1954) stated that “a major function of the perceptual machinery is to strip away some of the redundancy of stimulation, to describe or encode information in a form more economical than that in which it impinges on the receptors”. In fact, the idea of “economy of thoughts” could be even traced back to Mach (1886) and Pearson (1892).

Later on, Laughlin (1981) gave the first experimental evidence that information theory can be used to understand neural coding strategies. More precisely, Laughlin showed that the contrast-response function of large monopolar cells (LMC’s) of fly visual system follows the cumulative probability function for natural contrasts. In this way, the coding strategy maximizes the information capacity of the LMC cells by setting the maximal slope and the inflexion point of their contrast-function so that all levels of responses are used with equal frequency.

Consistently with the idea of redundancy reduction, Linsker (1988) proposed a self-organization principle of the sensory networks which is now called the *infomax* principle:

Given a layer L of cells, and the stationary ensemble statistical properties of the signal activity values in the layer and given that layer L is to provide input

to another cell layer M , the transformation of activity values from L to M is to be chosen such that the rate R of information transmission from L to M is maximized, subject to constraints and/or additional cost terms.

This information maximization principle, which can be mapped to Barlow's redundancy reduction hypothesis (Nadal et al. 1998), has been shown to have several computational properties such as blind source separation (Nadal and Parga 1994; Bell and Sejnowski 1995). The papers III and V in this thesis, show that this infomax principle leads as well to some realistic biological properties such as STDP.

When developing such kind of models, it is important to keep in mind that those models should help us to understand the functional role of the synapse, but the optimality approach does not imply that all experiments can be reproduced by such models. As Barlow himself made the point clear "I do not regard these ideas as moulds into which all experimental facts must be forced".

BCM Model

Another interesting consequence of the infomax principle is that it can be mapped to the Bienenstock-Cooper-Munro (BCM) learning rule (Bienenstock et al. 1982) (see (Nadal and Parga 1997) as well as paper IV). The aim of the BCM learning rule is to maximize the selectivity of a given neuron. Let $\mathcal{R} = \{\rho^1, \dots, \rho^N\} \in \mathbb{R}^M$ denote a set of N linearly independent inputs presented with probability p_j , $j = 1, \dots, N$ to a linear postsynaptic neuron, i.e. the response of this neuron is given by $y = w^T \rho$ where $w \in \mathbb{R}^M$ is the weight vector. The postsynaptic neuron is called selective to the input pattern $\rho^k \in \mathcal{R}$ if its response is maximal when ρ^j is presented and is zero whenever any other pattern $\rho^i \in \mathcal{R}$, $i \neq j$ is presented. Formally, the selectivity of a neuron with synapses characterized by w can be defined as

$$\text{Sel}(w) = 1 - \frac{\langle w\rho \rangle_{\rho \in \mathcal{R}}}{\max_{\rho \in \mathcal{R}}(w\rho)} \quad (1.2.12)$$

where $\langle w\rho \rangle_{\rho \in \mathcal{R}}$ denotes the average response of the neuron over all inputs $\rho \in \mathcal{R}$. Therefore, if the postsynaptic neuron gives the same response for all inputs, the average response

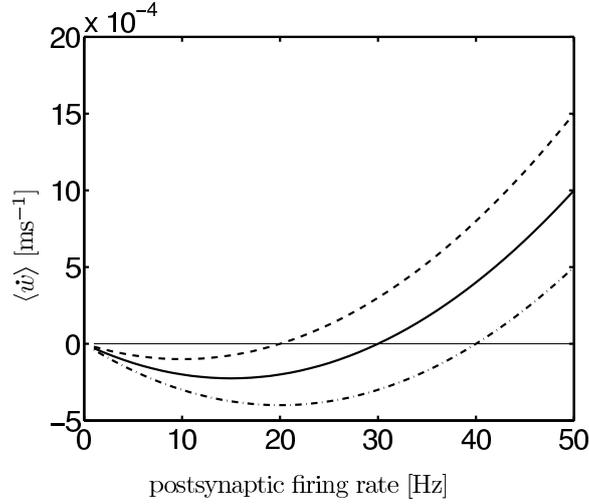


Figure 1.6: Weight change given by the BCM learning rule. Here the function ϕ is given by $\phi(y, \theta) = -y\theta + y^2$ and $\theta = 20$ Hz (dashed line), $\theta = 30$ Hz (solid line) and $\theta = 40$ Hz (dot-dashed line)

will be equal to the maximal response, therefore the selectivity will be $\text{Sel}(w) = 0$. Conversely, if the postsynaptic neuron's firing rate is minimal for all inputs except one to which it responds maximally, the neuron is maximally selective. One of the main results of Bienenstock, Cooper, and Munro (1982) is that the following learning rule maximizes the input selectivity:

$$\dot{w} = \eta \phi(w\rho^j, \theta) \rho^j \quad (1.2.13)$$

Here η is the learning rate and $\phi(y, \theta)$ is a function of the postsynaptic cell activity y (see Fig. 1.6) such that

$$\phi(0, \theta) = 0 \quad (1.2.14)$$

$$\text{sgn}(\phi(y, \theta)) = \text{sgn}(y - \theta) \quad (1.2.15)$$

θ is the threshold between depression and potentiation. If the threshold θ is fixed, it corresponds to an unstable fixed point of the dynamic, and therefore the weights could possibly

all decrease to zero or increase up to infinity. In order to prevent such an uninteresting case, the threshold is sliding according to the following non-linear dynamic⁵

$$\theta = \alpha \langle y \rangle^p \quad (1.2.16)$$

where $\alpha > 0$ and $p > 1$ and $\text{sgn}(s)$ is the sign function, i.e. $\text{sgn}(s) = -1$ if $s < 0$, $\text{sgn}(s) = 1$, if $s > 0$ and $\text{sgn}(0) = 0$. With this learning rule it is possible to show that the weights w^* such that the selectivity is maximal correspond to a stable fixed point of the dynamic given by Eq. (1.2.13). Therefore, in an environment of N input patterns, each postsynaptic neuron specializes automatically to a given input. The nature of the input to which it specializes depends on the initial condition. Because of this property, the BCM learning rule has been used to show the emergence of orientation selectivity and the binocular interaction in the visual cortex (Hubel and Wiesel 1959; Hubel and Wiesel 1962).

1.3 Aims and achievements

General aim

The aim of this thesis is to develop new models of Spike-Timing-Dependent Plasticity in order to better understand the functional role and the phenomenology of the synapse.

Specific aims

1. To choose a suitable probabilistic description of a neuron which can be used to define optimality criteria.
2. To develop optimality criteria in the context of supervised learning that lead to Spike-Timing-Dependent Plasticity.
3. To develop models of synaptic plasticity, in the context of unsupervised learning, that result from the maximization of mutual information with or without constraints.

⁵Note that another version of the sliding threshold proposed by Intrator and Cooper (1992) is given by $\theta = \alpha \langle y^p \rangle$. In this way, the result is slightly more general and requires less assumptions.

4. To show possible links between Spike-Timing-Dependent Plasticity and the BCM learning rule.
5. To develop a phenomenological model of STDP which can reproduce as many experiments as possible.

Achievements

This thesis is based on the following articles, which will be referred to by their Roman numerals in the text:

- I. **Pfister JP and Barber D and Gerstner W.** Optimal Hebbian Learning: A Probabilistic Point of View. *Artificial Neural Networks and Neural Information Processing - ICANN/ICONIP 2003*, edited by O. Kaynak, E. Alpaydin, E. Oja and L. Xu. Berlin: Springer-Verlag, 92-98, 2003.
- II. **Pfister JP and Toyozumi T and Aihara K and Gerstner W.** Optimal Spike-Timing Dependent Plasticity for Precise Action Potential Firing in Supervised Learning. *Neural Computation*, 18, 1309-1339, 2006.
- III. **Toyozumi T and Pfister JP and Aihara K and Gerstner W.** Spike-Timing Dependent Plasticity and Mutual Information Maximization for a Spiking Neuron Model. *Advances in Neural Information Processing Systems 17*, edited by L.K. Saul and Y. Weiss and L. Bottou, MIT Press, Cambridge MA, 1409-1416, 2005.
- IV. **Toyozumi T and Pfister JP and Aihara K and Gerstner W.** Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission. *Proceedings of the National Academy of Science USA*. 102, 5239-5244, 2005.
- V. **Toyozumi T and Pfister JP and Aihara K and Gerstner W.** Optimality Model of Unsupervised Spike-Timing Dependent Plasticity: synaptic memory and weight distribution. *Accepted to Neural Computation*.
- VI. **Pfister JP and Gerstner W.** Beyond Pair-Based STDP: a Phenomenological Rule for Spike Triplet and Frequency Effects. *Advances in Neural Information Processing*

Systems 18, edited by Y. Weiss and B. Schölkopf and J. Platt, MIT Press, Cambridge MA, 1083-1090, 2006.

VII. **Pfister JP and Gerstner W.** Why Triplets of Spike are Necessary in Models of Spike-Timing-Dependent Plasticity. *Submitted to The Journal of Neuroscience.*

More specifically, I deals with points 1,2 of the *specific aims* list. II deals with point 2. III, IV and V develop models of synaptic plasticity that follow point 3. IV, VI and VII address the question 4. VI and VII give an example of approach for point 5.

In the next chapter, all those seven papers are presented with an introduction. The order of the papers follows less the chronology of publication than the chronology of submission. In this way, papers I to V deal with the first point of the *general aim*, i.e. develop optimality models, and the papers VI and VII develop phenomenological models which is the second point of the *general aim*.

This chapter presents a compilation of all the publications that came out as a product of this thesis. Each paper is preceded by a short introductory summary that links it to the rest of the work and reference details.

2.1 Paper I

Summary

The first five papers of this thesis (I-V) develop synaptic learning rules that can be classified as *optimal models* (see section 1.2.3 of the introduction). Papers I and II deal with supervised learning while papers III-V consider the case of unsupervised learning.

Paper I is a first attempt to define an objective function that synapses should maximize. The objective function is here the logarithm of the probability of observing a single post-synaptic spike at time t^{post} (and no other spikes from 0 to T) when a presynaptic spike is elicited at time t^{pre} . This gives

$$L = \log P(t^{post}|t^{pre}) \tag{2.1.1}$$

The resulting weight change $\Delta w = \partial L / \partial w$ depends on the time difference between the pre- and postsynaptic spike and can be expressed as a learning window. With this objective

function, potentiation occurs if the presynaptic spike is just before the postsynaptic one and depression otherwise. The interesting point about the potentiation part is that it follows the shape of an excitatory postsynaptic potential (EPSP). Therefore, in this context, the potentiation time constant of the optimal learning window is given by the membrane time constant.

The acausal part of the learning window, i.e. when the presynaptic spike is after the postsynaptic one, is more problematic. Indeed, the strong asymmetry in the learning window occurs only in the presence of a depolarizing afterpotential which does not seem realistic. Another problem of this approach is the negative offset of the learning which is not present in the experimental results. Those open problems were the motivations for the paper II.

The interest of this paper is more in the developed framework than in the results themselves. Indeed, in order to have an objective function which depends explicitly on the synaptic strength w , we used (1) the Spike-Response Model (SRM) which formulates the membrane potential as an explicit function of the weight and (2) the escape noise theory which allows us to calculate explicitly the probability of emitting a spike (and therefore the probability of generating a whole spike train) for a given membrane potential. This combination of SRM and escape noise has been used in the papers II-V as well as in other papers such as (Bohte and Mozer 2005).

Reference

The content of this paper was presented at the *EPFL Latsis Symposium* (February 17-22, 2003, Lausanne, Switzerland) and more formally at the *Joint International Conference ICANN/ICONIP 2003* (June 26 - 29, 2003, Istanbul, Turkey). This paper was published after peer-review process in the proceedings of the *ICANN/ICONIP* conference under the following reference:

Pfister JP and Barber D and Gerstner W. Optimal Hebbian Learning: A Probabilistic Point of View. *Artificial Neural Networks and Neural Information Processing - ICANN/ICONIP 2003*, edited by O. Kaynak, E. Alpaydin, E. Oja and L. Xu. Berlin: Springer-Verlag, 92-98, 2003.

This paper has been cited by (Suri 2004).

Optimal Hebbian Learning: a Probabilistic Point of View

Jean-Pascal Pfister¹, David Barber², Wulfram Gerstner¹

¹ Laboratory of Computational Neuroscience, EPFL
CH-1005 Lausanne, Switzerland

{jean-pascal.pfister, wulfram.gerstner}@epfl.ch

² Institute for Adaptive and Neural Computation, Edinburgh University, 5 Forrest
Hill, Edinburgh, EH1 2QL, U.K.
d.barber@anc.ed.ac.uk

Preprint version of
Artificial Neural Networks and Neural Information Processing
ICANN/ICONIP, 92-98 (2003)

Abstract

Many activity dependent learning rules have been proposed in order to model long-term potentiation (LTP). Our aim is to derive a spike time dependent learning rule from a probabilistic optimality criterion. Our approach allows us to obtain quantitative results in terms of a learning window. This is done by maximising a given likelihood function with respect to the synaptic weights. The resulting weight adaptation is compared with experimental results.

1 Introduction

Since synaptic changes are most likely to underly memory and learning processes, it is crucial to determine the causes and underlying laws describing this adaptation process. Among the enormous number of models, there are mainly two categories: rate-based and spike-based learning rule. In this paper, we want to present a new way to derive a spike-time dependent learning rule. Existing models of spike-timing dependent plasticity are either phenomenological [12] or, in contrast, mechanistic [1].

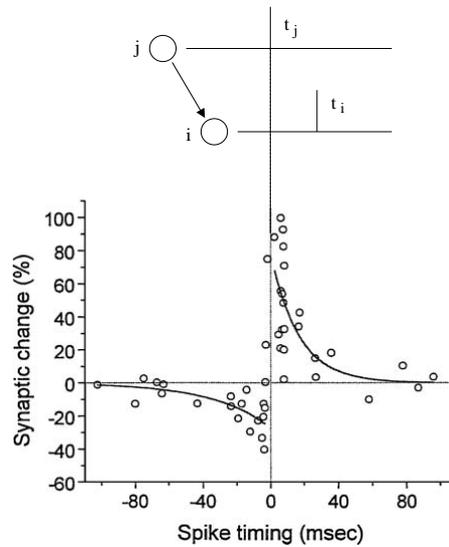


Figure 1: Critical window for synaptic modifications. Long-term potentiation (LTP)/long-term depression (LTD) were induced by correlated pre- and postsynaptic spiking at synapses between hippocampal glutamatergic neurons in culture. Figure adapted from [4].

Our model is derived from a probabilistic point of view in the sense that the learning rule should optimise the likelihood of observing a postsynaptic spike train with a desired timing, given the postsynaptic membrane potential at the location of the synapse.

A significant part of the synaptic plasticity models are based on Hebb's postulate [14]:

When an axon of cell A is near enough to excite cell B or repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.

In fact models rephrase this postulate by saying that the adaptation of the synaptic weights is driven by a simultaneous activity of the pre- and the postsynaptic neuron. This simultaneity has to be defined in a time window. Recent experiments [3] have shown the influence of a pair of a single pre- and postsynaptic spike on the synaptic strength (figure 1).

The aim of this paper is to show that it is possible to get a similar learning window as a result from an optimal learning rule. Recently Barber [2] studied this question with neurons discrete in time. Here we want to extend this study to the continuous case and discuss the results in relation with the experiments of Bi and Poo [3].

2 Spike Response Model

For the sake of simplicity, let us consider here a single presynaptic neuron j and a postsynaptic neuron i . Those two neurons are considered as Poisson neurons, i.e. their firing times depend only on the present value of the membrane potential. Let w be the synaptic weight between those neurons. Finally, let $\{t_j^{f'}\}$ and $\{t_i^f\}$ denote respectively the pre- and postsynaptic firing times.

The fundamental hypothesis in this article is to assume that the instantaneous firing rate of the postsynaptic neuron is given by an increasing function of the membrane potential $u(t)$:

$$\rho(t) = g(u(t)). \quad (1)$$

This firing rate can be also termed escape rate [11]. The membrane potential model we take is the Spike Response Model (SRM). The simplest SRM is called SRM₀ [9] and defines the membrane potential $u(t)$ as follow:

$$u(t) = u_{\text{rest}} + \eta(t - \hat{t}_i) + w \sum_{f'} \epsilon(t - t_j^{f'}), \quad (2)$$

where, in our case, $\eta(s)$ is a kernel describing the spike-afterpotential, $\epsilon(s)$ is the kernel representing the excitatory post-synaptic potential (EPSP) and \hat{t}_i is the last firing time of neuron i , i.e. $\hat{t}_i = \max\{t_i^f | t_i^f < t\}$. The goal is now to maximise the probability that the postsynaptic spike train $S_i(t) = \sum_{f'} \delta(t - t_i^{f'})$ has been generated by the firing rate $\rho(t)$.

3 Calculation of the likelihood \mathcal{L}

In order to calculate the likelihood of a spike train given a firing rate, it is useful to first make a time discretization before coming back to the continuous case. Let $\bar{\rho}(t)$ be the discretised version of $\rho(t)$ on the interval $I = [0, T]$ where $\bar{\rho}(t) = \rho(t_n)$, $\forall t \in [t_n, t_n + \Delta t]$ and $t_0 = 0$, $t_N = N\Delta t = T$.

The probability that a neuron produces a spike at time $t \in [\tilde{t}, \tilde{t} + \Delta t]$ given its firing rate $\bar{\rho}(s)$ is simply given by the probability of spiking between \tilde{t} and $\tilde{t} + \Delta t$ multiplied by the probability of not spiking at any other time:

$$\bar{P}(t \in [\tilde{t}, \tilde{t} + \Delta t] | \bar{\rho}(s)) \Delta t = \bar{\rho}(\tilde{t}) \Delta t \prod_{t_n \neq \tilde{t}} (1 - \bar{\rho}(t_n) \Delta t) \quad (3)$$

To extend this result to the case of M spikes, we need to define $\mathbf{t} = (t^1, \dots, t^M)$ a M -dimensional time variable ordered chronologically, i.e. $t^f < t^{f+1}$. Let $\tilde{\mathbf{t}}$ be the M desired firing times and $\Omega(\tilde{\mathbf{t}}) = \prod_n [\tilde{t}_n, \tilde{t}_n + \Delta t]$ a M -dimensional bin. The probability of firing at the M given times \mathbf{t} is

$$\bar{P}(\mathbf{t} \in \Omega(\tilde{\mathbf{t}}) | \bar{\rho}(s)) \Delta t^M = \prod_f \bar{\rho}(\tilde{t}^f) \Delta t^M \prod_{t_n \neq \tilde{t}^f} (1 - \bar{\rho}(t_n) \Delta t)$$

$$\begin{aligned}
&= \prod_f \frac{\bar{\rho}(\tilde{t}^f) \Delta t^M}{1 - \bar{\rho}(\tilde{t}^f) \Delta t} \prod_n (1 - \bar{\rho}(t_n) \Delta t) \\
&= \prod_f \frac{\bar{\rho}(\tilde{t}^f) \Delta t^M}{1 - \bar{\rho}(\tilde{t}^f) \Delta t} \exp \left(\sum_n \log(1 - \bar{\rho}(t_n) \Delta t) \right). \quad (4)
\end{aligned}$$

Now we can come back to the continuous case. By taking the limit $\Delta t \rightarrow 0$, we have $\bar{\rho}(t) \rightarrow \rho(t)$, $\bar{P}(\mathbf{t} \in \Omega(\tilde{t}) | \bar{\rho}(s)) \rightarrow P(\mathbf{t} = \mathbf{t}' | \rho(s))$, $\sum_n \log(1 - \bar{\rho}(t_n) \Delta t) \rightarrow -\int_0^T \rho(t) dt$ and $1 - \bar{\rho}(\tilde{t}^f) \Delta t \rightarrow 1$. Therefore we can define the log-likelihood $\mathcal{L}(\mathbf{t}_i | u(s))$ of the postsynaptic spike train given the membrane potential $u(s)$ by simply taking the logarithm of $P(\mathbf{t} = \mathbf{t}' | \rho(s))$:

$$\mathcal{L}(\mathbf{t}_i | u(s)) = \sum_f \log(g(u(t_i^f))) - \int_0^T g(u(t)) dt. \quad (5)$$

4 Learning rule

The goal of our study is to find a learning rule which tends to optimise the weight w in order to maximise the likelihood of getting postsynaptic firing times given the firing rate. This means that those weights must evolve in the direction of the gradient of \mathcal{L} :

$$w^{\text{new}} = w + \kappa \frac{\partial \mathcal{L}}{\partial w}, \quad (6)$$

with

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial w}(\mathbf{t}_i | u(s)) &= \sum_f \frac{\frac{dg(u(t_i^f))}{du(t_i^f)} \frac{\partial u(t_i^f)}{\partial w}}{g(u(t_i^f))} - \int_0^T \frac{dg(u(t))}{du(t)} \frac{\partial u(t)}{\partial w} dt \\
&= \sum_f \sum_{f'} \frac{\epsilon(t_i^f - t_j^{f'})}{g(u(t_i^f))} \frac{dg(u(t_i^f))}{du(t_i^f)} - \int_0^T \frac{dg(u(t))}{du(t)} \sum_{f'} \epsilon(t - t_j^{f'}) dt \quad (7)
\end{aligned}$$

and κ is the learning rate. Since $g(u(t)) = \exp(\beta(u(t) - \theta))$ is a reasonable choice [11], we can use it to evaluate the gradient of \mathcal{L} for a pre- and a postsynaptic spike train:

$$\frac{\partial \mathcal{L}}{\partial w}(\mathbf{t}_i | u(s)) = \beta \sum_f \sum_{f'} \epsilon(t_i^f - t_j^{f'}) - \beta \int_0^T \exp(\beta(u(t) - \theta)) \sum_{f'} \epsilon(t - t_j^{f'}) dt. \quad (8)$$

Let us now study the restricted case with only one pre- and one postsynaptic spike and $\beta = 1$:

$$\frac{\partial \mathcal{L}}{\partial w}(t_i | u(s)) = \epsilon(t_i - t_j) - \int_0^T \exp(u(t) - \theta) \epsilon(t - t_j) dt. \quad (9)$$

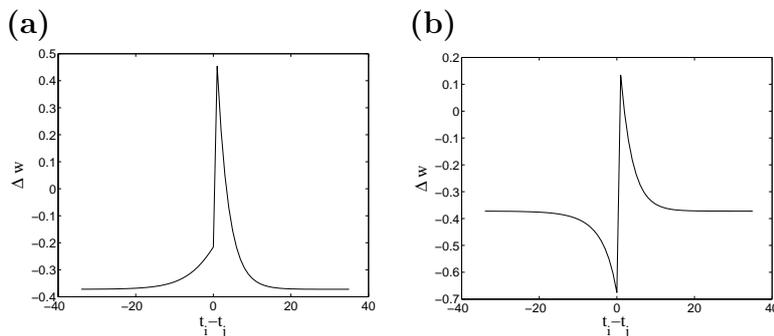


Figure 2: First step of the adaptation of weights $\Delta w = w^{\text{new}} - w = \frac{\partial \mathcal{L}}{\partial w}$. The parameters for this simulation are : $w = 0.2, \theta - u_{\text{rest}} = -2, \beta = 1, \epsilon_0 = 1, \tau_\epsilon = 3, \tau_\eta = 5$. The amplitude of the spike-afterpotential is given by $\eta_0 = -1$ (HAP) for (a) and $\eta_0 = 1$ (DAP) for (b). Note the different vertical scales.

In order to represent the gradient of the log-likelihood function \mathcal{L} (figure 2), it is necessary to choose determine specific kernels for $\eta(s)$ and $\epsilon(s)$. For simplicity sake, we take

$$\eta(s) = \eta_0 e^{-\frac{s}{\tau_\eta}} \Theta(s), \quad (10)$$

$$\epsilon(s) = \epsilon_0 e^{-\frac{s}{\tau_\epsilon}} \Theta(s), \quad (11)$$

where Θ is the usual Heaviside step function with $\Theta(s) = 1$ for $s > 0$ and $\Theta(s) = 0$ else. If $\eta_0 > 0$, the neuron exhibits a depolarizing afterpotential (DAP). In reverse, if $\eta_0 < 0$, is exhibits a hyperpolarizing afterpotential (HAP).

It is interesting to note that the qualitative shape of this learning window is similar to the one obtained by Bi and Poo [3] only in presence of DAP which could be consistent with DAP observed by Connors et al. in neocortical neurons [6].

5 Discussion

One can note that a major difference between the result of Bi and Poo and our model is the negative offset. This offset is related to the integral of the kernel $\epsilon(s)$. Indeed, if the postsynaptic spike occurs a long time after the presynaptic spike and if $w \simeq 0$, the first term of equation (9) can be neglected and the membrane potential can be approximated by its resting potential in the range where $\epsilon(t - t_j)$ is significant:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial w} &\simeq -\exp(u_{\text{rest}} - \theta) \int_{t_j}^T \epsilon(t - t_j) dt \\ &\simeq -\exp(u_{\text{rest}} - \theta) \int_0^\infty \epsilon(s) ds. \end{aligned} \quad (12)$$

This is of course valid only if $T \gg t_j$. In fact, this offset is related to the probability of not having a spike at time $t \neq t_i$ (c.f. last term of eq. (3)). In order to increase the likelihood of not having a spike the weight needs to be reduced.

A possible way to solve the problem of negative bias of Δw is to consider a slightly different scenario where we still impose the postsynaptic spike at a given time, but instead of imposing no spike at all the other times, we choose the desired number of spikes that are allowed to occur stochastically over the period T . This number of spikes has to be related to the spontaneous firing rate.

One can also note that the shape of the positive peak on figure 2a is dominated by the kernel $\epsilon(s)$ (c.f. first term of eq. (9)). This is due to the choice of an exponential for the function $g(u)$.

Let us note that we are looking at the gradient of the likelihood function \mathcal{L} and not at the optimal solution given by $\frac{\partial \mathcal{L}}{\partial w} = 0$. Indeed, it is straightforward to notice that there is no fixed point for w if $t_i < t_j$. For $t_j < t_i$, there is a fixed point and it is stable since $\frac{\partial^2 \mathcal{L}}{\partial w^2} < 0$.

We have shown a new framework for deriving a spike-time dependent learning rule. The interesting feature of this learning rule is the similarity to the one obtained by Bi and Poo. This similarity is valid only in presence of DAP. The duration of the DAP determines the width of the negative phase of the learning window.

As a consequence we could speculate that the form of the learning window changes according to the type of neuron since in this framework the learning window strongly depends on the spike-afterpotential.

It is of course possible to make the model more complex by not using the SRM₀ model but more realistic models. Even if our study was restricted to a single pre- and postsynaptic spike, equation (8) remains totally general for spike trains and is also valid for an entire neural network.

References

- [1] H.D.I. Abarbanel, R. Huerta and M.I. Rabinovich: Dynamical model of long-term synaptic plasticity. Proc. Natl. Acad. Sci. USA, 2002, Vol. 99, Issue 15, 10132-10137.
- [2] D. Barber: Learning in Spiking Neural Assemblies. To appear in proceedings of NIPS 2002.
- [3] G.Q. Bi and M.M. Poo: Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J. Neurosci., 1998, Vol. 18, 10464-10472.
- [4] G.Q. Bi and M.M. Poo: Synaptic modifications by correlated activity: Hebb's postulate revisited. Annu. Rev. Neurosci., 2001, Vol. 24, 139-166.

- [5] G. Bugmann and C. Christodoulou and J.G. Taylor: Role of temporal integration and fluctuation detection in the highly irregular firing of leaky integrator neuron model with partial reset. *Neural Computation*, 1997, Vol. 9, 985-1000.
- [6] B.W. Connors, M.J. Gutnick and D.A. Prince: Electrophysiological Properties of Neocortical Neurons in Vitro. *J. Neurophysiol.* 1982, Vol. 48, 1302-1320.
- [7] D. Debanne and B.H. Gähwiler and S.M. Thompson: Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol.*, 1998, Vol. 507, 237-247.
- [8] R. Froemke and Y. Dan: Spike-timing dependent plasticity induced by natural spike trains. *Nature*, 2002, Vol. 416, 433-438.
- [9] W. Gerstner: Time Structure of the Activity in Neural Network Models. *Phys. Rev. E*, 1995, 51, Vol. 1, 738-758.
- [10] W. Gerstner, R. Kempter, J.L. van Hemmen and H. Wagner: A neuronal learning rule for sub-millisecond temporal coding, *Nature*, 1996, Vol. 383, 76-78.
- [11] W. Gerstner and W.M. Kistler: *Spiking Neuron Models*. Cambridge University Press, 2002.
- [12] W. Gerstner and W.M. Kistler: Mathematical Formulations of Hebbian Learning. *Biological Cybernetics*, 2002, 87, 404-415.
- [13] R. Gütiğ and R. Aharonov and S. Rotter and H. Sompolinsky: Learning input correlations through non-linear temporally asymmetry Hebbian plasticity. To appear.
- [14] D.O. Hebb: *The Organization of Behavior*. Wiley, 1949, New York.
- [15] R. Kempter and W. Gerstner and J. L. van Hemmen: Hebbian learning and spiking neurons. *Phys. Rev. E*, 1999, Vol. 59, 4, 4498-4514.
- [16] W.M. Kistler and J. Leo van Hemmen: Modeling Synaptic Plasticity in Conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput.* 2000, Vol 12, 385-405.
- [17] H. Markram and B. Sakmann: Action potentials propagating back into dendrites trigger changes in efficacy of single-axon synapses between layer V pyramidal neurons. *Soc. Neurosci. Abstr.*”, 1995, Vol. 21, 2007.
- [18] M. Rapp, Y. Yarom and I. Siegev: Modeling back propagating action potential in weakly excitable dendrites of neocortical pyramidal cells. *Proc. Natl. Acad. Sci USA*, 1996, Vol. 93, 11985-11990.
- [19] P.D. Roberts and C.C. Bell: Spike timing dependent synaptic plasticity in biological systems. *Biol. Cybernetics*, 2002, Vol. 87, 392-403.

- [20] W. Senn, H. Markram and M. Tsodyks: An Algorithm for Modifying Neurotransmitter Release Probability Based on Pre- and Postsynaptic Spike Timing. *Neural Comput.* 2000, Vol. 13, 35-67.
- [21] H.Z. Shouval, M.F. Bear and L.N. Cooper: A unified model of NMDA receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA*, 2002, Vol. 99, 10831-10836.
- [22] P.J. Sjöström, G.G. Turrigiano and S.B. Nelson; Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 2001, Vol. 32, 1149-1164.
- [23] S. Song and K.D. Miller and L.F. Abbott: Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience*, 2000, Vol. 3, 919-926
- [24] G.J. Stuart and B. Sakmann: Active propagation of somatic action-potentials into neocortical pyramidal cell dendrites. *Nature*, 1994, Vol. 367, 69-72.
- [25] T.W. Troyer and K.D. Miller: Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation*, 1997, Vol. 9, 971-983.
- [26] M.C.W. van Rossum and G.Q. Bi and G.G. Turrigiano: stable Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience*, 2000, Vol. 20, 8812-8821.

2.2 Paper II

Summary

This second paper follows the same line as the first one in the sense that it looks for an objective function L that can be maximized by the weight changes in the hope that the resulting weight change can be compared to the experimental results. In fact, we propose three different objective functions L^{A_c} , L^{B_c} and L^{C_c} in three different scenarios (A , B and C) and discuss the resulting learning windows. The motivation behind those scenarios is to solve the two main problems raised by paper I, i.e. the negative offset in the learning window and the difficulty to get an acausal LTD part in the learning window. This latter point was achieved in paper I with the use of an unrealistic depolarizing afterpotential (DAP).

Scenario A proposes a solution for the LTD problem, but does not solve the offset problem. The idea is that, since we consider a supervised scenario, the postsynaptic spike has to be triggered by a supervisor (or teacher). This is why we included a teaching signal, which is present in the *in vitro* experiments. By taking a depolarizing current as a teaching signal, there is no need anymore to assume a neuronal dynamic with a depolarizing afterpotential in order to get LTD.

The second scenario (B) gives a solution to remove the negative offset while still having an acausal LTD part. The LTD part is obtained with the same trick as the one used in scenario A , i.e. the teaching signal. The offset problem has been solved in the following way. In the first scenario we imposed a postsynaptic spike at time t^{post} and no one else from 0 to T . Since this constraint of “no other spikes” induced the negative offset, we relaxed this constraint in scenario B and allowed a fixed spontaneous firing rate.

Finally, scenario C , which is a generalization of scenario A , solves the offset and the LTD problem in a different way. Since the outcome of the simplest scenario (A) is a learning window with an EPSP shaped potentiation and a negative offset, the idea is to impose an (artificial) locality constraint which penalizes stronger weight changes when $|t^{pre} - t^{post}|$ is bigger. In this way, the EPSP shaped potentiation remains, the negative offset for t^{pre} just after t^{post} is transformed into LTD and the offset for large $|t^{pre} - t^{post}|$ vanishes because of this locality constraint.

In order to make the link between the learning rules derived in this paper (and also the

one in paper I) which depend on the desired firing time t^{des} (and not the actual one) and the experimentally observed STDP function, it is necessary to interpret correctly those learning rules. A first possibility is to see those rules in a supervised scenario where the desired firing times t^{des} are forced by the experimentalist himself or by a group of neurons and therefore the desired firing time corresponds to the actual one. A second interpretation would place the learning rules of papers I and II in a reinforcement learning context (see section 4.1 in paper II). From this point of view, the weights change only when the reward is non-zero. Therefore, if a given spatio-temporal pattern of input leading to the desired postsynaptic spikes t^{des} triggers some positive reward, the actual firing times t^{post} match the desired one t^{des} . Therefore in both the supervised scenario and the reinforcement learning scenario, the desired firing times can be matched to the actual firing times.

A stable result across the three scenarios A , B and C (and in agreement with paper I) is that the shape of the potentiation part of the learning window remains given by the shape of an EPSP. The situation for the acausal part of the learning window is different. Indeed, the presence of a depression in the post-before-pre region is caused by different mechanisms and therefore does not have a unique interpretation.

Reference

In each of the following conferences, some parts of this paper have been presented: the *CNS Meeting* (July 6-10, 2003, Alicante, Spain), *Les Houches summer school: Methods and Models in Neurophysics* (July 28 - August 29, 2003, Les Houches, France), the *The Monte Verità Workshop on Spike-Timing Dependent Plasticity* (February 29 - March 5, 2004, Ascona, Switzerland), the *Arc Lémanique Neuroscience Retreat* (September 27 - 28, 2004, Les Diablerets, Switzerland), the *CNS Meeting*, (July 17 - 22, 2004, Baltimore, USA) the *Cosyne conference* (March 17 - March 22, 2005, Salt-Lake City, USA).

This paper is published in *Neural Computation* under the following reference:

Pfister JP and Toyozumi T and Aihara K and Gerstner W. Optimal Spike-Timing Dependent Plasticity for Precise Action Potential Firing in Supervised Learning. *Neural Computation*, 18, 1309-1339, 2006.

An earlier version of this manuscript has been published on the arXiv server, the 24th of

February 2005. It can be found at the following address:

<http://arxiv.org/abs/q-bio.NC/0502037>.

Optimal Spike-Timing Dependent Plasticity for Precise Action Potential Firing in Supervised Learning

Jean-Pascal Pfister, Taro Toyozumi*, David Barber[†], Wulfram Gerstner
 Laboratory of Computational Neuroscience,
 School of Computer and Communication Sciences
 and Brain-Mind Institute,
 Ecole Polytechnique Fédérale de Lausanne (EPFL)
 CH-1015 Lausanne
 {jean-pascal.pfister, wulfram.gerstner}@epfl.ch
 taro@sat.t.u-tokyo.ac.jp, david.barber@idiap.ch

Preprint version of:
Neural Computation 18, 1309-1339 (2006)

Abstract

In timing-based neural codes, neurons have to emit action potentials at precise moments in time. We use a supervised learning paradigm to derive a synaptic update rule that optimizes via gradient ascent the likelihood of postsynaptic firing at one or several desired firing times. We find that the optimal strategy of up- and downregulating synaptic efficacies depends on the relative timing between presynaptic spike arrival and desired postsynaptic firing. If the presynaptic spike arrives before the desired postsynaptic spike timing, our optimal learning rule predicts that the synapse should become potentiated. The dependence of the potentiation on spike timing directly reflects the time course of an excitatory postsynaptic potential. However, our approach gives no unique reason for synaptic depression under reversed spike-timing. In fact, the presence and amplitude of depression of synaptic efficacies for reversed spike timing depends on how constraints are implemented in the optimization problem. Two different constraints, i.e., control of postsynaptic rates or control of temporal locality, are studied. The relation of our results to Spike-Timing Dependent Plasticity (STDP) and reinforcement learning is discussed.

*Current address: Department of Complexity Science and Engineering, Graduate School of Frontier Sciences, The University of Tokyo.

[†]Current address: IDIAP, Rue du Simplon 4, Case Postale 592, CH-1920 Martigny.

1 Introduction

Experimental evidence suggests that precise timing of spikes is important in several brain systems. In the barn owl auditory system, for example, coincidence detecting neurons receive volleys of temporally precise spikes from both ears (Carr and Konishi 1990). In the electrosensory system of mormyrid electric fish, medium ganglion cells receive input at precisely timed delays after electric pulse emission (Bell et al. 1997). Under the influence of a common oscillatory drive as present in the rat hippocampus or olfactory system, the strength of a constant stimulus is coded in the relative timing of neuronal action potentials (Hopfield 1995; Brody and Hopfield 2003; Mehta et al. 2002). In humans precise timing of first spikes in tactile afferents encode touch signals at the finger tips (Johansson and Birznieks 2004). Similar codes have also been suggested for rapid visual processing (Thorpe et al. 2001), and for the rat’s whisker response (Panzeri et al. 2001).

The precise timing of neuronal action potentials also plays an important role in Spike-Timing Dependent Plasticity (STDP). If a presynaptic spike arrives at the synapse *before* the postsynaptic action potential, the synapse is potentiated; if the timing is reversed the synapse is depressed (Markram et al. 1997; Zhang et al. 1998; Bi and Poo 1998; Bi and Poo 1999; Bi and Poo 2001). This biphasic STDP function is reminiscent of a temporal contrast or temporal derivative filter and suggests that STDP is sensitive to the temporal features of a neural code. Indeed, theoretical studies have shown that, given a biphasic STDP function, synaptic plasticity can lead to a stabilization of synaptic weight dynamics (Kempster et al. 1999; Song et al. 2000; Kempster et al. 2001; van Rossum et al. 2000; Rubin et al. 2001) while the neuron remains sensitive to temporal structure in the input (Gerstner et al. 1996; Roberts 1999; Kempster et al. 1999; Kistler and van Hemmen 2000; Rao and Sejnowski 2001; Gerstner and Kistler 2002a).

While the relative firing time of pre- and postsynaptic neurons, and hence temporal aspects of a neural code, play a role in STDP, it is, however, less clear whether STDP is useful to *learn* a temporal code. In order to further elucidate the computational function of STDP, we ask in this paper the following question: What is the ideal form of a STDP function in order to generate action potentials of the postsynaptic neuron with high temporal precision?

This question naturally leads to a supervised learning paradigm - i.e., the task to be learned by the neuron is to fire at a *predefined desired firing time* t^{des} . Supervised paradigms are common in machine learning in the context of classification and prediction problems (Minsky and Papert 1969; Haykin 1994; Bishop 1995), but have more recently also been studied for spiking neurons in feedforward and recurrent networks (Legenstein et al. 2005; Rao and Sejnowski 2001; Barber 2003; Gerstner et al. 1993; Izhikevich 2003). Compared to unsupervised or reward-based learning paradigms, supervised paradigms on the level of single spikes are obviously less relevant from a biological point, since it is questionable what type of signal could tell the neuron about the ‘desired’ firing time. Nevertheless, we think it is worth addressing the problem of supervised learning - firstly as a problem in it’s own right, and secondly as a starting point of spike base reinforcement learning (Xie and Seung 2004; Seung 2003). Reinforcement learning in a temporal coding paradigm implies that certain sequences of firing times are rewarded whereas others are not. The

“desired firing times” are hence defined *indirectly via the presence or absence of a reward signal*. The exact relation of our supervised paradigm to reward-based reinforcement learning will be presented in section 4. Section 2 introduces the stochastic neuron model and coding paradigm which are used to derive the results presented in section 3.

2 Model

2.1 Coding Paradigm

In order to explain our computational paradigm, we focus on the example of temporal coding of human touch stimuli (Johansson and Birznieks 2004), but the same ideas would apply analogously to the other neuronal systems with temporal codes mentioned above (Carr and Konishi 1990; Bell et al. 1997; Hopfield 1995; Brody and Hopfield 2003; Mehta et al. 2002; Panzeri et al. 2001). For a given touch stimulus, spikes in an ensemble of N tactile afferents occur in a precise temporal order. If the same touch stimulus with identical surface properties and force vector is repeated several times, the relative timing of action potentials is reliably reproduced whereas the spike timing in the same ensemble of afferents is different for other stimuli (Johansson and Birznieks 2004). In our model, we assume that all input lines, labeled by the index j with $1 \leq j \leq N$ converge onto one or several postsynaptic neurons. We think of the postsynaptic neuron as a detector for a given spatio-temporal spike patterns in the input. The full spike pattern detection paradigm will be used in Section 3.3. As a preparation and first steps towards the full coding paradigm we will also consider the response of a postsynaptic neuron to a single presynaptic spike (Section 3.1) or to one given spatio-temporal firing pattern (Section 3.2).

2.2 Neuron Model

Let us consider a neuron i which is receiving input from N presynaptic neurons. Let us denote the ensemble of all spikes of neuron j by $x_j = \{t_j^1, \dots, t_j^{N_j}\}$ where t_j^k denotes the time when neuron j fired its k^{th} spike. The spatio-temporal spike pattern of all presynaptic neurons $1 \leq j \leq N$ will be denoted by boldface $\mathbf{x} = \{x_1, \dots, x_N\}$.

A presynaptic spike elicited at time t_j^f evokes an excitatory postsynaptic potential (EPSP) of amplitude w_{ij} and time course $\epsilon(t - t_j^f)$. For the sake of simplicity, we approximate the EPSP time course by a double exponential

$$\epsilon(s) = \epsilon_0 \left[\exp\left(-\frac{s}{\tau_m}\right) - \exp\left(-\frac{s}{\tau_s}\right) \right] \Theta(s) \quad (1)$$

with a membrane time constant of $\tau_m = 10$ ms and a synaptic time constant of $\tau_s = 0.7$ ms which yields an EPSP rise time of 2 ms. Here $\Theta(s)$ denotes the Heaviside step function with $\Theta(s) = 1$ for $s > 0$ and $\Theta(s) = 0$ else. We set $\epsilon_0 = 1.3$ mV such that a spike at a synapse with $w_{ij} = 1$ evokes an EPSP with amplitude of approximately 1 mV. Since the EPSP amplitude is a measure of the strength of a synapse, we refer to w_{ij} also as the efficacy (or “weight”) of the synapse between neuron j and i .

Let us further suppose that the postsynaptic neuron i receives an additional input $I(t)$ that could either arise from a second group of neurons or from intracellular current injection. We think of the second input as a ‘teaching’ input that increases the probability that the neuron fires at or close to the desired firing time t^{des} . For the sake of simplicity we model the teaching input as a square current pulse $I(t) = I_0 \Theta(t - t^{\text{des}} + 0.5\Delta T) \Theta(t^{\text{des}} + 0.5\Delta T - t)$ of amplitude I_0 and duration ΔT . The effect of the teaching current on the membrane potential is

$$u_{\text{teach}}(t) = \int_0^\infty k(s) I(t-s) ds \quad (2)$$

with $k(s) = k_0 \exp(-s/\tau_m)$ where k_0 is a constant that is inversely proportional to the capacitance of the neuronal membrane.

In the context of the human touch paradigm discussed in section 2.1, the teaching input could represent some preprocessed visual information (‘object touched by fingers starts to slip now’), feedback from muscle activity (‘strong counterforce applied now’), cross-talk from other detector neurons in the same population (‘your colleagues are active now’), or unspecific modulatory input due to arousal or reward (‘be aware - something interesting happening now’).

In the context of training of recurrent networks (e.g. Rao and Sejnowski 2001), the teaching input consists of a short pulse of an amplitude that guarantees action potential firing.

The membrane potential of the postsynaptic neuron i (Spike Response Model, Gerstner and Kistler 2002b) is influenced by the EPSPs evoked by all afferent spikes of stimulus \mathbf{x} , the ‘teaching’ signal and the refractory effects generated by spikes t_i^f of the postsynaptic neuron

$$u_i(t|\mathbf{x}, y_i^i) = u_{\text{rest}} + \sum_{j=1}^N w_{ij} \sum_{t_j^f \in x_j} \epsilon(t - t_j^f) + \sum_{t_i^f \in y_i^i} \eta(t - t_i^f) + u_{\text{teach}}(t) \quad (3)$$

where $u_{\text{rest}} = -70$ mV is the resting potential, $y_i^i = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ is the set of postsynaptic spikes that occurred before t and t_i^F always denotes the last postsynaptic spike before t . On the right-hand side of Eq. (3), $\eta(s)$ denotes the spike-afterpotential generated by an action potential. We take

$$\eta(s) = \eta_0 \exp\left(-\frac{s}{\tau_m}\right) \Theta(s) \quad (4)$$

where $\eta_0 < 0$ is a reset parameter that describes how much the voltage is reset after each spike; for the relation to integrate-and-fire neurons see (Gerstner and Kistler 2002b). The spikes themselves are not modeled explicitly but reduced to formal firing times. Unless specified otherwise, we take $\eta_0 = -5$ mV.

In a deterministic version of the model, output spikes would be generated whenever the membrane potential u_i reaches a threshold ϑ . In order to account for intrinsic noise, and also for a small amount of ‘synaptic noise’ generated by stochastic spike arrival from additional excitatory and inhibitory presynaptic neurons which are not modeled explicitly we replace the strict threshold by a stochastic one. More precisely we adopt the following procedure (Gerstner and Kistler 2002b). Action

potentials of the postsynaptic neuron i are generated by a point process with time dependent stochastic intensity $\rho_i(t) = g(u_i(t))$ that depends non-linearly upon the membrane potential u_i . Since the membrane potential in turn depends on both the input and the firing history of the postsynaptic neuron, we write:

$$\rho_i(t|\mathbf{x}, y_t^i) = g(u_i(t|\mathbf{x}, y_t^i)). \quad (5)$$

We take an exponential to describe the stochastic escape across threshold, i.e. $g(u) = \rho_0 \exp\left(\frac{u-\vartheta}{\Delta u}\right)$ where $\vartheta = -50$ mV is the formal threshold, $\Delta u = 3$ mV is the width of the threshold region and therefore tunes the stochasticity of the neuron, and $\rho_0 = 1/\text{ms}$ is the stochastic intensity at threshold (see Fig. 1). Other choices of the escape function g are possible with no qualitative change of the results. For $\Delta u \rightarrow 0$, the model is identical to the deterministic leaky integrate-and-fire model with synaptic current injection (Gerstner and Kistler 2002b).

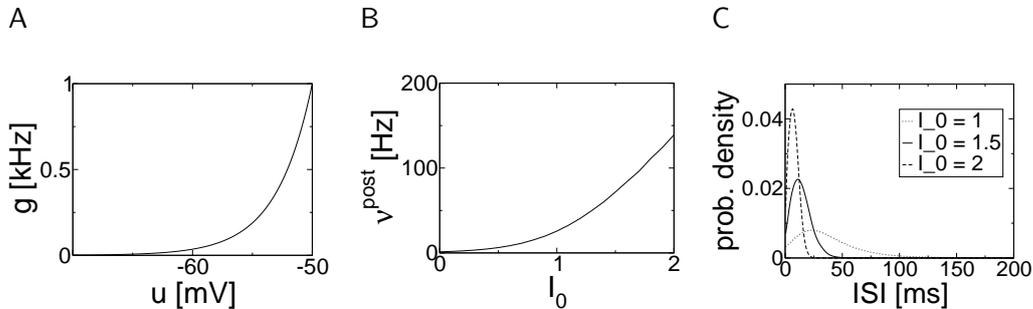


Figure 1: **A.** Escape rate $g(u) = \rho_0 \exp\left(\frac{u-\vartheta}{\Delta u}\right)$. **B.** Firing rate of the postsynaptic neuron as a function of the amplitude I_0 of a constant stimulation current (arbitrary units). **C.** Interspike interval (ISI) distribution for different input currents.

We note that the stochastic process, defined in Eq. (5) is similar to, but different from a Poisson process since the stochastic intensity depends on the set y_t of the previous spikes of the postsynaptic neuron. Thus the neuron model has some ‘memory’ of previous spikes.

2.3 Stochastic Generative Model

The advantage of the probabilistic framework introduced above via the noisy threshold is that it is possible to describe the probability density¹ $P_i(y|\mathbf{x})$ of an entire spike

¹For the sake of simplicity, we denoted the set of postsynaptic spikes from 0 to T by y instead of y_T .

train² $Y(t) = \sum_{t_i^f \in y} \delta(t - t_i^f)$ (see appendix A for details):

$$\begin{aligned} P_i(y|\mathbf{x}) &= \left(\prod_{t_i^f \in y} \rho_i(t_i^f|\mathbf{x}, y_{t_i^f}) \right) \exp \left(- \int_0^T \rho_i(s|\mathbf{x}, y_s) ds \right) \\ &= \exp \left(\int_0^T \log(\rho_i(s|\mathbf{x}, y_s)) Y(s) - \rho_i(s|\mathbf{x}, y_s) ds \right) \end{aligned} \quad (6)$$

Thus we have a generative model that allows us to describe explicitly the likelihood $P_i(y|\mathbf{x})$ of emitting a set of spikes y for a given input \mathbf{x} . Moreover, since the likelihood in Eq. (6) is a smooth function of its parameters, it is straightforward to differentiate it with respect to any variable. Let us differentiate $P_i(y|\mathbf{x})$ with respect to the synaptic efficacy w_{ij} , since this is a quantity that we will use later on:

$$\frac{\partial \log P_i(y|\mathbf{x})}{\partial w_{ij}} = \int_0^T \frac{\rho'_i(s|\mathbf{x}, y_s)}{\rho_i(s|\mathbf{x}, y_s)} [Y(s) - \rho_i(s|\mathbf{x}, y_s)] \sum_{t_j^f \in x_j} \epsilon(s - t_j^f) ds. \quad (7)$$

where $\rho'_i(s|\mathbf{x}, y_s) = \frac{dg}{du}|_{u=u_i(s|\mathbf{x}, y_s)}$.

In this paper, we propose three different *optimal models* called A, B and C (cf. Table 1). The models differ in the stimulation paradigm and the specific task of the neuron. In section 3, the task and hence the optimality criteria are supposed to be given explicitly. However, the task in model C could also be defined indirectly by the presence or absence of a reward signal as discussed in section 4.1. The common idea behind all three approaches (A-C) is the notion of optimal performance. Optimality is defined by an objective function L that is directly related to the likelihood formula of Eq. (6) and that can be maximized by changes of the synaptic weights. Throughout the paper, this optimization is done by a standard technique of gradient ascent:

$$\Delta w_{ij} = \alpha \frac{\partial L}{\partial w_{ij}} \quad (8)$$

with a learning rate α . Since the three models correspond to three different tasks, they have a slightly different objective function. Therefore, gradient ascent yields slightly different strategies for synaptic update. In the following we start with the simplest model with the aim to illustrate the basic principles that generalize to the more complex models.

3 Results

In this section we present synaptic updates rules derived by optimizing the likelihood of postsynaptic spike firing at some desired firing time t^{des} . The essence of the argument is introduced in a particularly simple scenario, where the neuron is stimulated by one presynaptic spike and the neuron is inactive except at the desired firing time t^{des} . This is the raw scenario that is further developed in several different directions.

²Capital Y is the spike train generated by the ensemble (lower case) y .

Firstly, we may ask the question of how the postsynaptic spike at the desired time t^{des} is generated: (i) the spike could simply be given by a supervisor. As always in maximum likelihood approaches, we then optimize the likelihood that this spike could have been generated by the neuron model (i.e., the generative model) given the known input; (ii) the spike could have been generated by an strong current pulse of short duration applied by the supervisor (teaching input). In this case the a priori likelihood that the generative model fires at or close to the desired firing time is much higher. The two conceptual paradigms give slightly different results as discussed in scenario A.

Secondly, we may, in addition to the spike at the desired time t^{des} allow for other postsynaptic spikes generated spontaneously. The consequences of spontaneous activity for the STDP function are discussed in scenario B. Thirdly, instead of imposing a single postsynaptic spike at a desired firing time t^{des} , we can think of a temporal coding scheme where the postsynaptic neuron responds to one (out of M) presynaptic spike patterns with a desired output spike train containing several spikes while staying inactive for the other $M - 1$ presynaptic spike patterns. This corresponds to a pattern classification task which is the topic of scenario C.

Moreover, optimization can be performed in an unconstrained fashion or under some constraint. As we will see in this section, the specific form of the constraint influences the results on STDP, in particular the strength of synaptic depression for *post-before-pre* timing. To emphasize this aspect, we discuss two different constraints. The first constraint is motivated by the observation that neurons have a preferred working point defined by a typical mean firing rate that is stabilized by homeostatic synaptic processes (Turrigiano and Nelson 2004). Penalizing deviations from a target firing rate is the constraint that we will use in scenario B. For very low target firing rate, the constraint reduces to the condition of ‘no activity’ which is the constraint implemented in scenario A.

The second type of constraint is motivated by the notion of STDP itself: changes of synaptic plasticity should depend on the relative timing of pre- and postsynaptic spike firing and not on other factors. If STDP is to be implemented by some physical or chemical mechanisms with finite time constants, we must require the STDP function to be local in time, i.e., the amplitude of the STDP function approaches zero for large time differences. This is the temporal locality constraint used in scenario C. While the unconstrained optimization problems are labeled with the subscript u (A_u, B_u, C_u), the constrained problems are marked by the subscript c (A_c, B_c, C_c) (c.f Table 1).

3.1 Scenario A: One Postsynaptic Spike Imposed

Let us start with a particularly simple model which consists of one presynaptic neuron and one postsynaptic neuron (c.f. Fig. 2A). Let us suppose that the task of the postsynaptic neuron i is to fire a single spike at time t^{des} in response to the input which consists of a single presynaptic spike at time t^{pre} , i.e. the input is $x = \{t^{\text{pre}}\}$ and the desired output of the postsynaptic neuron is $y = \{t^{\text{des}}\}$. Since there is only a single pre- and a single postsynaptic neuron involved, we drop in this section the indices j and i of the two neurons.

Unconstrained scenarios		Constrained scenarios	
A_u	Postsynaptic spike imposed $L^{A_u} = \log(\rho(t^{\text{des}}))$	A_c	No activity $L^{A_c} = L^{A_u} - \int_0^T \rho(t) dt$
B_u	Postsynaptic spike imposed + spontaneous activity $L^{B_u} = \log(\bar{\rho}(t^{\text{des}}))$	B_c	Stabilized activity $L^{B_c} = L^{B_u} - \frac{1}{T\sigma^2} \int_0^T (\bar{\rho}(t) - \nu_0)^2 dt$
C_u	Postsynaptic spike patterns imposed $L^{C_u} = \log \left(\prod_i P_i(y^i \mathbf{x}^i) \prod_{k \neq i} P_i(0 \mathbf{x}^k)^{\frac{\gamma}{M-1}} \right)$	C_c	Temporal locality constraint $L^{C_c} = L^{C_u}, P_{\Delta\Delta'} = a\delta_{\Delta\Delta'} (\Delta - \bar{T}_0)^2$

Table 1: Summary of the optimality criterion L for the three unconstrained scenarios (A_u , B_u , C_u) and the three constrained scenarios (A_c , B_c , C_c). The constraint for scenario C is not included in the likelihood function L^{C_c} itself, but rather in the deconvolution with a matrix P that penalizes quadratically the terms that are non-local in time. See appendix C for more details.

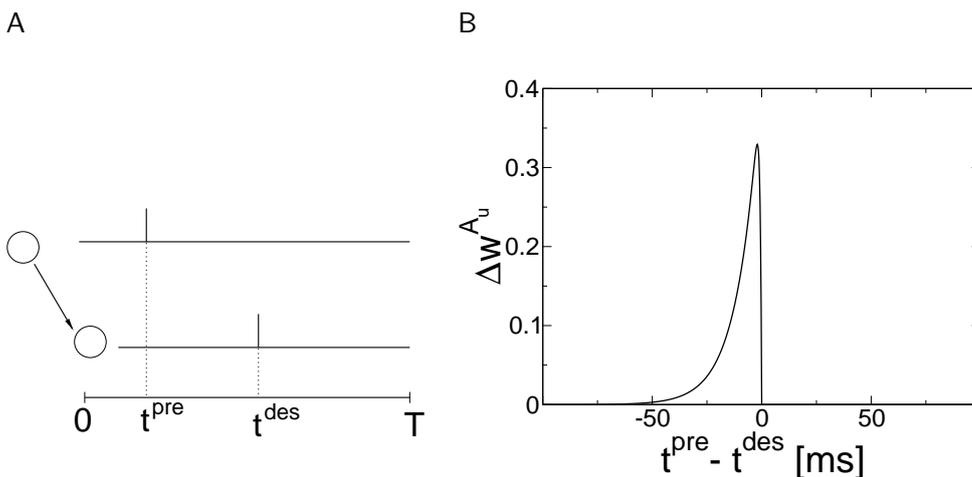


Figure 2: **A.** Scenario A: a single presynaptic neuron connected to a postsynaptic neuron with a synapse of weight w . **B.** Optimal weight change given by Eq. (10) for the scenario A_u . This weight change is exactly the mirror image of an EPSP.

3.1.1 Unconstrained scenario A_u : One Spike at t^{des}

In this subsection, we assume that the postsynaptic neuron has not been active in the recent past, i.e. refractory effects are negligible. In this case, we have $\rho(t|\mathbf{x}, y_t) = \rho(t|\mathbf{x})$ because of the absence of previous spikes. Moreover, since there is only a single presynaptic spike (i.e. $\mathbf{x} = \{t^{\text{pre}}\}$), we write $\rho(t|t^{\text{pre}})$ instead of $\rho(t|\mathbf{x})$.

Since the task of the postsynaptic neuron is to fire at time t^{des} , we can define the optimality criterion L^{A_u} as the log-likelihood of the firing intensity at time t^{des} , i.e.

$$L^{A_u} = \log \left(\rho(t^{\text{des}} | t^{\text{pre}}) \right) \quad (9)$$

The gradient ascent on this function leads to the following STDP function:

$$\Delta w^{A_u} = \alpha \frac{\partial L^{A_u}}{\partial w} = \alpha \frac{\rho'(t^{\text{des}}|t^{\text{pre}})}{\rho(t^{\text{des}}|t^{\text{pre}})} \epsilon(t^{\text{des}} - t^{\text{pre}}) \quad (10)$$

where $\rho'(t|t^{\text{pre}}) \equiv \frac{dg}{du}|_{u=u(t|t^{\text{pre}})}$. Since this optimal weight change Δw^{A_u} can be calculated for any presynaptic firing time t^{pre} , we get a STDP function which depends on the time difference $\Delta t = t^{\text{pre}} - t^{\text{des}}$ (c.f. Fig. 2B). As we can see directly from Eq. (10), the shape of the potentiation is exactly a mirror image of an EPSP. This result is independent of the specific choice of the function $g(u)$.

The drawback of this simple model becomes apparent, if the STDP function given by Eq. 10 is iterated over several repetitions of the experiment. Ideally, it should converge to an optimal solution given by $\Delta w^{A_u} = 0$ in Eq. (10). However, the optimal solution given by $\Delta w^{A_u} = 0$ is problematic: for $\Delta t < 0$, the optimal weight tends towards ∞ whereas for $\Delta t \geq 0$, there is no unique optimal weight ($\Delta w^{A_u} = 0, \forall w$). The reason of this problem is, of course, that the model describes only potentiation and includes no mechanisms for depression.

3.1.2 Constrained Scenario A_c : No Other Spikes Than at t^{des}

In order to get some insight of where the depression could come from, let us consider a small modification of the previous model. In addition to the fact that the neuron has to fire at time t^{des} , let us suppose that it should not fire anywhere else. This condition can be implemented by an application of Eq. (6) to the case of a single input spike $x = \{t^{\text{pre}}\}$ and a single output spike $y = \{t^{\text{des}}\}$. In terms of notation we set $P(y|x) = P(t^{\text{des}}|t^{\text{pre}})$ and similarly $\rho(s|x, y) = \rho(s|t^{\text{pre}}, t^{\text{des}})$ and use Eq. (6) to find:

$$P(t^{\text{des}}|t^{\text{pre}}) = \rho(t^{\text{des}}|t^{\text{pre}}) \exp \left[- \int_0^T \rho(s|t^{\text{pre}}, t^{\text{des}}) ds \right]. \quad (11)$$

Note that for $s \leq t^{\text{des}}$, the firing intensity does not depend on t^{des} , hence $\rho(s|t^{\text{pre}}, t^{\text{des}}) = \rho(s|t^{\text{pre}})$ for $s \leq t^{\text{des}}$. We define the objective function L^{A_c} as the log-likelihood of generating a single output spike at time t^{des} , given a single input spike at t^{pre} . Hence, with Eq. (11):

$$\begin{aligned} L^{A_c} &= \log(P(t^{\text{des}}|t^{\text{pre}})) \\ &= \log(\rho(t^{\text{des}}|t^{\text{pre}})) - \int_0^T \rho(s|t^{\text{pre}}, t^{\text{des}}) ds \end{aligned} \quad (12)$$

and the gradient ascent $\Delta w^{A_c} = \alpha \partial L^{A_c} / \partial w$ rule yields

$$\Delta w^{A_c} = \alpha \frac{\rho'(t^{\text{des}}|t^{\text{pre}})}{\rho(t^{\text{des}}|t^{\text{pre}})} \epsilon(t^{\text{des}} - t^{\text{pre}}) - \alpha \int_0^T \rho'(s|t^{\text{pre}}, t^{\text{des}}) \epsilon(s - t^{\text{pre}}) ds \quad (13)$$

Since we have a single postsynaptic spike at t^{des} , Eq. (13) can directly be plotted as a STDP function. In Fig. 3 we distinguish two different cases. In Fig. 3A we optimize the likelihood L^{A_c} in the absence of any teaching input. To understand this scenario we may imagine that a postsynaptic spike has occurred spontaneously at the desired firing time t^{des} . Applying the appropriate weight update calculated

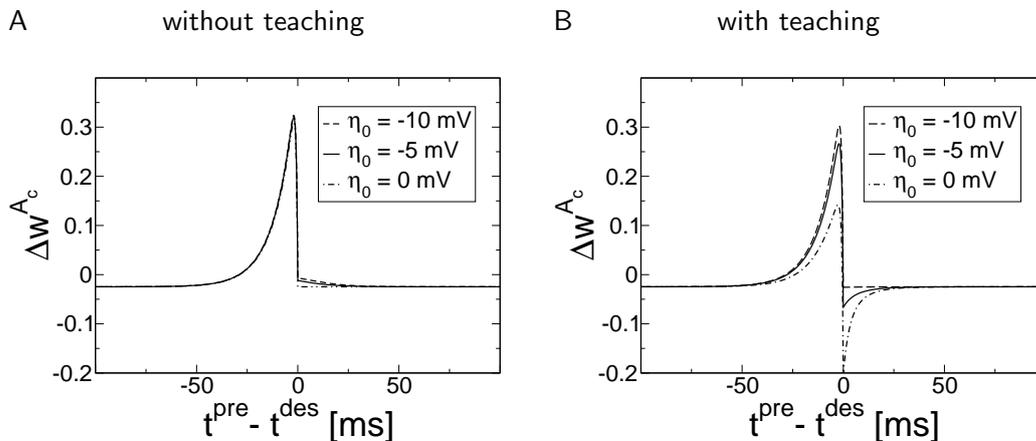


Figure 3: Optimal weight adaptation for scenario A_c given by Eq. (7) in the absence of a teaching signal (**A**) and in the presence of a teaching signal (**B**). The weight change in the *post-before-pre* is governed by the spike afterpotential $u_{AP}(t) = \eta(t) + u_{\text{teach}}(t)$. The duration of the teaching input is $\Delta T = 1$ ms. The amplitude of the current I_0 is chosen so that $\max_t u_{\text{teach}}(t) = 5$ mV. u_{rest} is chosen such that the spontaneous firing rate $g(u_{\text{rest}})$ matches the desired firing rate $1/T$, i.e. $u_{\text{rest}} = \Delta u \log \frac{1}{T\rho_0} + \theta \simeq -60$ mV. The weight strength is $w = 1$.

from Eq. (13) will make such a timing more likely the next time the presynaptic stimulus is repeated. The reset amplitude η_0 has only a small influence.

In Fig. 3B we consider a case where firing of the postsynaptic spike at the appropriate time was made highly likely by a teaching input of duration $\Delta T = 1$ ms centered around the desired firing t^{des} . The form of the STDP function depends on the amount η_0 of the reset. If there is no reset $\eta_0 = 0$, the STDP function shows strong synaptic depression of synapses that become active *after* the postsynaptic spike. This is due to the fact that the teaching input causes an increase of the membrane potential that decays back to rest with the membrane time constant τ_m . Hence the window of synaptic depression is also exponential with the same time constant. Qualitatively the same is true, if we include a weak reset. The form of the depression window remains the same, but its amplitude is reduced. The inverse of the effect occurs only for strong reset to or below resting potential. A weak reset is standard in applications of integrate-and-fire models to in vivo data and is one of the possibilities to explain the high coefficient of variation of neuronal spike trains in vivo (Bugmann, Christodoulou, and Taylor 1997; Troyer and Miller 1997).

A further property of the STDP functions in Fig. 3 is a negative offset for $|t^{\text{pre}} - t^{\text{des}}| \rightarrow \infty$. The amplitude of the offset can be calculated for $w \simeq 0$ and $\Delta t > 0$, i.e. $\Delta w_0 \simeq -\rho'(u_{\text{rest}}) \int_0^\infty \epsilon(s) ds$. This offset is due to the fact that we do not want spikes at other times than t^{des} . As a result, the optimal weight w^* (i.e. solution of $\Delta w^{A_u} = 0$), should be as negative as possible ($w^* \rightarrow -\infty$ or $w^* \rightarrow w^{\text{min}}$ in the presence of a lower bound) for $\Delta t > 0$ or $\Delta t \ll 0$.

3.2 Scenario B: Spontaneous Activity

The constraint in Scenario A_c of having strictly no other postsynaptic spikes than the one at time t^{des} may seem artificial. Moreover, it is this constraint which leads to the negative offset of the STDP function discussed at the end of the previous paragraph. In order to relax the constraint of “no spiking”, we allow in scenario B for a reasonable spontaneous activity. As above, we start with an unconstrained scenario B_u before we turn to the constrained scenario B_c .

3.2.1 Unconstrained scenario B_u : Maximize the Firing Rate at t^{des}

Let us start with the simplest model which includes spontaneous activity. Scenario B_u is the analog of the model A_u , but with two differences.

First, we include spontaneous activity in the model. Since $\rho(t|\mathbf{x}, y_t)$ depends on the spiking history for any given trial, we have to define a quantity which is independent of the specific realizations y of the postsynaptic spike train.

Secondly, instead of considering only one presynaptic neuron, we consider $N = 200$ presynaptic neurons, each of them emitting a single spike at time $t_j = j\delta t$, where $\delta t = 1$ ms (see Fig. 4A). The input pattern will therefore be described by the set of delayed spikes $\mathbf{x} = \{x_j = \{t_j\}, j = 1, \dots, N\}$. As long as we consider only a single spatio-temporal spike pattern in the input, it is always possible to relabel neurons appropriately so that neuron $j + 1$ fires after neuron j .

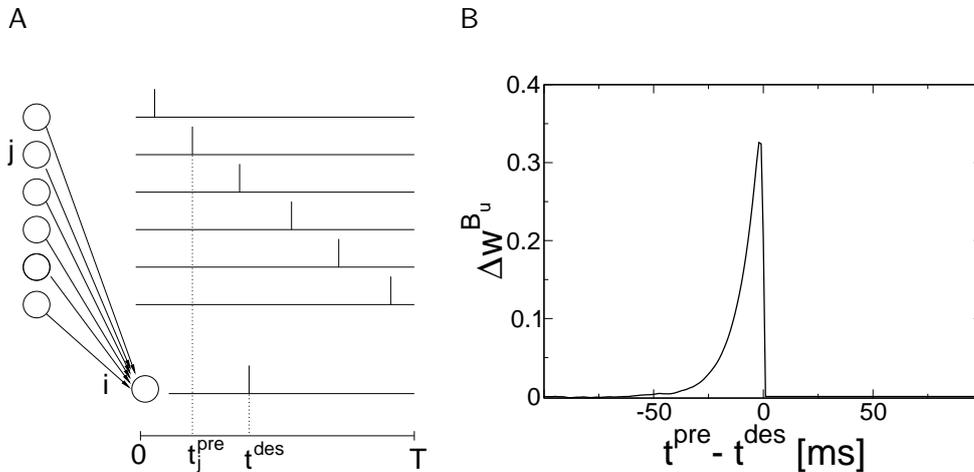


Figure 4: Scenario B. **A.** $N = 200$ presynaptic neurons are firing one after the other at time $t_j = j\delta t$ with $\delta t = 1$ ms. **B.** The optimal STDP function of scenario B_u .

Let us define the instantaneous firing rate $\bar{\rho}(t)$ that can be calculated by averaging $\rho(t|y_t)$ over all realizations of postsynaptic spike trains:

$$\bar{\rho}(t|\mathbf{x}) = \langle \rho(t|\mathbf{x}, y_t) \rangle_{y_t|\mathbf{x}}. \quad (14)$$

Here the notation $\langle \cdot \rangle_{y_t|\mathbf{x}}$ means taking the average over all possible configuration of postsynaptic spikes up to t for a given input \mathbf{x} . In analogy to a Poisson process, a specific spike train with firing times $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ is generated with

probability $P(y_t|\mathbf{x})$ given by Eq. (6). Hence, the average $\langle \cdot \rangle_{y_t|\mathbf{x}}$ of Eq. (14) can be written as follows (see appendix B for numerical evaluation of $\bar{\rho}(t)$):

$$\bar{\rho}(t|\mathbf{x}) = \sum_{F=0}^{\infty} \frac{1}{F!} \int_0^t \cdots \int_0^t \rho(t|\mathbf{x}, y_t) P(y_t|\mathbf{x}) dt_i^F \cdots dt_i^1. \quad (15)$$

Analogously to the model A_u , we can define the quality criterion as the log-likelihood L^{B_u} of firing at the desired time t^{des} :

$$L^{B_u} = \log(\bar{\rho}(t^{\text{des}}|\mathbf{x})). \quad (16)$$

Thus the optimal weight adaptation of synapse j is given by

$$\Delta w_j^{B_u} = \alpha \frac{\partial \bar{\rho}(t^{\text{des}}|\mathbf{x}) / \partial w_j}{\bar{\rho}(t^{\text{des}}|\mathbf{x})} \quad (17)$$

where $\frac{\partial \bar{\rho}(t|\mathbf{x})}{\partial w_j}$ is given by

$$\frac{\partial \bar{\rho}(t|\mathbf{x})}{\partial w_j} = \bar{\rho}'(t|\mathbf{x}) \epsilon(t - t_j) + \left\langle \rho(t|\mathbf{x}, y_t) \frac{\partial}{\partial w_j} \log P(y_t|\mathbf{x}) \right\rangle_{y_t|\mathbf{x}}, \quad (18)$$

$\frac{\partial}{\partial w_j} \log P(y_t|\mathbf{x})$ is given by Eq. (7) and $\bar{\rho}'(t|\mathbf{x}) = \left\langle \frac{dg}{du} \Big|_{u=u(t|\mathbf{x}, y_t)} \right\rangle_{y_t|\mathbf{x}}$.

Figure 4B shows that, for our standard set of parameters, the differences to scenario A_u are negligible.

Figure 5A depicts the STDP function for various values of the parameter Δu at a higher postsynaptic firing rate. We can see a small undershoot in the *pre-before-post* region. The presence of this small undershoot can be understood as follows: enhancing a synapse of a presynaptic neuron that fires too early would induce a postsynaptic spike that arrives before the desired firing time and therefore, because of refractoriness, would prevent the generation of a spike at the desired time. The depth of this undershoot decreases with the stochasticity of the neuron and increases with the amplitude of the refractory period (if $\eta_0 = 0$, there is no undershoot). In fact, correlations between pre- and postsynaptic firing reflect the shape of an EPSP in the high-noise regime, whereas they show a trough for low noise (Poliakov et al. 1997; Gerstner 2001). Our theory shows that the *pre-before-post* region of the optimal plasticity function is a mirror image of these correlations.

3.2.2 Constrained scenario B_c : Firing Rate Close to ν_0

In analogy to model A_c we now introduce a constraint. Instead of imposing strictly no spikes at times $t \neq t^{\text{des}}$, we can relax the condition and minimize deviations of the instantaneous firing rate $\bar{\rho}(t|\mathbf{x}, t^{\text{des}})$ from a reference firing rate ν_0 . This can be done by introducing into Eq. (16) a penalty term P_B given by

$$P_B = \exp \left(-\frac{1}{T} \int_0^T \frac{(\bar{\rho}(t|\mathbf{x}, t^{\text{des}}) - \nu_0)^2}{2\sigma^2} dt \right). \quad (19)$$

For small σ , deviations from the reference rate yields a large penalty. For $\sigma \rightarrow \infty$, the penalty term has no influence. The optimality criterion is a combination of a

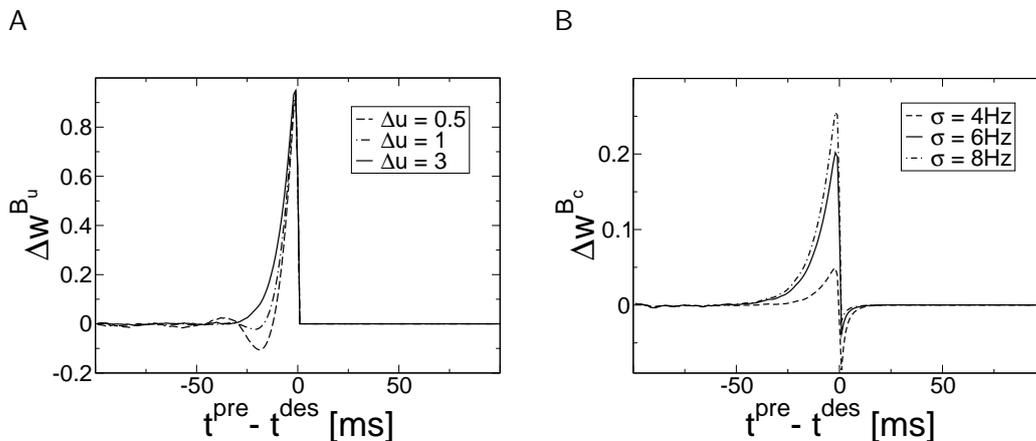


Figure 5: **A.** The optimal STDP functions of scenario B_u for different level of stochasticity described by the parameter Δu . The standard value ($\Delta u = 3$ mV) is given by the solid line, decreased noise ($\Delta u = 1$ mV and $\Delta u = 0.5$ mV) are indicated by dot-dashed and dashed lines respectively. In the low-noise regime, enhancing a synapse which fires slightly too early can prevent the firing at the desired firing time t^{des} due to refractoriness. To increase the firing rate at t^{des} it is thence advantageous to decrease the firing probability some time before t^{des} . Methods: For each value of Δu , the initial weight w_0 are set such that the spontaneous firing rate is $\bar{\rho} = 30$ Hz. In all three cases, Δw has been multiplied by Δu in order to normalize the amplitude of the STDP function. Reset: $\eta_0 = -5$ mV. **B.** Scenario B_c . Optimal STDP function for scenario B_c given by Eq. (21) for a teaching signal of duration $\Delta T = 1$ ms. The maximal increase of the membrane potential after 1 ms of stimulation with the teaching input is $\max_t u_{\text{teach}}(t) = 5$ mV. Synaptic efficacies w_{ij} are initialized such that $u_0 = -60$ mV which gives a spontaneous rate of $\bar{\rho} = \nu_0 = 5$ Hz. Standard noise level: $\Delta u = 3$ mV.

high firing rate $\bar{\rho}$ at the desired time under the constraint of small deviations from the reference rate ν_0 . If we impose the penalty as a multiplicative factor and take as before the logarithm, we get:

$$L^{B_c} = \log \left(\bar{\rho}(t^{\text{des}}|\mathbf{x}) P_B \right) \quad (20)$$

Hence the optimal weight adaptation is given by

$$\Delta w_j^{B_c} = \alpha \frac{\partial \bar{\rho}(t^{\text{des}}|\mathbf{x}) / \partial w_j}{\bar{\rho}(t^{\text{des}}|\mathbf{x})} - \frac{\alpha}{T\sigma^2} \int_0^T (\bar{\rho}(t|\mathbf{x}, t^{\text{des}}) - \nu_0) \frac{\partial}{\partial w_j} \bar{\rho}(t|\mathbf{x}, t^{\text{des}}) dt. \quad (21)$$

Since in scenario B each presynaptic neuron j fires exactly once at time $t_j = j\delta t$ and the postsynaptic neuron is trained to fire at time t^{des} , we can interpret the weight adaptation $\Delta w_j^{B_c}$ of Eq. (21) as a STDP function Δw^{B_c} which depends on the time difference $\Delta t = t^{\text{pre}} - t^{\text{des}}$. Fig. 5 shows this STDP function for different values of the free parameter σ of Eq. (19). The higher the standard deviation σ , the less effective is the penalty term. In the limit of $\sigma \rightarrow \infty$, the penalty term can be ignored and the situation is identical to that of scenario B_u .

3.3 Scenario C: Pattern Detection

3.3.1 Unconstrained Scenario C_u : Spike Pattern Imposed

This last scenario is a generalization of the scenario A_c . Instead of restricting the study to a single pre- and postsynaptic neuron, we consider N presynaptic neurons and M postsynaptic neurons (see Fig. 6). The idea is to construct M independent *detector neurons*. Each detector neuron $i = 1, \dots, M$, should respond best to a specific prototype stimulus, say \mathbf{x}^i , by producing a desired spike train y^i , but should not respond to other stimuli, i.e. $y^i = 0, \forall \mathbf{x}^k, k \neq i$ (see Fig. 7). The aim is to find a set of synaptic weights that maximize the probability that neuron i produces y^i when \mathbf{x}^i is presented and produces no output when $\mathbf{x}^k, k \neq i$ is presented. Let the likelihood function L^{C_u} be

$$L^{C_u} = \log \left(\prod_{i=1}^M P_i(y^i | \mathbf{x}^i) \prod_{k=1, k \neq i}^M P_i(0 | \mathbf{x}^k)^{\gamma} \right) \quad (22)$$

where $P_i(y^i | \mathbf{x}^i)$ (c.f Eq. (6)) is the probability that neuron i produces the spike train y^i when the stimulus \mathbf{x}^i is presented. The parameter γ characterizes the relative importance of the patterns that should not be learned compared to those that should be learned. We get

$$L^{C_u} = \sum_{i=1}^M \log(P_i(y^i | \mathbf{x}^i)) + \gamma \left\langle \log(P_i(0 | \mathbf{x}^k)) \right\rangle_{\mathbf{x}^k \neq \mathbf{x}^i} \quad (23)$$

where the notation $\langle \cdot \rangle_{\mathbf{x}^k \neq \mathbf{x}^i} \equiv \frac{1}{M-1} \sum_{k \neq i}^M$ means taking the average over all patterns other than \mathbf{x}^i . The optimal weight adaptation yields

$$\Delta w_{ij}^C = \alpha \frac{\partial}{\partial w_{ij}} \log(P_i(y^i | \mathbf{x}^i)) + \alpha \gamma \left\langle \frac{\partial}{\partial w_{ij}} \log(P_i(0 | \mathbf{x}^k)) \right\rangle_{\mathbf{x}^k \neq \mathbf{x}^i} \quad (24)$$

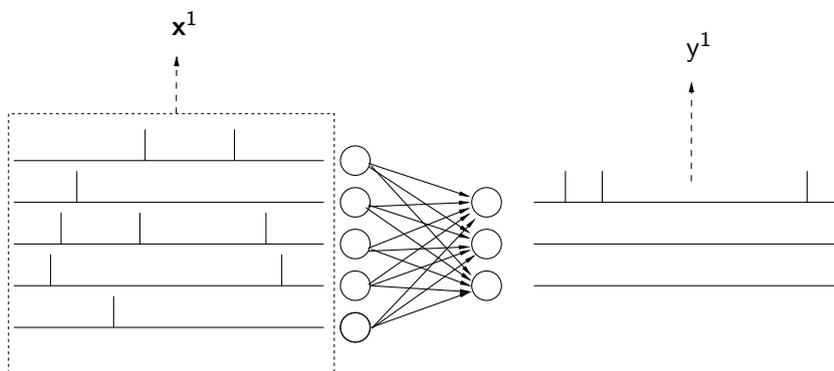


Figure 6: Scenario C. N presynaptic neurons are fully connected to M postsynaptic neurons. Each postsynaptic neuron is trained to respond to a specific input pattern and not respond to $M - 1$ other patterns as described by the objective function of Eq. (22).

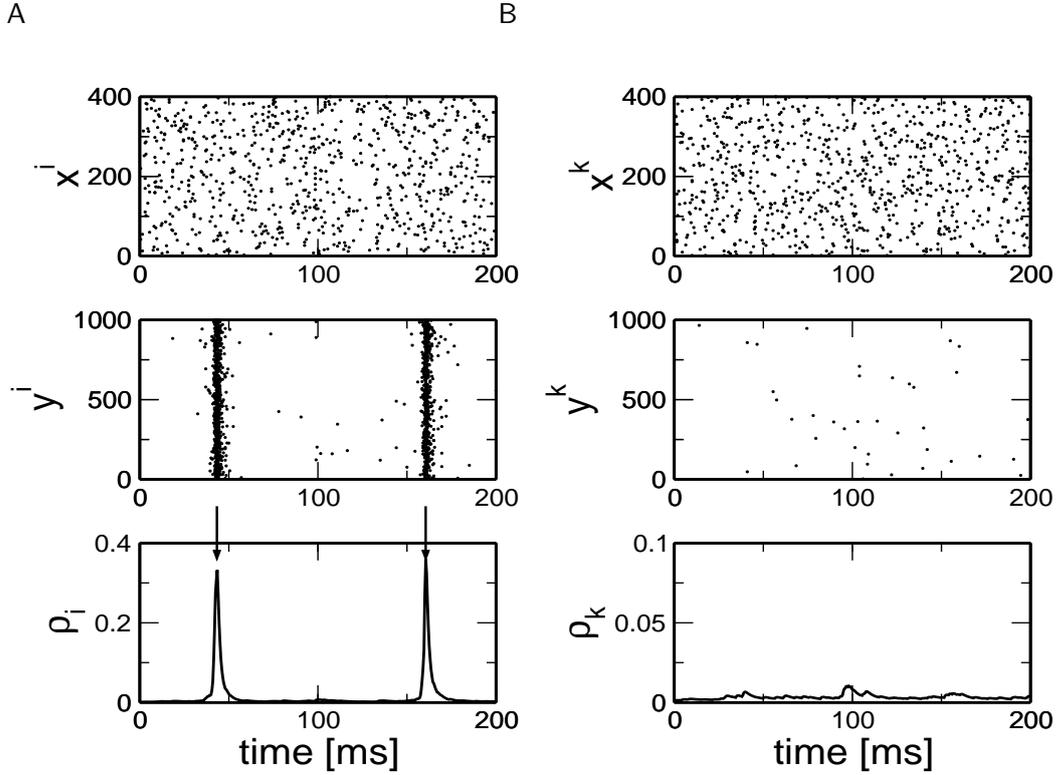


Figure 7: Pattern detection after learning. **Top.** The *left* raster plot represents the input pattern the i^{th} neuron has to be sensitive to. Each line corresponds to one of the $N = 400$ presynaptic neurons. Each dot represents an action potential. The *right* figure represents one of the patterns the i^{th} neuron should not respond to. **Middle.** The *left* raster plot corresponds to 1000 repetitions of the output of neuron i when the corresponding pattern \mathbf{x}^i is presented. The *right* plot is the response of neuron i to one of the pattern it should not respond to. **Bottom.** The *left* graph represents the probability density of firing when pattern \mathbf{x}^i is presented. This plot can be seen as the PSTH of the middle graph. Arrows indicate the supervised timing neuron i learned. The *right* graph describes the probability density of firing when pattern \mathbf{x}^k is presented. Note the different scales of vertical axis.

The learning rule of Eq. (24) gives the optimal weight change for each synapse and can be evaluated after presentation of all pre- and postsynaptic spike patterns, i.e. it is a “batch” update rule. Since each pre- and postsynaptic neuron emit many spikes in the interval $[0, T]$, we can not directly interpret the result of Eq. (24) as a function of the time difference $\Delta t = t^{\text{pre}} - t^{\text{des}}$ as we did in scenario A or B.

Ideally, we would like to write the total weight change of the optimal rule given by Eq. (24) as a sum of contributions

$$\Delta w_{ij}^C = \sum_{\substack{t^{\text{pre}} \in x_j^i \\ t^{\text{des}} \in y^i}} \Delta W^{\text{Cu}}(t^{\text{pre}} - t^{\text{des}}), \quad (25)$$

where $\Delta W^{\text{Cu}}(t^{\text{pre}} - t^{\text{des}})$ is a STDP function and the summation runs over all pairs of pre- and postsynaptic spikes. The number of pairs of pre- and postsynaptic spikes with a given time shift is given by the correlation function which is best defined in discrete time. We assume time steps of duration $\delta t = 0.5$ ms. Since the correlation will depend on the presynaptic neuron j and the postsynaptic neuron i under consideration, we introduce a new index $k = N(i - 1) + j$. We define the correlation in discrete time by its matrix elements $C_{k\Delta}$ that describe the correlation between the presynaptic spike train $X_j^i(t)$ and the postsynaptic spike train $Y^i(t - T_0 + \Delta\delta t)$. For example, $C_{3\Delta} = 7$ implies that 7 spike pairs of presynaptic neuron $j = 3$ with postsynaptic neuron $i = 1$ have a relative time shift of $T_0 - \Delta\delta t$. With this definition, we can rewrite Eq. (25) in vector notation (see appendix C.1 for more details):

$$\Delta \mathbf{w}^{\text{C}} \stackrel{!}{=} C \Delta \mathbf{W}^{\text{Cu}} \quad (26)$$

where $\Delta \mathbf{w}^{\text{C}} = (\Delta w_{11}^{\text{C}}, \dots, \Delta w_{1N}^{\text{C}}, \Delta w_{21}^{\text{C}}, \dots, \Delta w_{MN}^{\text{C}})^T$ is the vector containing all the optimal weight change given by Eq. (24) and $\Delta \mathbf{W}^{\text{Cu}}$ is the vector containing the discretized STDP function with components $\Delta W_{\Delta}^{\text{Cu}} = \Delta W^{\text{Cu}}(-T_0 + \Delta\delta t)$ for $1 \leq \Delta \leq 2\tilde{T}_0$ with $\tilde{T}_0 = T_0/\delta$. In particular, the center of the STDP function (i.e. $t^{\text{pre}} = t^{\text{des}}$) corresponds to the index $\Delta = \tilde{T}_0$. The symbol $\stackrel{!}{=}$ expresses the fact that we want to find $\Delta \mathbf{W}^{\text{Cu}}$ such that $\Delta \mathbf{w}^{\text{C}}$ is as close as possible to $C \Delta \mathbf{W}^{\text{Cu}}$. By taking the pseudo-inverse $C^+ = (C^T C)^{-1} C^T$ of C , we can invert Eq. (26) and get

$$\Delta \mathbf{W}^{\text{Cu}} = C^+ \Delta \mathbf{w}^{\text{C}} \quad (27)$$

The resulting STDP function is plotted in Fig. 8A. As it was the case for the scenario A_u , the STDP function exhibits a negative offset. In addition to the fact the postsynaptic neuron i should not fire at other times than the ones given by y^i , it should also not fire whenever pattern \mathbf{x}^k , $k \neq i$ is presented. The presence of the negative offset is due to those two factors.

3.3.2 Constrained Scenario C_c : Temporal Locality

In the previous paragraph, we obtained a STDP function with a negative offset. This negative offset does not seem realistic because it implies that the STDP function is not localized in time. In order to impose temporal locality (finite memory span of the learning rule) we modify Eq. (27) in the following way (see appendix C.2 for more details):

$$\Delta \mathbf{W}^{\text{Cc}} = (C^T C + P)^{-1} C^T \Delta \mathbf{w}^{\text{C}} \quad (28)$$

where P is a diagonal matrix which penalizes non-local terms. In this paper, we take a quadratic suppression of terms that are non-local in time. With respect to a postsynaptic spike at t^{des} , the penalty term is proportional to $(t - t^{\text{des}})^2$. In matrix notation, and using our convention that the postsynaptic spike corresponds to $\Delta = \tilde{T}_0$, we have:

$$P_{\Delta\Delta'} = a \delta_{\Delta\Delta'} \left(\Delta - \tilde{T}_0 \right)^2 \quad (29)$$

The resulting STDP functions for different values of a are plotted in Fig. 8B. The higher the parameter a , the more non-local terms are penalized, the narrower is the STDP function.

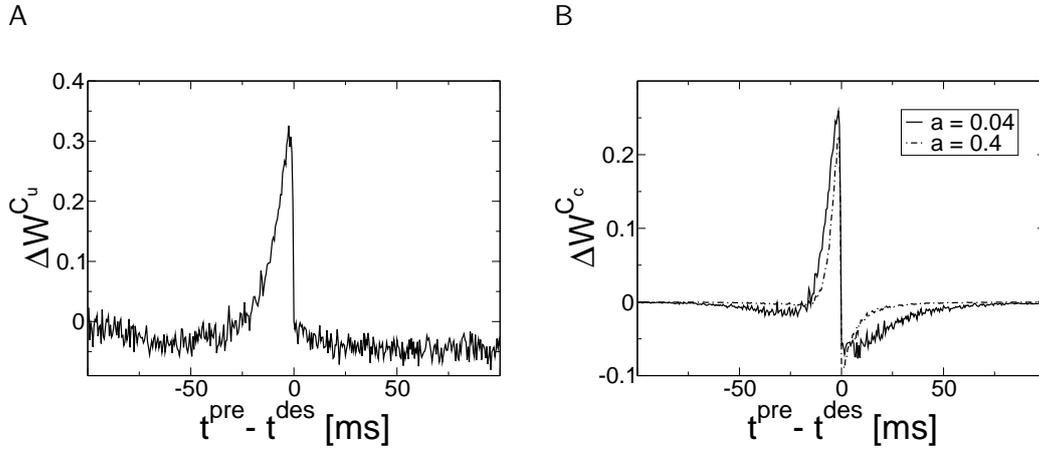


Figure 8: **A.** Optimal weight change for scenario C_u . In this case, no locality constraint is imposed and the result is similar to the STDP function of scenario A_c (with $\eta_0 = 0$ and $u_{\text{teach}}(t) = 0$) represented on Fig. 3. **B.** Optimal weight change for scenario C_c as a function of the locality constraint characterized by a . The stronger the importance of the locality constraint, the narrower is the spike-spike interaction. For **A** and **B**: $M = 20$, $\eta_0 = -5$ mV. The initial weights w_{ij} are chosen so that the spontaneous firing rate matches the imposed firing rate.

Fig. 9A shows the STDP functions for various number of patterns M . No significant change can be observed for different numbers of input patterns M . This is due to the appropriately chosen normalization factor $1/(M - 1)$ in the exponent of Eq. (22).

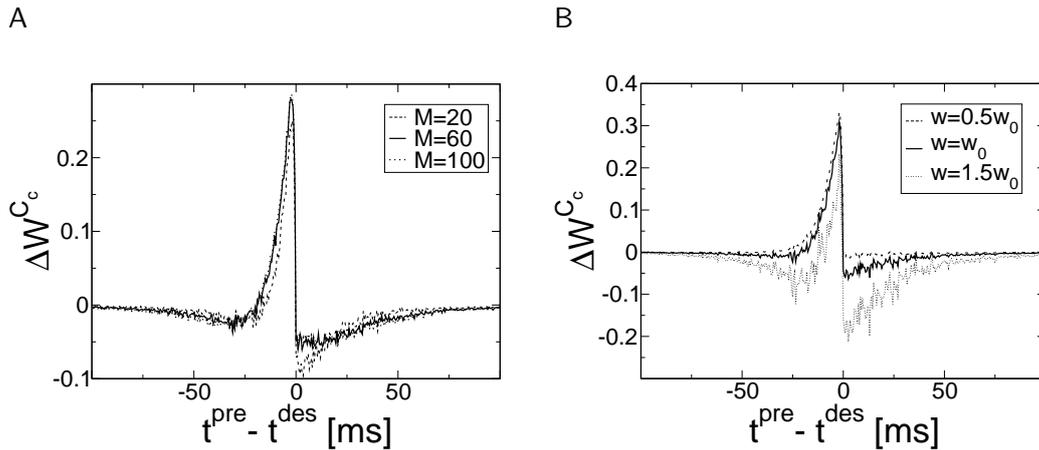


Figure 9: **A.** Optimal STDP function as a function of the number of input patterns M . ($a = 0.04$, $N = 400$) **B.** Optimal weight change as a function of the weight w . If the weights are small (*dashed line*) potentiation dominates whereas if they are big (*dotted line*) depression dominates.

The target spike trains y^i have a certain number of spikes during the time window T , i.e. they set a target value for the mean rate. Let $\nu^{\text{post}} = \frac{1}{TM} \sum_{i=1}^M \int_0^T y^i(t) dt$ be the imposed firing rate. Let w_0 denote the amplitude of the synaptic strength such that the firing rate $\bar{\rho}_{w_0}$ given by those weights is identical to the imposed firing rate: $\bar{\rho}_{w_0} = \nu^{\text{post}}$. If the actual weights are smaller than w_0 , almost all the weights should increase whereas if they are bigger than w_0 , depression should dominate (c.f. Fig 9B). Thus the exact form of the optimal STDP function depends on the initial weight value w_0 . Alternatively, homeostatic process could assure that the mean weight value is always in the appropriate regime.

In Eqs. (25) and (26) we imposed that the total weight change should be generated as a sum over *pairs* of pre- and postsynaptic spikes. This is an assumption which has been made in order to establish a link to standard STDP theory and experiments where spike pairs have been in the center of interest (Gerstner et al. 1996; Kempter et al. 1999; Kistler and van Hemmen 2000; Markram et al. 1997; Bi and Poo 1998; Zhang et al. 1998). It is, however, clear by now that the timing of spike pairs is only one of several factors contributing to synaptic plasticity. We therefore asked how much we miss if we attribute the ‘optimal’ weight changes calculated in Eq. (24) to spike-pair effects only. To answer this question we compared the optimal weight change Δw_{ij}^C from Eq. (24) with that derived from the pair-based STDP rule $\Delta w_{ij}^{\text{rec}} = \sum_{t^{\text{pre}} \in x_j^i} \sum_{t^{\text{des}} \in y^i} \Delta W^{C_c}(t^{\text{pre}} - t^{\text{des}})$ with or without locality constraint, i.e. for different values of the locality parameter ($a = 0, 0.04, 0.4$): see Fig. 10. More precisely, we simulate $M = 20$ detector neurons, each one of them having $N = 400$ presynaptic inputs, so each subplot of Fig. 10 contains 8000 points. Each point in a graph corresponds to the optimal change of one weight for one detector neuron (x axis) compared to the weight change of the very same weight due to pair based STDP (y axis). We found that in the absence of a locality constraint the pair-wise contributions are well correlated with the optimal weight changes. With strong locality constraints the quality of the correlation drops significantly. However, for a weak locality constraint that corresponds to a STDP function with reasonable potentiation and depression regimes, the correlation of the pair-based STDP rule with the optimal update is still good. This suggests that synaptic updates with a STDP-function based on pairs of pre- and postsynaptic spikes is close to optimal in the ‘pattern detection’ paradigm.

4 Discussion

4.1 Supervised versus Unsupervised and Reinforcement Learning

Our approach is based upon the maximization of the probability of firing at desired times t^{des} with or without constraints. From the point of view of machine learning, this is a supervised learning paradigm implemented as a maximum likelihood approach using the spike response model with escape noise as a generative model. Our work can be seen as a continuous-time extension of the maximum likelihood approach proposed in Barber (2003).

The starting point of all supervised paradigms is the comparison of a desired output with the actual output a neuron has, or would have, generated. The difference between the desired and actual output is then used as the driving signal for synaptic

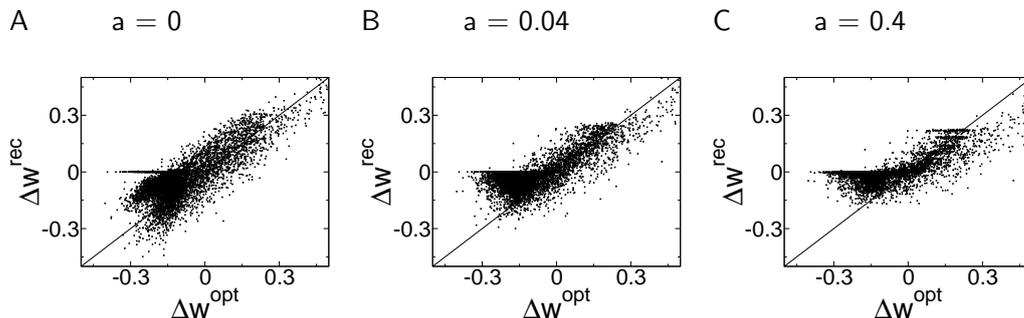


Figure 10: Correlation plot between the optimal synaptic weight change $\Delta \mathbf{w}^{\text{opt}} = \Delta \mathbf{w}^{\text{Cu}}$ and the reconstructed weight change $\Delta \mathbf{w}^{\text{rec}} = C \Delta \mathbf{W}^{\text{Cc}}$ using the temporal locality constraint. **A.** No locality constraint, i.e. $a = 0$. Deviations from the diagonal are due to the fact that the optimal weight change given by Eq. (24) can not be perfectly accounted for the sum of pair effects. The mean deviations are given by Eq. (43). **B.** A weak locality constraint ($a = 0.04$) almost does not change the quality of the weight change reconstruction. **C.** Strong locality constraint ($a = 0.4$). The horizontal lines arise since most synapses are subject to a few strong updates induced by pairs of pre- and postsynaptic spike times with small time shifts.

Unconstrained scenarios		Constrained scenarios	
A_u	<i>pre-before-post</i> LTP \sim EPSP	A_c	<i>post-before-pre</i> LTD (or LTP) \sim spike afterpot.
B_u	<i>pre-before-post</i> LTP/LTD \sim reverse correlation	B_c	<i>post-before-pre</i> LTD \sim increased firing rate
C_u	<i>pre-before-post</i> LTP \sim EPSP LTD \sim background patterns	C_c	<i>post-before-pre</i> LTD \sim background patterns \sim temporal locality

Table 2: Main results for each scenario.

updates in typical model approaches (Minsky and Papert 1969; Haykin 1994; Bishop 1995). How does this compare to experimental approaches? Experiments focusing on STDP have been mostly performed in vitro (Markram et al. 1997; Magee and Johnston 1997; Bi and Poo 1998). Since in typical experimental paradigms firing of the postsynaptic neuron is enforced by strong pulses of current injection, the neuron is not in a natural ‘unsupervised’ setting; on the other hand, the situation is also not fully supervised, since there is never a conflict between the desired and actual output of a neuron. In one of the rare in vivo experiments to STDP (Frégnac et al. 1988; Frégnac et al. 1992), the spikes of the postsynaptic neuron are also imposed by current injection. Thus, a classification of STDP experiments in terms of supervised, unsupervised, or reward based, is not as clearcut as it may seem at a first glance.

From the point of view of neuroscience, paradigms of unsupervised or reinforcement are probably much more relevant than the supervised scenario discussed here. However, most of our results from the supervised scenario analyzed in this paper, can be reinterpreted in the context of reinforcement learning following the approach

proposed by Xie and Seung (2004). To illustrate the link between reinforcement learning and supervised learning, we define a global reinforcement signal $R(\mathbf{x}, y)$ which depends on the spike timing of the presynaptic neurons \mathbf{x} and the postsynaptic neuron y . The quantity optimized in reinforcement learning is the expected reward $\langle R \rangle_{\mathbf{x}, y}$ averaged over all pre- and postsynaptic spike trains, i.e.

$$\langle R \rangle_{\mathbf{x}, y} = \sum_{\mathbf{x}, y} R(\mathbf{x}, y) P(y|\mathbf{x}) P(\mathbf{x}) \quad (30)$$

If the goal of learning is to maximize the expected reward, we can define a learning rule which achieves this goal by changing synaptic efficacies in direction of the gradient of the expected reward $\langle R \rangle_{\mathbf{x}, y}$:

$$\langle \Delta w \rangle_{\mathbf{x}, y} = \alpha \left\langle R(\mathbf{x}, y) \frac{\partial \log P(y|\mathbf{x})}{\partial w} \right\rangle_{\mathbf{x}, y} \quad (31)$$

where α is a learning parameter and $\frac{\partial \log P(y|\mathbf{x})}{\partial w}$ is the quantity we discussed in this paper. Thus the quantities optimized in our supervised paradigm re-appear naturally in a reinforcement learning paradigm.

For an intuitive interpretation of the link between reinforcement learning and supervised learning consider a postsynaptic spike that (spontaneously) occurred at time t_0 . If no reward is given, no synaptic change takes place. However, if the postsynaptic spike at t_0 is linked to a rewarding situation, the synapse will try to recreate in the next trial a spike at the same time, i.e., t_0 has the role of the desired firing time t^{des} introduced in this paper. Thus the STDP function with respect to a postsynaptic spike at t^{des} derived in this paper, can be seen as the spike timing dependence which maximizes the expected reward in a spike-based reinforcement learning paradigm.

4.2 Interpretation of STDP Function

Let us now summarize and discuss our results in a broader context. In all three scenarios, we found an STDP function with potentiation for *pre-before-post* timing. Thus this result is structurally stable and independent of model details. However, depression for *post-before-pre* timing does depend on model details.

In scenario A, we have seen that the behavior of the *post-before-pre* region is determined by the spike afterpotential (see table 2 for a result summary of the three models). In the presence of a teaching input and firing rate constraints, a weak reset of the membrane potential after the spike means that the neuron effectively has a depolarizing spike after potential (DAP). In experiments, DAPs have been observed by Feldman (2000), Markram et al. (1997) and Bi and Poo (1998) for strong presynaptic input. Other studies, however, have shown that the level of depression does not depend on the postsynaptic membrane potential (Sjöström et al. 2001). In any case, a weak reset (i.e., to a value below threshold rather than to the resting potential) is consistent with the findings of other researchers that used integrate-and-fire models to account for the high coefficient of variation of spike trains in vivo (Bugmann et al. 1997; Troyer and Miller 1997).

In the presence of spontaneous activity (c.f. scenario B), a constraint on the spontaneous firing rate causes the optimal weight change to elicit a depression of

presynaptic spikes that arrive immediately after the postsynaptic one. In fact, the reason of the presence of the depression in scenario B_c is directly related to the presence of a DAP caused by the strong teaching stimulus. In both scenarios A and B, depression occurs in order to compensate the increased firing probability due to the DAP.

In scenario C, it has been shown that the best way to adapt the weights (in a task where the postsynaptic neuron has to detect a specific input pattern among others) can be described as a STDP function. This task is similar to the one in Izhikevich (2003) in the sense that a neuron is designed to be sensitive to a specific input pattern, but different since our work does not assume any axonal delays. The depression part in this scenario arises from a locality constraint. We impose that weight changes are explained by a sum of pair-based STDP functions.

There are various ways of defining objective functions and we have used three different objective functions in this paper. The formulation of an objective function gives a mathematical expression of the ‘functional role’ we assign to a neuron. The functional role depends on the type of coding (temporal coding or rate coding) and hence on the information the postsynaptic neurons will read out. The functional role also depends on the task or context in which a neuron is embedded. It might seem that different tasks and coding schemes could thus give rise to a huge number of objective functions. However the reinterpretation of our approach in the context of reinforcement learning provides a unifying viewpoint: even if the functional role of some neurons in a specific region of the brain can be different from other neurons of a different region, it is still possible to see the different objective functions as different instantiations of the same underlying concept - the maximization of the reward, where the reward is task specific.

More specifically, all objective functions used in this paper maximized the firing probability at a desired firing time t^{des} - reflecting the fact that in the framework of timing based codes, the task of a neuron is to fire at precise moments in time. With a different assumption on the neuron’s role on signal processing, different objective functions need to be used. An extreme case is a situation, where the neuron’s task is to avoid firing at time t^{des} . A good illustration is given by the experiments done in the electrosensory lobe (ELL) of the electric fish (Bell et al. 1997). These cells receive two sets of input: the first one contains the pulses coming from the electric organ while the second input conveys information about the sensory stimulus. Since a large fraction of the sensory stimulus can be predicted by the information coming from the electric organ, it is computationally interesting to subtract the predictable contribution and focus only on the unpredictable part of the sensory stimulus. In this context, a reasonable task would be precisely to ask the neuron not to fire at time t^{des} where t^{des} is the time where the predictable simulation arrives and this task could be defined indirectly via an appropriate reward signal. An objective function of this type would, in the end, reverse the sign of the weight change of the causal part (LTD for the *pre-before-post* region), and this is precisely what is seen experimentally (Bell et al. 1997).

In our framework, the definition of the objective function is closely related to the neuronal coding. In scenario C, we postulate that neurons emit a precise spike train whenever the “correct” input is presented and be silent otherwise. This coding scheme is clearly not the most efficient one. Another possibility is to impose

postsynaptic neurons to produce a specific but different spike train for each input pattern and not only for the “correct” input. Such a modification of the scenario does not dramatically change the results. The only effect is to reduce the amount of depression and increase the amount of potentiation.

4.3 Optimality Approaches versus Mechanistic Models

Theoretical approaches to neurophysiological phenomena in general, and to synaptic plasticity in particular, can be roughly grouped into three categories: first, biophysical models that aim at explaining the STDP function from principles of ion channel dynamics and intracellular processes (Senn et al. 2001; Shouval et al. 2002; Abarbanel et al. 2002; Karmarkar and Buonomano 2002); second, mathematical models that start from a given STDP function and analyze computational principles such as intrinsic normalization of summed efficacies or sensitivity to correlations in the input (Kempster et al. 1999; Roberts 1999; Roberts and Bell 2000; van Rossum et al. 2000; Kistler and van Hemmen 2000; Song et al. 2000; Song and Abbott 2001; Kempster et al. 2001; Gütig et al. 2003); finally, models that derive ‘optimal’ STDP properties for a given computational task (Chechik 2003; Dayan and Häusser 2004; Hopfield and Brody 2004; Bohte and Mozer 2005; Bell and Parra 2005; Toyozumi et al. 2005a; Toyozumi et al. 2005b). Optimizing the likelihood of postsynaptic firing in a predefined interval, as we did in this paper, is only one possibility amongst others of introducing concepts of optimality (Barlow 1961; Atick and Redlich 1990; Bell and Sejnowski 1995) into the field of STDP. Chechik (2003) uses concepts from information theory, but restricts his study to the classification of stationary patterns. The paradigm considered in Bohte and Mozer (2005) is similar to our scenario B_c, in that they use a fairly strong teaching input to make the postsynaptic neuron fire. Bell and Parra (2005) and Toyozumi et al. (2005a) are also using concepts from information theory, but they are applying them to the pre- and postsynaptic spike trains. The work of Toyozumi et al. (2005a) is a clearcut *unsupervised* learning paradigm and hence distinct from the present approach. Dayan and Häusser (2004) use concepts of optimal filter theory, but are not interested in precise firing of the postsynaptic neuron. The work of Hopfield and Brody (2004) is similar to our approach in that it focuses on recognition of temporal input patterns, but in our approach we are in addition interested in triggering postsynaptic firing with precise timing. Hopfield and Brody emphasize the repair of disrupted synapses in a network that has previously acquired its function of temporal pattern detector.

Optimality approaches, such as ours, will never be able to make strict predictions about the properties of neurons or synapses. Optimality criteria may, however, help to elucidate computational principles and provide insights about potential tasks of electrophysiological phenomena such as STDP.

Acknowledgments

This work was supported by the Swiss National Science Foundation (200020-103530/1 and 200020-108093/1). T.T was supported by the Research Fellowships of the Japan Society for the Promotion of Science for Young Scientists and a Grant-in-Aid for JSPS Fellows.

A Probability Density of a Spike Train

The probability density of generating a spike train $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ with the stochastic process defined by Eq. (5) can be expressed as follows:

$$P(y_t) = P(t_i^1, \dots, t_i^F)R(t|y_t) \quad (32)$$

where $P(t_i^1, \dots, t_i^F)$ is the probability density of having F spikes at times t_i^1, \dots, t_i^F and $R(t|y_t) = \exp\left(-\int_{t_i^F}^t \rho(t'|y_{t'})dt'\right)$ corresponds to the probability of having no spikes from t_i^F to t . Since the joint probability $P(t_i^1, \dots, t_i^F)$ can be expressed as a product of conditional probabilities

$$P(t_i^1, \dots, t_i^F) = P(t_i^1) \prod_{f=2}^F P(t_i^f | t_i^{f-1}, \dots, t_i^1) \quad (33)$$

Eq. (32) becomes

$$\begin{aligned} P(y_t) &= \rho(t_i^1 | y_{t_i^1}) \exp\left(-\int_0^{t_i^1} \rho(t'|y_{t'})dt'\right) \\ &\cdot \left\{ \prod_{f=2}^F \rho(t_i^f | y_{t_i^f}) \exp\left(-\int_{t_i^{f-1}}^{t_i^f} \rho(t'|y_{t'})dt'\right) \right\} \exp\left(-\int_{t_i^F}^t \rho(t'|y_{t'})dt'\right) \\ &= \left(\prod_{t_i^f \in y_t} \rho(t_i^f | y_{t_i^f}) \right) \exp\left(-\int_0^t \rho(t'|y_{t'})dt'\right) \end{aligned} \quad (34)$$

B Numerical Evaluation of $\bar{\rho}(t)$

Since it is impossible to numerically evaluate the instantaneous firing rate $\bar{\rho}(t)$ with the analytical expression given by Eq. (14), we have to do it in a different way. In fact, there are two different ways to evaluate $\bar{\rho}(t)$. Before going into the details, let us first recall that from the law of large numbers, the instantaneous firing rate is equal to the empirical density of spikes at time t :

$$\langle \rho(t|y_t) \rangle_{y_t} = \langle Y(t) \rangle_{Y(t)} \quad (35)$$

where $Y(t) = \sum_{t_i^f \in y_t} \delta(t - t_i^f)$ is one realization of the postsynaptic spike train. Thus the first and simplest method based on the r.h.s of Eq. (35) is to build a PSTH by counting spikes in small time bins $[t, t + \delta t]$ over, say $K = 10'000$ repetitions of an experiment. The second, and more advanced method, consists in evaluating the l.h.s. of Eq. (35) by Monte-Carlo sampling: instead of averaging over all possible spike trains y_t , we generate $K = 10'000$ spike trains by repetition of the same stimulus. A specific spike train $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ will automatically appear with appropriate probability given by Eq. (6). The Monte-Carlo estimation $\tilde{\rho}(t)$ of $\bar{\rho}(t)$ can be written as

$$\tilde{\rho}(t) = \frac{1}{P} \sum_{m=1}^P \rho(t|y_t^m) \quad (36)$$

where y_t^m is the m^{th} spike train generated by the stochastic process given by Eq. (5). Since we use the analytical expression of $\rho(t|y_t^m)$, we will call Eq. (36) a semi-analytical estimation. Let us note that the semi-analytical estimation $\tilde{\rho}(t)$ converges more rapidly to the true value $\bar{\rho}(t)$ than the empirical estimation based on the PSTH.

In the limit of a Poisson process, i.e. $\eta_0 = 0$, the semi-analytical estimation $\tilde{\rho}(t)$ given by Eq. (36) is equal to the analytical expression of Eq. (14), since the instantaneous firing rate ρ of a Poisson process is independent of the firing history $y_t = \{t_i^1, t_i^2, \dots, t_i^F < t\}$ of the postsynaptic neuron.

C Deconvolution

C.1 Deconvolution for Spike Pairs

With a learning rule such as Eq. (24), we know the optimal weight change Δw_{ij} for each synapse, but we still do not know the corresponding STDP function.

Let us first define the correlation function $c_k(\tau)$, $k = N(i-1) + j$ between the presynaptic spike train $X_j^i(t) = \sum_{t^{\text{pre}} \in x_j^i} \delta(t - t^{\text{pre}})$ and the postsynaptic spike train $Y^i(t) = \sum_{t^{\text{des}} \in y^i} \delta(t - t^{\text{des}})$:

$$c_k(\tau) = \int_0^T X_j^i(s) Y^i(s + \tau) ds, \quad k = 1, \dots, NM \quad (37)$$

where we allow a range $-T_0 \leq \tau \leq T_0$, with $T_0 \ll T$. Since the sum of the pair based weight change ΔW should be equal to the total adaptation of weights Δw_k , we can write:

$$\int_{-T_0}^{T_0} c_k(s) \Delta W(s) ds \stackrel{!}{=} \Delta w_k \quad k = 1, \dots, NM \quad (38)$$

If we want to express Eq. (37) in a matrix form, we need to discretize time in small bins δt and define the matrix element

$$C_{k\Delta} = \int_{\Delta\delta t - T_0}^{(\Delta+1)\delta t - T_0} c_k(s) ds \quad (39)$$

Now Eq. (38) becomes

$$\Delta \mathbf{w} \stackrel{!}{=} C \Delta \mathbf{W} \quad (40)$$

where $\Delta \mathbf{w} = (\Delta w_{11}, \dots, \Delta w_{1N}, \Delta w_{21}, \dots, \Delta w_{MN})^T$ is the vector containing all the optimal weight change and $\Delta \mathbf{W}$ is the vector containing the discretized STDP function, i.e. $\Delta W_\Delta = \Delta W(-T_0 + \Delta\delta t)$, for $\Delta = 1, \dots, 2\tilde{T}_0$ with $\tilde{T}_0 = T_0/\delta t$.

In order to solve the last matrix equation, we have to compute the inverse of the non-square $NM \times 2\tilde{T}_0$ matrix C , which is known as the Moore-Penrose inverse (or the pseudo-inverse):

$$C^+ = (C^T C)^{-1} C^T \quad (41)$$

which exists only if $(C^T C)^{-1}$ exists. In fact, the solution given by

$$\Delta \mathbf{W} = C^+ \Delta \mathbf{w} \quad (42)$$

minimizes the square distance

$$D = \frac{1}{2} (C \Delta \mathbf{W} - \Delta \mathbf{w})^2 \quad (43)$$

C.2 Temporal Locality Constraint

If we want to impose a constraint of locality, we can add a term in the minimization process of Eq. (43) and define the following:

$$E = D + \frac{1}{2} \Delta \mathbf{W}^T P \Delta \mathbf{W} \quad (44)$$

where P is a diagonal matrix which penalizes non-local terms. In this paper, we take a quadratic suppression of terms that are non-local in time:

$$P_{\Delta\Delta'} = a \delta_{\Delta\Delta'} \left(\Delta - \tilde{T}_0 \right)^2 \quad (45)$$

\tilde{T}_0 corresponds to the index of the vector $\Delta \mathbf{W}$ in Eqs. (40) and (44) for which $t^{\text{pre}} - t^{\text{des}} = 0$. Calculating the gradient of E given by Eq. (44) with respect to $\Delta \mathbf{W}$ yields

$$\nabla_{\Delta \mathbf{W}} E = C^T (C \Delta \mathbf{W} - \Delta \mathbf{w}) + P \Delta \mathbf{W} \quad (46)$$

By looking at the minimal value of E , i.e. $\nabla_{\Delta \mathbf{W}} E = 0$, we have

$$\Delta \mathbf{W} = (C^T C + P)^{-1} C^T \Delta \mathbf{w} \quad (47)$$

By setting $a = 0$, we recover the previous case.

References

- Abarbanel, H., R. Huerta, and M. Rabinovich (2002). Dynamical model of long-term synaptic plasticity. *Proc. Natl. Academy of Sci. USA* 59, 10137–10143.
- Atick, J. and A. Redlich (1990). Towards a theory of early visual processing. *Neural Computation* 4, 559–572.
- Barber, D. (2003). Learning in spiking neural assemblies. In *Advances in Neural Information Processing Systems 15*, pp. 149–156. MIT Press.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith (Ed.), *Sensory Communication*, pp. 217–234. MIT Press.
- Bell, A. and T. Sejnowski (1995). An information maximization approach to blind separation and blind deconvolution. *Neural Computation* 7, 1129–1159.
- Bell, A. J. and L. C. Parra (2005). Maximising sensitivity in a spiking network. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 121–128. Cambridge, MA: MIT Press.
- Bell, C., V. Han, Y. Sugawara, and K. Grant (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387, 278–281.
- Bi, G. and M. Poo (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bi, G. and M. Poo (1999). Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature* 401, 792–796.

- Bi, G. and M. Poo (2001). Synaptic modification of correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci.* *24*, 139–166.
- Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Oxford: Clarendon Press.
- Bohte, S. M. and M. C. Mozer (2005). Reducing spike train variability: A computational theory of spike-timing dependent plasticity. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 201–208. Cambridge, MA: MIT Press.
- Brody, C. and J. Hopfield (2003). Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron* *37*, 843–852.
- Bugmann, G., C. Christodoulou, and J. G. Taylor (1997). Role of temporal integration and fluctuation detection in the highly irregular firing of leaky integrator neuron model with partial reset. *Neural Computation* *9*, 985–1000.
- Carr, C. E. and M. Konishi (1990). A circuit for detection of interaural time differences in the brain stem of the barn owl. *J. Neurosci.* *10*, 3227–3246.
- Chechik, G. (2003). Spike-timing-dependent plasticity and relevant mutual information maximization. *Neural Computation* *15*, 1481–1510.
- Dayan, P. and M. Häusser (2004). Plasticity kernels and temporal statistics. In *Advances in Neural Information Processing Systems 16*. MIT Press.
- Feldman, D. (2000). Timing-based LTP and LTD and vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* *27*, 45–56.
- Frégnac, Y., D. E. Shulz, S. Thorpe, and E. Bienenstock (1988). A cellular analogue of visual cortical plasticity. *Nature* *333(6171)*, 367–370.
- Frégnac, Y., D. E. Shulz, S. Thorpe, and E. Bienenstock (1992). Cellular analogs of visual cortical epigenesis. I plasticity of orientation selectivity. *Journal of Neuroscience* *12(4)*, 1280–1300.
- Gerstner, W. (2001). Coding properties of spiking neurons: reverse- and cross-correlations. *Neural Networks* *14*, 599–610.
- Gerstner, W., R. Kempter, J. L. van Hemmen, and H. Wagner (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* *383*, 76–78.
- Gerstner, W. and W. K. Kistler (2002a). Mathematical formulations of Hebbian learning. *Biological Cybernetics* *87*, 404–415.
- Gerstner, W. and W. K. Kistler (2002b). *Spiking Neuron Models*. Cambridge University Press.
- Gerstner, W., R. Ritz, and J. L. van Hemmen (1993). Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.* *69*, 503–515.
- Gütig, R., R. Aharonov, S. Rotter, and H. Sompolinsky (2003). Learning input correlations through non-linear temporally asymmetric hebbian plasticity. *J. Neuroscience* *23*, 3697–3714.
- Haykin, S. (1994). *Neural Networks*. Upper Saddle River, NJ: Prentice Hall.

- Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376, 33–36.
- Hopfield, J. J. and C. D. Brody (2004). Learning rules and network repair in spike-timing-based computation networks. *Proc. Natl. Acad. Sci. USA* 101, 337–342.
- Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks* 14, 1569–1572.
- Johansson, R. and I. Birznieks (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience* 7, 170–177.
- Karmarkar, U. and D. Buonomano (2002). A model of spike-timing dependent plasticity: one or two coincidence detectors. *J. Neurophysiology* 88, 507–513.
- Kempster, R., W. Gerstner, and J. L. van Hemmen (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514.
- Kempster, R., W. Gerstner, and J. L. van Hemmen (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Computation* 13, 2709–2741.
- Kistler, W. M. and J. L. van Hemmen (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput.* 12, 385–405.
- Legenstein, R., C. Naeger, and W. Maass (2005). What can a neuron learn with spike-timing-dependent plasticity? *Neural Computation* 17, 2337–2382.
- Magee, J. C. and D. Johnston (1997). A synaptically controlled associative signal for hebbian plasticity in hippocampal neurons. *Science* 275, 209–213.
- Markram, H., J. Lübke, M. Frotscher, and B. Sakmann (1997). Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSP. *Science* 275, 213–215.
- Mehta, M. R., A. K. Lee, and M. A. Wilson (2002). Role of experience of oscillations in transforming a rate code into a temporal code. *Nature* 417, 741–746.
- Minsky, M. L. and S. A. Papert (1969). *Perceptrons*. Cambridge Mass.: MIT Press.
- Panzeri, S., R. Peterson, S. Schultz, M. Lebedev, and M. Diamond (2001). The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron* 29, 769–777.
- Poliakov, A. V., R. K. Powers, and M. C. Binder (1997). Functional identification of the input-output transforms of motoneurons in the rat and cat. *J. Physiology* 504, 401–424.
- Rao, R. P. N. and T. J. Sejnowski (2001). Spike-timing dependent Hebbian plasticity as temporal difference learning. *Neural Computation* 13, 2221–2237.
- Roberts, P. (1999). Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. *J. Computational Neuroscience* 7, 235–246.

- Roberts, P. and C. Bell (2000). Computational consequences of temporally asymmetric learning rules: II. Sensory image cancellation. *Computational Neuroscience* 9, 67–83.
- Rubin, J., D. D. Lee, and H. Sompolinsky (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Physical Review Letters* 86, 364–367.
- Senn, W., M. Tsodyks, and H. Markram (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation* 13, 35–67.
- Seung, S. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* 40, 1063–1073.
- Shouval, H. Z., M. F. Bear, and L. N. Cooper (2002). A unified model of nmda receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA* 99, 10831–10836.
- Sjöström, P., G. Turrigiano, and S. Nelson (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164.
- Song, S. and L. Abbott (2001). Column and map development and cortical remapping through spike-timing dependent plasticity. *Neuron* 32, 339–350.
- Song, S., K. Miller, and L. Abbott (2000). Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience* 3, 919–926.
- Thorpe, S., A. Delorme, and R. Van Rullen (2001). Spike-based strategies for rapid processing. *Neural Networks* 14, 715–725.
- Toyoizumi, T., J.-P. Pfister, K. Aihara, and W. Gerstner (2005b). Generalized bienenstock-cooper-munro rule for spiking neurons that maximizes information transmission. *Proc. National Academy Sciences (USA)* 102, 5239–5244.
- Toyoizumi, T., J.-P. Pfister, K. Aihara, and W. Gerstner (2005a). Spike-timing dependent plasticity and mutual information maximization for a spiking neuron model. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 1409–1416. Cambridge, MA: MIT Press.
- Troyer, T. W. and K. Miller (1997). Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation* 9, 971–983.
- Turrigiano, G. and S. Nelson (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience* 5, 97–107.
- van Rossum, M. C. W., G. Q. Bi, and G. G. Turrigiano (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience* 20, 8812–8821.
- Xie, X. and S. Seung (2004). Learning in neural networks by reinforcement of irregular spiking. *Phys. Rev. E* 69, 041909.
- Zhang, L., H. Tao, C. Holt, W.A.Harris, and M.-M. Poo (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395, 37–44.

2.3 Paper III

Summary

Like the two first papers of this thesis, this third paper also aims to develop an optimal model of STDP. The context is however completely different. Indeed, this paper and the two following ones derive the synaptic learning in an unsupervised scenario. Here, we take the mutual information between the input spike train and the output spike train as the objective function.

The particularity of this paper, compared to the two next ones, is that we consider small perturbations around the membrane potential. This assumption allows us to express the mutual information as a function of the Fisher information.

With these assumptions, the potentiation part of the learning window is determined by the $\epsilon(s)$ kernel which reflects the trajectory of an EPSP. More precisely, the causal part of the learning window depends on $\epsilon^2(s)$ as opposed to the results of the first two papers where the LTP part of the learning window is given by $\epsilon(s)$.

The negative part of the STDP function comes essentially from the refractoriness. Indeed, a presynaptic spike has no effect on the postsynaptic neuron if it occurs just after the postsynaptic spike, i.e. during the refractory period. Note that this result is present in all three unsupervised papers (papers III - V). Since refractoriness plays a crucial role in those three papers, we included a strong refractoriness¹.

It should be noted that the offset problem is not absent from this paper. The optimal weight change resulting from mutual information maximization is decomposed into a constant term, a pre- and a postsynaptic term and a correlation term. Those four terms corresponds to the first Volterra expansion terms of Eqs. (1.2.2), (1.2.3) and (1.2.4). Here, the offset is absorbed in the constant term, the pre- and the postsynaptic terms, so that the correlation term can be compared to the experiments.

Curiously, the online learning rule given by Eq. (15) of this paper contains only pre-post

¹In papers III and IV, refractoriness is modeled as a multiplicative factor (the instantaneous firing rate $\rho(t)$ is given by $\rho(t) = g(u)R(t)$, where $R(t)$ models refractoriness) which contains an absolute refractory period as well as a relative refractory period. In paper V, refractoriness is included the EPSP suppression with a relatively long recovery time constant ($\tau_a = 50$ ms). In this way, it is possible to get a reasonably long depression window.

effects, i.e. synaptic weights increase whenever there is a postsynaptic spike at time t . So where does the post-pre depression come from? In short, the LTD part comes from averaging. Indeed, the calculated weight change corresponds to the expected weight change when t^{pre} and t^{post} are given. This means that other pre- or postsynaptic spikes can be generated by the statistics of the pre- and postsynaptic neuron. Therefore, because of refractoriness, the autocorrelation of the postsynaptic spike train is smaller than average for small interspike intervals. This lead to the LTD part of the correlation term.

Reference

This paper has been presented by Taro Toyoizumi at the NIPS 2004 conference (December 13-18, 2004, Vancouver, Canada) and has been published after a peer-review process under the following reference:

Toyoizumi T and Pfister JP and Aihara K and Gerstner W. Spike-Timing Dependent Plasticity and Mutual Information Maximization for a Spiking Neuron Model. *Advances in Neural Information Processing Systems 17*, edited by L.K. Saul and Y. Weiss and L.Bottou, MIT Press, Cambridge MA, 1409-1416, 2005.

This paper has been cited by (Bohte and Mozer 2005).

Spike-Timing Dependent Plasticity and Mutual Information Maximization for a Spiking Neuron Model

Taro Toyozumi^{†‡}, **Jean-Pascal Pfister[‡]**

Kazuyuki Aihara^{§ *}, **Wulfram Gerstner[‡]**

[†] Department of Complexity Science and Engineering,
The University of Tokyo, 153-8505 Tokyo, Japan

[‡] Ecole Polytechnique Fédérale de Lausanne (EPFL),
School of Computer and Communication Sciences and
Brain-Mind Institute, 1015 Lausanne, Switzerland

[§] Graduate School of Information Science and Technology,
The University of Tokyo, 153-8505 Tokyo, Japan

taro@sat.t.u-tokyo.ac.jp, jean-pascal.pfister@epfl.ch
aihara@sat.t.u-tokyo.ac.jp, wulfram.gerstner@epfl.ch

Abstract

We derive an optimal learning rule in the sense of mutual information maximization for a spiking neuron model. Under the assumption of small fluctuations of the input, we find a spike-timing dependent plasticity (STDP) function which depends on the time course of excitatory postsynaptic potentials (EPSPs) and the autocorrelation function of the postsynaptic neuron. We show that the STDP function has both positive and negative phases. The positive phase is related to the shape of the EPSP while the negative phase is controlled by neuronal refractoriness.

1 Introduction

Spike-timing dependent plasticity (STDP) has been intensively studied during the last decade both experimentally and theoretically (for reviews see [1, 2]). STDP is a variant of Hebbian learning that is sensitive not only to the spatial but also to the temporal correlations between pre- and postsynaptic neurons. While the exact time course of the STDP function varies between different types of neurons, the functional consequences of these differences are largely unknown. One line of modeling research takes a given STDP rule and analyzes the evolution of synaptic efficacies [3–5]. In this article, we take a different

*Alternative address: ERATO Aihara Complexity Modeling Project, JST, 45-18 Oyama, Shibuya-ku, 151-0065 Tokyo, Japan

approach and start from first principles. More precisely, we ask what is the optimal synaptic update rule so as to maximize the mutual information between pre- and postsynaptic neurons.

Previously information theoretical approaches to neural coding have been used to quantify the amount of information that a neuron or a neural network is able to encode or transmit [6–8]. In particular, algorithms based on the maximization of the mutual information between the output and the input of a network, also called infomax principle [9], have been used to detect the principal (or independent) components of the input signal, or to reduce the redundancy [10–12]. Although it is a matter of discussion whether neurons simply ‘transmit’ information as opposed to classification or task-specific processing [13], strategies based on information maximization provide a reasonable starting point to construct neuronal networks in an unsupervised, but principled manner.

Recently, using a rate neuron, Chechik applied information maximization to detect static input patterns from the output signal, and derived the optimal temporal learning window; the learning window has a positive part due to the effect of the postsynaptic potential and has flat negative parts with a length determined by the memory span [14].

In this paper, however, we employ a stochastic spiking neuron model to study not only the effect of postsynaptic potentials generated by synaptic input but also the effect of the refractory period of the postsynaptic neuron on the shape of the optimal learning window. We discuss the relation of mutual information and Fisher information for small input variance in Sec. 2. Optimization of the Fisher information by gradient ascent yields an optimal learning rule as shown in Sec. 3

2 Model assumptions

2.1 Neuron model

The model we are considering is a stochastic neuron with refractoriness. The instantaneous firing rate ρ at time t depends on the membrane potential $u(t)$ and refractoriness $R(t)$:

$$\rho(t) = g(\beta u(t))R(t), \quad (1)$$

where $g(\beta u) = g_0 \log_2[1 + e^{\beta u}]$ is a smoothed piecewise linear function with a scaling variable β and a constant $g_0 = 85\text{Hz}$. The refractory variable is $R(t) = \frac{(t - \hat{t} - \tau_{\text{abs}})^2}{\tau_{\text{refr}}^2 + (t - \hat{t} - \tau_{\text{abs}})^2} \Theta(t - \hat{t} - \tau_{\text{abs}})$ and depends on the time elapsed since the last firing time \hat{t} , the absolute refractory period $\tau_{\text{abs}} = 3$ ms, and the time constant of relative refractoriness $\tau_{\text{refr}} = 10$ ms. The Heaviside step function Θ takes a value of 1 for positive arguments and zero otherwise. The postsynaptic potential depends on the input spike trains of N presynaptic neurons. A presynaptic spike of neuron $i \in \{1, 2, \dots, N\}$ emitted at time t_i^f evokes a postsynaptic potential with time course $\epsilon(t - t_i^f)$. The total membrane potential is

$$u(t) = \sum_{i=1}^N w_i \sum_f \epsilon(t - t_i^f) = \sum_{i=1}^N w_i \int \epsilon(s) x_i(t - s) ds \quad (2)$$

where $x_i(t) = \sum_f \delta(t - t_i^f)$ denotes the spike train of the presynaptic neuron i . The above model is a special case of the spike response model with escape noise [2]. For vanishing refractoriness $\tau_{\text{refr}} \rightarrow 0$ and $\tau_{\text{abs}} \rightarrow 0$, the above model reduces to an inhomogeneous Poisson process.

For a given set of presynaptic spikes in an interval $[0, T]$, hence for a given time course of

membrane potential $\{u(t)|t \in [0, T]\}$, the model generates an output spike train

$$y(t) = \sum_f \delta(t - t^f) \quad (3)$$

with firing times $\{t^f|f = 1, \dots, n\}$ with a probability density

$$P(y|u) = \exp \left[\int_0^T (y(t) \log \rho(t) - \rho(t)) dt \right]. \quad (4)$$

where $\rho(t)$ is given by Eq. (1), i.e., $\rho(t) = g(\beta u(t)) R(t)$. Since the refractory variable R depends on the firing time \hat{t} of the *previous* output spike, we sometimes write $\rho(t|\hat{t})$ instead of $\rho(t)$ in order to make this dependence explicit. Equation (4) can then be re-expressed in terms of the survivor function $S(t|\hat{t}) = e^{-\int_{\hat{t}}^t \rho(s|\hat{t}) ds}$ and the interval distribution $Q(t|\hat{t}) = \rho(t|\hat{t})S(t|\hat{t})$ in a more transparent form:

$$P(y|u) = \left(\prod_{f=1}^n Q(t^f|t^{f-1}) \right) S(T|t^n), \quad (5)$$

where $t^0 = 0$ and n is the number of postsynaptic spikes in $[0, T]$. In words, the probability that a specific output spike train y occurs can be calculated from the interspike intervals $Q(t^f|t^{f-1})$ and the probability that the neuron ‘survives’ from the last spike at time t^n to time T without further firing.

2.2 Fisher information and mutual information

Let us consider input spike trains with stationary statistics. These input spike trains generate an input potential $u(t)$ with an average value u_0 and standard deviation σ . Assuming a weak dependence of g on the membrane potential u , i.e., for small β , we expand g around $g_0 = g(0)$ to obtain $g(\beta u(t)) = g_0 + g'_0 \beta u(t) + g''_0 [\beta u(t)]^2/2 + O(\beta^3)$ where g_0 is the value of g in the absence of input and the next terms describe the influence of the input. Here and in the following, all calculations will be done to order β^2 .

In the limit of small β , the mutual information is given by [15]

$$I(Y; X) = \frac{\beta^2}{2} \int_0^T dt \int_0^T dt' \Sigma(t-t') J_0(t-t') + O(\beta^3), \quad (6)$$

with the autocovariance function of the membrane potential

$$\Sigma(t-t') = \langle \Delta u(t) \Delta u(t') \rangle_X, \quad (7)$$

with $\Delta u(t) = u(t) - u_0$ and Fisher information

$$J_0(t-t') = - \left\langle \frac{\partial^2 \log P(y|u)}{\partial \beta u(t) \partial \beta u(t')} \Big|_{\beta=0} \right\rangle_{Y|\beta=0}, \quad (8)$$

with $\langle \cdot \rangle_{Y|\beta=0} = \int \cdot P(y|\beta = 0) dy$ and $\langle \cdot \rangle_X = \int \cdot P(x) dx$. Note that the Fisher information (8) is to be evaluated at the constant g_0 , i.e., at the value $\beta u = 0$, whereas the autocovariance in Eq. (7) is calculated with respect to the *mean* membrane potential $u_0 = \langle u(t) \rangle_X$ which is in general different from zero. The derivation of (6) is based on the assumption that the variability of the output signal is small and $g(\beta u)$ does not deviate much from g_0 , i.e., it corresponds to the regime of small signal-to-noise ratio. It is well known that the information capacity of the Gaussian channel is given by the log of the signal-to-noise ratio [16], and the mutual information is proportional to the

signal-to-noise ratio when it is small. The relation between the Fisher information, the mutual information, and optimal tuning curves has previously been established in the regime of large signal-to-noise ratio [17].

We introduce the following notation: Let $\mu_0 = \langle y(t) \rangle_{Y|\beta=0} = \langle \rho(t) \rangle_{Y|\beta=0}$ be the spontaneous firing rate in the absence of input and $\mu_0^{-1} \langle y(t)y(t') \rangle_{Y|\beta=0} = \delta(t-t') + \mu_0[1 + \phi(t-t')]$ be the postsynaptic firing probability at time t given a postsynaptic spike at t' , i.e., the autocorrelation function of Y . From the theory of stationary renewal processes [2]

$$\begin{aligned} \mu_0 &= \left[\int s Q_0(s) ds \right]^{-1}, \\ \mu_0[1 + \phi(s)] &= Q_0(|s|) + \int Q_0(s') \mu_0[1 + \phi(|s| - s')] \Theta(|s| - s') ds', \end{aligned} \quad (9)$$

where $Q_0(s) = g_0 R(s) e^{-g_0[(s-\tau_{\text{abs}}) - \tau_{\text{refr}} \arctan(s-\tau_{\text{abs}})/\tau_{\text{refr}}]}$ is the interval distribution for constant $g = g_0$. The interval distribution vanishes during the absolute refractory time τ_{abs} ; cf. Fig. 1.

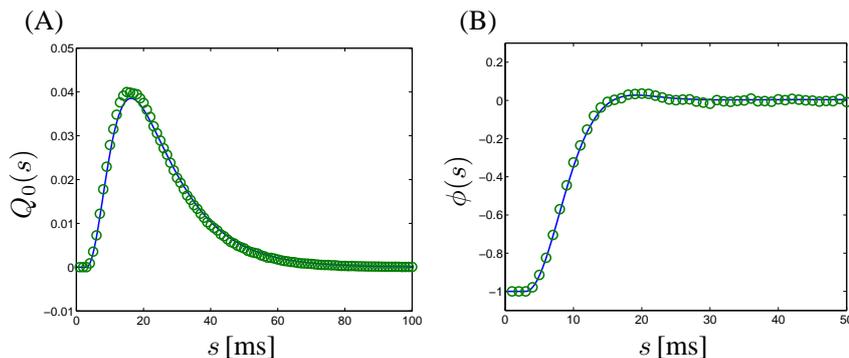


Figure 1: Interspike interval distribution Q_0 and normalized autocorrelation function ϕ . The circles show numerical results, the solid line the theory.

The Fisher information of (8) is calculated from (4) to be

$$J_0(t-t') = \delta(t-t') \left(\frac{g'_0}{g_0} \right)^2 \langle \rho_0(t) \rangle_{Y|\beta=0} \quad (10)$$

with the instantaneous firing rate $\rho_0(t) = g_0 R(t)$. Hence the mutual information is

$$I(Y; X) = \frac{\beta^2}{2} \left(\frac{g'_0}{g_0} \right)^2 \int_0^T dt \mu_0 \sigma^2 \quad (11)$$

$$= \frac{\beta^2}{2} \left(\frac{g'_0}{g_0} \right)^2 T \mu_0 \sigma^2. \quad (12)$$

For an interpretation of Eq. (11) we note that $\sigma^2 = \Sigma(0)$ is the variance of the membrane potential and depends on the statistics of the *presynaptic* input whereas μ_0 is the spontaneous firing rate which characterizes the output of the *postsynaptic* neuron. Hence, Equation (11) contains both pre- and postsynaptic factors.

3 Results: Optimal spike-timing dependent learning rule

In the previous section we have calculated the mutual information between presynaptic input spike trains and the output of the postsynaptic neuron under the assumption of small

fluctuations of g . The mutual information depends on parameters of the model neuron, in particular the synaptic weights that characterize the efficacy of the connections between pre- and postsynaptic neurons. In this section, we will optimize the mutual information by changing the synaptic weights in an appropriate fashion. To do so we will proceed in several steps.

First, based on gradient ascent we derive a batch learning rule of synaptic weights that maximizes the mutual information. In a second step, we transform the batch rule into an online rule that reduces to the batch version when averaged. Finally, in subsection 3.2, we will see that the online learning rule shares properties with STDP, in particular a biphasic dependence upon the relative timing of pre- and postsynaptic spikes.

3.1 Learning rule for spiking model neuron

In order to keep the analysis as simple as possible, we suppose that the input spike trains are independent Poisson trains, i.e., $\langle \Delta x_i(t) \Delta x_j(t') \rangle_X = \nu_i \delta(t - t') \delta_{ij}$, where $\Delta x_i(t) = x_i(t) - \nu_i$ with rate $\nu_i = \langle x_i(t) \rangle_X$. Then we obtain the variance of the membrane potential

$$\sigma^2 = \langle [\Delta u(t)]^2 \rangle_X = \epsilon_2 \sum_j w_j^2 \nu_j \quad (13)$$

with $\epsilon_2 = \int \epsilon^2(s) ds$.

Applying gradient ascent to (11) with an appropriate learning rate α , we obtain the batch learning rule of synaptic weights as

$$\Delta w_i = \alpha \frac{\partial I(Y; X)}{\partial w_i} \approx \alpha \frac{\beta^2}{2} \left(\frac{g'_0}{g_0} \right)^2 \int_0^T dt \mu_0 \frac{\partial \sigma^2}{\partial w_i}. \quad (14)$$

The derivative of μ_0 with respect to w_i vanishes, since μ_0 is the spontaneous firing rate in the absence of input. We note that both μ_0 and σ^2 are defined by an ensemble averages, as is typical for a ‘batch’ rule.

While there are many candidates of online learning rule that give (14) on average, we are interested in rules that depend directly on neuronal spikes rather than mean rates. To proceed it is useful to write $\sigma^2 = \langle [\Delta u(t)]^2 \rangle_X$ with $\Delta u = \sum_i w_i \Delta \epsilon_i(t)$ where $\Delta \epsilon_i(t) = \int \epsilon(s) \Delta x_i(t - s) ds$. In this notation, one simple form of an online learning rule that depends on both the postsynaptic firing statistics and presynaptic autocorrelation is

$$\frac{dw_i}{dt} = \alpha \beta^2 \left(\frac{g'_0}{g_0} \right)^2 y(t) \Delta \epsilon_i(t) \Delta u(t), \quad (15)$$

Hence weights are updated with each postsynaptic spike with an amplitude proportional to an online estimate of the membrane potential variance calculated as the product of Δu and $\Delta \epsilon_i$. Indeed, to order β^0 , the input and the output spikes are independent; $\langle y(t) \Delta \epsilon_i(t) \Delta u(t) \rangle_{Y, X} = \langle y(t) \rangle_{Y | \beta=0} \langle \Delta \epsilon_i(t) \Delta u(t) \rangle_X$ and the average of (15) leads back to (14).

3.2 STDP function as a spike-pair effect

Application of the online learning rule (15) during a trial of duration T , yields a total change of the synaptic efficacy which depends on all the presynaptic spikes via the factor $\Delta \epsilon_i$; on the postsynaptic potential via the factor Δu ; and on the postsynaptic spike train $y(t)$. In order to extract the spike pair effect evoked by a given presynaptic spike at t_i^{pre} and a postsynaptic spike at t^{post} , we average over x and y given the pair of spikes. The spike pair effect up to the second order of β is therefore described as

$$\Delta w_i(t^{post} - t_i^{pre}) = \alpha \beta^2 \left(\frac{g'_0}{g_0} \right)^2 \int_0^T dt \langle y(t) \rangle_{Y | t^{post}, \beta=0} \langle \Delta \epsilon_i(t) \Delta u(t) \rangle_{X | t_i^{pre}}, \quad (16)$$

where $\langle \cdot \rangle_{Y|t^{post}, \beta=0} = \int dy \cdot P(y|t^{post}, \beta = 0)$ and $\langle \cdot \rangle_{X|t_i^{pre}} = \int dx \cdot P(x|t_i^{pre})$. Note that the leading factor of Eq. (16) is already of order β^2 , so that all other factors have to be evaluated to order β^0 . Suppressing all terms containing β , we obtain $P(y|t^{post}, u) \approx P(y|t^{post}, \beta = 0)$ and from the Bayes formula $P(x|t_i^{pre}, t^{post}) = \frac{P(t^{post}|x, t_i^{pre})}{\langle P(t^{post}|x, t_i^{pre}) \rangle_{X|t_i^{pre}}} P(x|t_i^{pre}) \approx P(x|t_i^{pre})$.

In order to see the contribution of t_i^{pre} and t^{post} , we think of separating the effects caused by spikes at t_i^{pre}, t^{post} from the mean weight evolution caused by all other spikes. Therefore we insert $\langle y(t) \rangle_{Y|t^{post}, \beta=0} = \delta(t - t^{post}) + \mu_0 [1 + \phi(t - t^{post})]$ and $\langle \Delta \epsilon_i(t) \Delta u(t) \rangle_{X|t_i^{pre}} = w_i [\epsilon^2(t - t^{pre}) + \epsilon_2 \nu_i]$ into Eq. (16) and decompose $\Delta w_i(t^{post} - t_i^{pre})$ into the following four terms: the drift term $\Delta w_i^0 = \alpha \beta^2 \left(\frac{g'_0}{g_0}\right)^2 T \mu_0 \epsilon_2 w_i \nu_i$ of the batch learning (14) that does not depend on t_i^{pre} or t^{post} ; the presynaptic component $\Delta w_i^{pre} = \alpha \beta^2 \left(\frac{g'_0}{g_0}\right)^2 \mu_0 \epsilon_2 w_i$ that is triggered by the presynaptic spike at t_i^{pre} ; the postsynaptic component $\Delta w_i^{post} = \alpha \beta^2 \left(\frac{g'_0}{g_0}\right)^2 \left[1 + \mu_0 \int_0^T \phi(t - t^{post}) dt\right] \epsilon_2 w_i \nu_i$ that is triggered by the postsynaptic spike at t^{post} ; and the correlation component

$$\Delta w_i^{corr} = \alpha \beta^2 \left(\frac{g'_0}{g_0}\right)^2 w_i \left[\epsilon^2(t^{post} - t_i^{pre}) + \mu_0 \int_0^T \phi(t - t^{post}) \epsilon^2(t - t_i^{pre}) dt \right] \quad (17)$$

that depends on the difference of the pre- and postsynaptic spike timing.

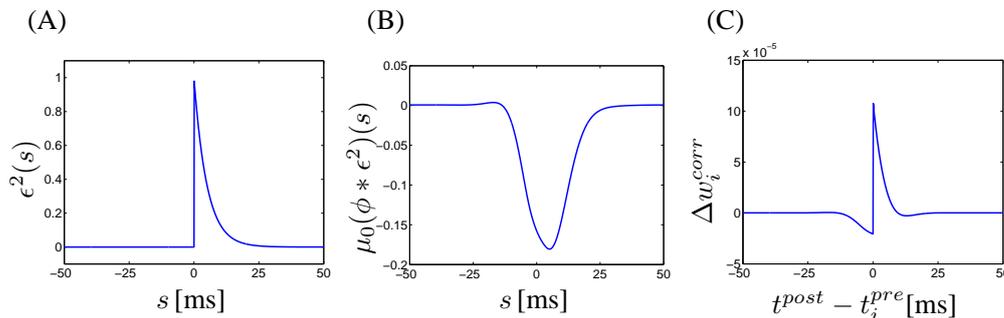


Figure 2: (A) The effect from EPSP: the first term in the square bracket of (17). (B) The effect from refractoriness: the second term in the square bracket of (17). (C) Temporal learning window Δw_i^{corr} of (17).

In the following, we choose a simple exponential EPSP $\epsilon(t) = \Theta(s)e^{-s/\tau_u}$ with a time constant $\tau_u = 10$ ms. The parameters are $N = 100$, $\nu_i = 40$ Hz for all i , $w_i = (N\tau_u\nu_i)^{-1}$, $\alpha = 1$ and $\beta = 0.1$.

Figure 2 shows Δw_i^{corr} of (17). The first term of (17) indicates the contribution of a presynaptic spike at t_i^{pre} to increase the online estimation of membrane potential variance at time t^{post} , whereas the second term represents the effect of the refractory period on postsynaptic firing intensity, i.e., the normalized autocorrelation function convolved with the presynaptic contribution term. Due to the averaging of $\langle \cdot \rangle_{Y|t^{post}, \beta=0}$ and $\langle \cdot \rangle_{X|t_i^{pre}}$ in (16), this optimal temporal learning window is local in time; we do not need to impose a memory span [14] to restrict the negative part of the learning window.

Figure 3 compares Δw_i of (16) with numerical simulations of (15). We note a good agreement between theory and simulation. We recall, that all calculations, and hence the STDP function of (17) are valid for small β , i.e., for small fluctuation of g .

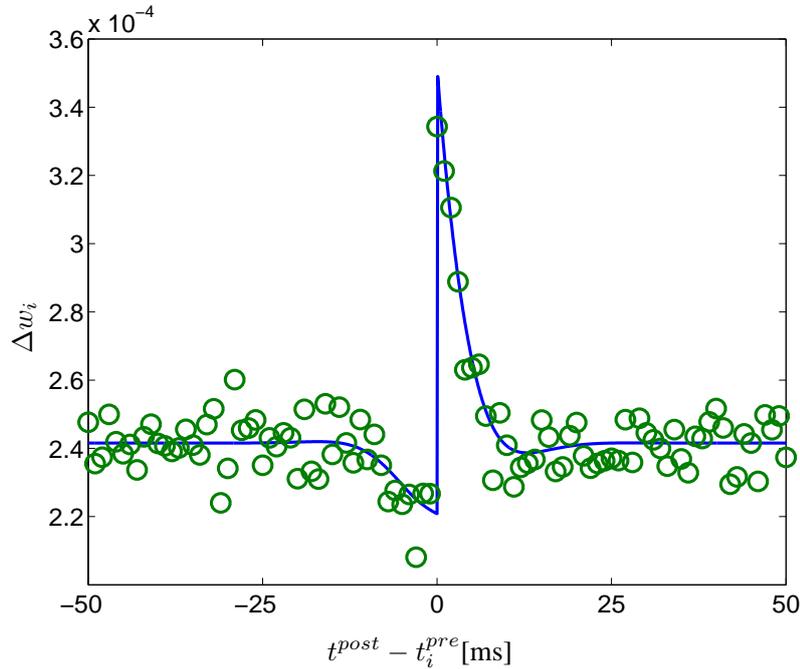


Figure 3: The comparison of the analytical result of (16) (solid line) and the numerical simulation of the online learning rule (15) (circles). For the simulation, the conditional average $\langle \Delta w_i \rangle_{X,Y|t_i^{pre}, t^{post}}$ is evaluated by integrating $\frac{dw_i}{dt}$ over 200 ms around spike pairs with the given interval $t^{post} - t_i^{pre}$;

4 Conclusion

It is important for neurons especially in primary sensory systems to send information from previous processing circuits to neurons in other areas while capturing the essential features of its input. Mutual information is a natural quantity to be maximized from this perspective. We introduced an online learning rule for synaptic weights that increases information transmission for small input fluctuation. Introduction of the temporal properties of the target neuron enables us to analyze the temporal properties of the learning rule required to increase the mutual information. Consequently, the temporal learning window is given in terms of the time course of EPSPs and the autocorrelation function of the postsynaptic neuron. In particular, neuronal refractoriness plays a major role and yields the negative part of the learning window. Though we restrict our analysis here to excitatory synapses with independent spike trains, it is straightforward to generalize the approach to a mixture of excitatory and inhibitory neurons with weakly correlated spike trains as long as the synaptic weights are small enough. The analytically derived temporal learning window is similar to the experimentally observed bimodal STDP window [1]. Since the effective time course of EPSPs and the autocorrelation function of output spike trains vary from one part of the brain to another, it is important to compare those functions with the temporal learning window in biological settings.

Acknowledgments

T.T. is supported by the Japan Society for the Promotion of Science and a Grant-in-Aid for JSPS Fellows; J.-P.P. is supported by the Swiss National Science Foundation. We thank Y. Aviel for discussions.

References

- [1] G. Bi and M. Poo. Synaptic modification of correlated activity: Hebb's postulate revisited. *Annu. Rev. Neurosci.*, 24:139–166, 2001.
- [2] W. Gerstner and W. M. Kistler. *Spiking Neuron Models*. Cambridge University Press, 2002.
- [3] R. Kempster, W. Gerstner, and J. L. van Hemmen. Hebbian learning and spiking neurons. *Phys. Rev. E*, 59:4498–4514, 1999.
- [4] W. Gerstner and W. M. Kistler. Mathematical formulations of hebbian learning. *Biol. Cybern.*, 87:404–415, 2002.
- [5] R. Gütiğ, R. Aharonov, S. Rotter, and H. Sompolinsky. Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *J. Neurosci.*, 23(9):3697–3714, 2003.
- [6] R. B. Stein. The information capacity of nerve cells using a frequency code. *Biophys. J.*, 7:797–826, 1967.
- [7] W. Bialek, F. Rieke, R. de Ruyter van Steveninck, and D. Warland. Reading a neural code. *Science*, 252:1854–1857, 1991.
- [8] F. Rieke, D. Warland, R. R. van Steveninck, and W. Bialek. *Spikes*. MIT Press, 1997.
- [9] R. Linsker. Self-organization in a perceptual network. *Computer*, 21:105–117, 1988.
- [10] J-P. Nadal and N. Parga. Nonlinear neurons in the low-noise limit: a factorial code maximizes information transfer. *Network: Comput. Neural Syst.*, 5:565–581, 1994.
- [11] J-P. Nadal, N. Brunel, and N. Parga. Nonlinear feedforward networks with stochastic outputs: infomax implies redundancy reduction. *Network: Comput. Neural Syst.*, 9:207–217, 1998.
- [12] A. J. Bell and T. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.*, 7(6):1004–1034, 1995.
- [13] J. J. Hopfield. Encoding for computation: recognizing brief dynamical patterns by exploiting effects of weak rhythms on action-potential timing. *Proc. Natl. Acad. Sci. USA*, 101(16):6255–6260, 2004.
- [14] G. Checkik. Spike-timing-dependent plasticity and relevant mutual information maximization. *Neural Comput.*, 15:1481–1510, 2003.
- [15] V. V. Prelov and E. C. van der Meulen. An asymptotic expression for the information and capacity of a multidimensional channel with weak input signals. *IEEE. Trans. Inform. Theory*, 39(5):1728–1735, 1993.
- [16] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. New York: Wiley, 1991.
- [17] N. Brunel and J-P. Nadal. Mutual information, fisher information, and population coding. *Neural Comput.*, 10:1731–1757, 1998.

2.4 Paper IV

Summary

This fourth paper has a similar approach to the last paper in the sense that it considers the weight dynamics resulting from the maximization of the mutual information. However, it differs from the previous one because it considers an additional term in the objective function. This cost term, or penalty term, implements an homeostatic constraint which penalizes large deviations from a fixed firing rate.

Clearly synapses can transmit more information if they increase their strength and thereby increase the postsynaptic firing rate, but this is costly from the point of view of biophysical implementation. It should be also noted that in the previous paper, where we only had the mutual information term, there was no overall depression. Indeed, depression in the acausal part of the learning window was visible only after subtracting the positive offset. For those reasons, we have added an extra homeostatic constraint.

One of the interesting outcomes of this study is that the optimal weight change that results from the maximization of the mutual information and the cost term lead the BCM learning rule if Poisson statistics is assumed for the pre- and postsynaptic neuron. In this case, the sliding threshold necessary for the BCM learning rule comes directly from the homeostatic constraint. Moreover, this sliding threshold has to depend non-linearly on the averaged postsynaptic firing rate. This is precisely the case whenever the homeostatic term is present (i.e. when $\gamma > 0$, c.f. Eq. (17) of paper IV).

The optimal learning rule derived in this paper does not have only rate-based properties, like the input selectivity of the BCM learning rule. It shows also interesting spike-based features. Indeed, if a first group of presynaptic neurons is spike-spike correlated and not the second group, synapses of the first group will be potentiated (see Fig. 5 of paper IV). This interesting property is a feature that the pure BCM learning rule cannot elicit.

Reference

Some parts of this paper have been presented by Taro Toyozumi at the *Cosyne conference* (March 17 - 22, 2005, Salt-Lake City, USA).

This paper has been published after a peer-review process in the *Proceedings of the*

National Academy of Science USA under the following reference:

Toyoizumi T and Pfister JP and Aihara K and Gerstner W. Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission. *Proceedings of the National Academy of Science USA*. 102, 5239-5244, 2005.

Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission

Taro Toyoizumi ^{*†}, Jean-Pascal Pfister^{*},
Kazuyuki Aihara ^{†‡} and Wulfram Gerstner^{§*}

^{*} Ecole Polytechnique Fédérale de Lausanne (EPFL)
School of Computer and Communication Sciences
and Brain-Mind Institute, 1015 Lausanne, Switzerland

[†] Department of Complexity Science and Engineering
The University of Tokyo,

4-6-1, Komaba, Meguro-ku, 153-8505 Tokyo, Japan

[‡]ERATO Aihara Complexity Modeling Project, JST,
45-18 Oyama, Shibuya-ku, 151-0065 Tokyo, Japan

Preprint version of:
PNAS 102, 5239-5244 (2005)

Abstract

Maximization of information transmission by a spiking neuron model predicts changes of synaptic connections that depend on timing of pre- and postsynaptic spikes as well as on the postsynaptic membrane potential. Under the assumption of Poisson firing statistics, the synaptic update rule exhibits all the features of the Bienenstock-Cooper-Munro rule, in particular regimes of synaptic potentiation and depression separated by a sliding threshold. Moreover, the new learning rule is also applicable to the more realistic case of neuron models with refractoriness and is sensitive to correlations between input spikes, even in the absence of presynaptic rate modulation. The learning rule is found by maximizing the mutual information between presynaptic and postsynaptic spike trains under the constraint that the postsynaptic firing rate stays close to some target firing rate. An interpretation of the synaptic update rule in terms of homeostatic synaptic processes and spike-timing dependent plasticity is discussed.

[§]Corresponding author, Telephone: 41-21-693 6713; FAX 41-21-693-9600; email wulfram.gerstner@epfl.ch

1 Introduction

The efficacy of synaptic connections between neurons in the brain is not fixed but varies depending on the firing frequency of presynaptic neurons [1,2], the membrane potential of the postsynaptic neuron [3], spike timing [4–6] as well as intra-cellular parameters such as the calcium concentration; for a review see [7]. During the last decades, a large number of theoretical concepts and mathematical models have emerged that have helped to understand the functional consequences of synaptic modifications, in particular long-term potentiation and depression (LTP and LTD) during development, learning, and memory; for reviews see [8–10]. Apart from the work of Hebb [11], one of the most influential theoretical concepts has been the model proposed by Bienenstock, Cooper, and Munro (BCM) originally developed to account for cortical organization and receptive field properties during development [12]. The model predicted (a) regimes of both LTD and LTP depending on the state of the postsynaptic neuron and (b) a sliding threshold that separates the two regimes. Both predictions *a* and *b* have subsequently been confirmed experimentally [2,13,14].

In this paper we construct a bridge between the BCM model and a seemingly unconnected line of research in theoretical neuroscience centered around the concept of optimality.

There are indications that several components of neural systems show close to optimal performance [15–17]. Instead of looking at a specific implementation of synaptic changes, defined by a rule such as in the BCM model, we therefore ask: what would be the optimal synaptic update rule so as to guarantee that a spiking neuron transmits as much information as possible? Information theoretic concepts have been used by several researchers, since they allow to compare performance of neural systems with a fundamental theoretical limit [16,17], but ‘optimal’ synaptic update rules have so far been mostly restricted to a pure rate description [18–21]. In the following we apply the concept of information maximization to a spiking neuron model with refractoriness. Mutual information is maximized under the constraint that the postsynaptic firing rate stays as close as possible to the neuron’s typical ‘target’ firing rate stabilized by homeostatic synaptic processes [22]. In the special case of vanishing refractoriness, we find that the optimal update rule has the two BCM properties, i.e., regimes of potentiation and depression separated by a sliding threshold. In contrast to the optimality approach of Intrator and Cooper [23], the sliding threshold follows automatically from our formulation of the optimality problem. Moreover, our extension of the BCM rule to spiking neurons with refractoriness shows that synaptic changes should naturally depend on spike timing, spike frequency, and PSP, in agreement with experimental results.

2 Methods and Models

Spiking Neuron Model. We consider a stochastically spiking neuron model with refractoriness. For simulations, and also for some parts of the theory, it is convenient to formulate the model in discrete time with step size Δt , i.e. $t^k = k \Delta t$.

However, for the ease of interpretation and with respect to a comparison with biological neurons, it is more practical to turn to continuous time by taking $\Delta t \rightarrow 0$. The continuous time limit is indicated in the following formulas by a right arrow (\rightarrow). The postsynaptic neuron receives input at N synapses. A presynaptic spike train at synapse j is described in discrete time as a sequence x_j^k ($k = 1, \dots, K$) of zeros (no spike) and ones (spike). The upper index k denotes time bin k . Thus $x_j^k = 1$ indicates that a presynaptic spike arrived at synapse j at a time $t_j^{(f)}$ with $t^{k-1} \leq t_j^{(f)} < t^k$. Each presynaptic spike evokes a postsynaptic potential (PSP) of amplitude w_j and exponential time course $\epsilon(t - t_j^{(f)})$ with time constant $\tau_m = 10$ ms. The membrane potential at time step t^k is denoted as $u(t^k)$ and calculated as the total postsynaptic potential

$$u(t^k) = u_r + \sum_{j=1}^N \sum_{n=1}^k w_j \epsilon(t^k - t_j^n) x_j^n \rightarrow u_r + \sum_j \sum_f w_j \epsilon(t - t_j^{(f)}) \quad [1]$$

where $u_r = -70$ mV is the resting potential. The probability ρ^k of firing in time step k is a function of the membrane potential u and the refractory state R of the neuron,

$$\rho^k = 1 - \exp[-g(u(t^k)) R(t^k) \Delta t] \approx g(u(t^k)) R(t^k) \Delta t. \quad [2]$$

where Δt is the time step and g is a smooth increasing function of u . Thus, the larger the membrane potential, the higher the firing probability. For $\Delta t \rightarrow 0$ we may think of $g(u) R(t)$ as the instantaneous firing rate, or hazard of firing, given knowledge about the previous firing history. We focus on non-adapting neurons where the refractoriness R depends only on the timing of the last postsynaptic spike, but the model can be easily generalized to include a dependence upon earlier spikes as well. More specifically, we take for the simulations

$$R(t) = \frac{(t - \hat{t} - \tau_{\text{abs}})^2}{\tau_{\text{refr}}^2 + (t - \hat{t} - \tau_{\text{abs}})^2} \Theta(t - \hat{t} - \tau_{\text{abs}}), \quad [3]$$

where \hat{t} denotes the last firing time of the postsynaptic neuron, $\tau_{\text{abs}} = 3$ ms is the absolute refractory time, and $\tau_{\text{refr}} = 10$ ms is a parameter characterizing the duration of relative refractoriness. The Heaviside function $\Theta(x)$ takes a value of one for positive arguments and vanishes otherwise. With a function $R(t)$ such as in Eq. (3) that depends only on the most recent postsynaptic spike, the above neuron model has renewal properties and can be mapped onto a Spike Response Model with escape noise [9]. Except for Fig. 2, we take throughout the paper $g(u) = r_0 \log\{1 + \exp[(u - u_0)/\Delta u]\}$ with $u_0 = -65$ mV, $\Delta u = 2$ mV, and $r_0 = 11$ Hz. This set of parameters corresponds to *in vivo* conditions with a spontaneous firing rate of about 1 Hz. The function $g(u)$ and the typical firing behavior of the neuron model are shown in Fig. 1A. For Fig. 2, we consider the case $\tau_{\text{abs}} = \tau_{\text{refr}} = 0$ and an instantaneous rate $g_2(u) = \{10\text{ms} + [1/g(u)]\}^{-1}$ with $g(u)$ as above i.e., the neuron model exhibits no refractoriness and is defined by an inhomogeneous Poisson

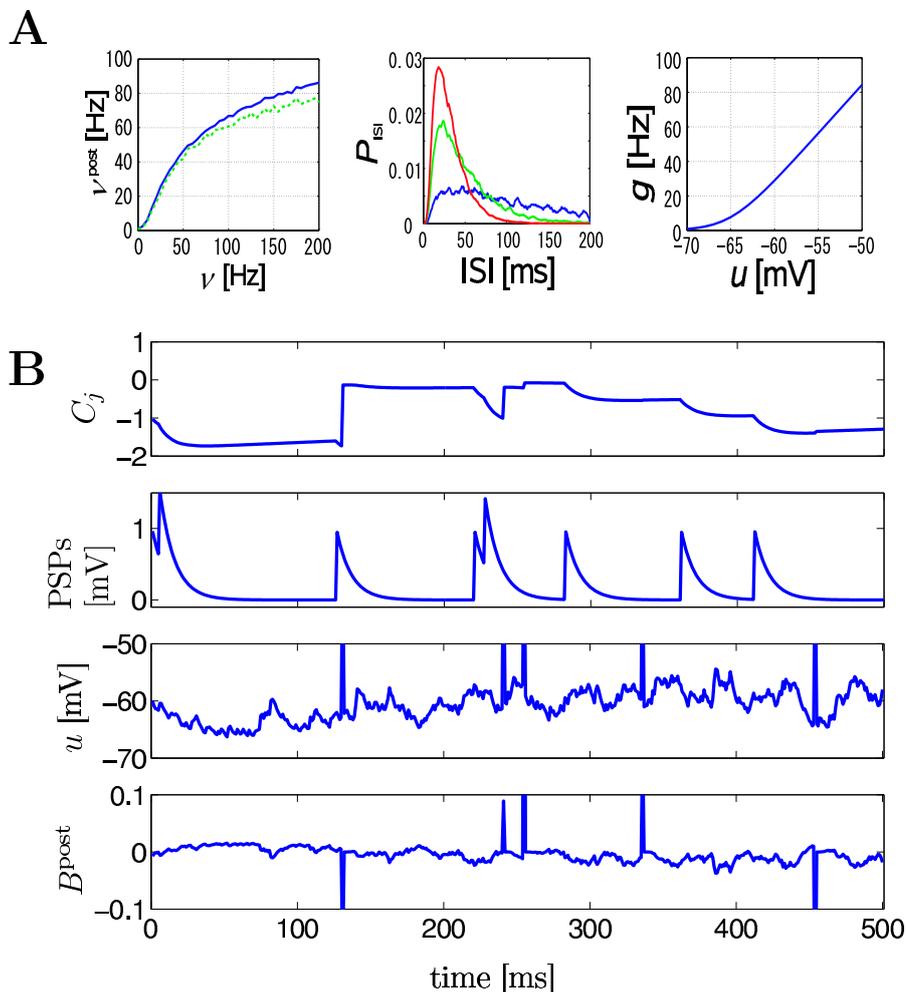


Figure 1: Neuron model. **A** *Left*: Output rate ν^{post} of the model neuron (solid line: spiking neuron model used in Figs. 3 - 5; green dashed line: Poisson model used in Fig. 2) as a function of presynaptic spike arrival rate at $N = 100$ synapses. All synapses have the same efficacy $w_j = 0.5$ and are stimulated by independent Poisson trains at the same rate ν . *Center*: Interspike interval distribution P_{ISI} of the spiking neuron model during firing at 10Hz (blue line), 20 Hz (green line), or 30 Hz (red line). Firing is impossible during the absolute refractory time of $\tau_{\text{abs}} = 3\text{ms}$. *Right*: The function $g(u)$ used to generate action potentials (see methods for details). **B** *From top to bottom*: The measure C_j that is sensitive to correlations between the state of the postsynaptic neuron and presynaptic spike arrival at synapse j ; the postsynaptic potentials caused by spike arrival at the same synapse j ; the membrane potential u ; and the postsynaptic factor B^{post} of Eq. (14) as a function of time. During postsynaptic action potentials, the postsynaptic factor B^{post} has marked peaks. Their amplitude and sign depend on the membrane potential at the moment of action potential firing. The coincidence measure C_j exhibits significant changes only during the duration of PSPs at synapse j .

process with maximum rate of 100Hz; cf. Fig. 1A. Integration of all equations is performed in Matlab on a standard personal computer using a time step $\Delta t = 1$ ms.

Spike Trains. The output of the postsynaptic neuron at time step t^k is denoted by a variable $y^k = 1$ if a postsynaptic spike occurred between t^{k-1} and t^k and zero otherwise. A specific output spike train up to time bin k is denoted by upper case letters $Y^k = \{y^1, y^2, \dots, y^k\}$. Since spikes are generated by a random process, we distinguish the random variable \mathbf{Y}^k by bold face characters from a specific realization Y^k . Note that the lower-case variable y^k refers always to a specific time bin, whereas the uppercase variable Y^k to a whole spike train. Similar remarks hold for the input: \mathbf{X} is the random variable characterizing the input at all synapse $1 \leq j \leq N$; X^k is a specific realization of all input spike trains up to time t^k ; $X_j^k = \{x_j^1, x_j^2, \dots, x_j^k\}$ a specific realization of an input spike train at synapse j , and x_j^k it's value in time bin k .

For given presynaptic spike trains X^k and postsynaptic spike history Y^{k-1} the probability of emitting a postsynaptic spike is described by the following binary distribution:

$$P(y^k | Y^{k-1}, X^k) = (\rho^k)^{y^k} (1 - \rho^k)^{(1-y^k)} \quad [4]$$

since $y^k \in \{0, 1\}$. Analogously, we find the marginal probability of y^k given Y^{k-1}

$$P(y^k | Y^{k-1}) = (\bar{\rho}^k)^{y^k} (1 - \bar{\rho}^k)^{(1-y^k)} \quad [5]$$

where $\bar{\rho}^k = \langle \rho^k \rangle_{\mathbf{X}^k | Y^{k-1}}$ and $\langle \cdot \rangle_{\mathbf{X}^k | Y^{k-1}} = \sum_{X^k} \cdot P(X^k | Y^{k-1})$.

From probability calculus we obtain the conditional probability of the output spike train Y^k given the presynaptic spike trains X^k ,

$$P(Y^k | X^k) = \prod_{l=1}^k P(y^l | Y^{l-1}, X^l), \quad [6]$$

and an analogous formula for the marginal probability distribution of $P(Y^k)$. With Eqs. (4) and (6) we have an expression for the probabilistic relation between an output spike train and an ensemble of input spike trains.

Mutual Information Optimization. Transmission of information between an ensemble of presynaptic spike trains \mathbf{X}^K of total duration $K \Delta t$ and the output train \mathbf{Y}^K of the postsynaptic neuron can be quantified by the mutual information [24]

$$I(\mathbf{Y}^K; \mathbf{X}^K) = \sum_{Y^K, X^K} P(Y^K, X^K) \log \frac{P(Y^K | X^K)}{P(Y^K)}, \quad [7]$$

While it is easier to transmit information if the postsynaptic neuron increases its firing rate, firing at high rates is costly from the point of view of energy consumption and also difficult to implement by the cells biophysical machinery. We therefore optimize information transmission under the condition that the firing statistics $P(Y^K)$ of the postsynaptic neuron stays as close as possible to a target distribution $\hat{P}(Y^K)$.

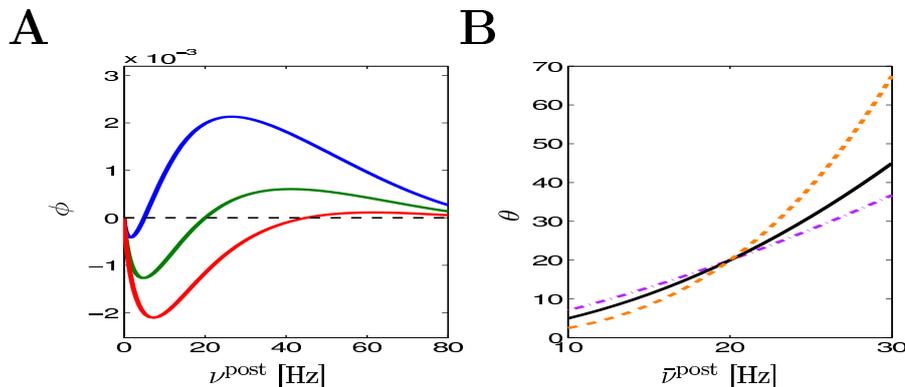


Figure 2: Relation to BCM rule. **A.** The function $\phi(\nu^{\text{post}})$ of the BCM learning rule (16) derived from our model under assumption of Poisson firing statistics of the postsynaptic neuron. A value of $\phi(\nu^{\text{post}}) > 0$ for a given postsynaptic rate ν^{post} means that synapses are potentiated when stimulated presynaptically. The transition from depression to potentiation occurs at a value θ that depends on the average firing rate $\bar{\nu}^{\text{post}}$ of the postsynaptic neuron (blue $\bar{\nu}^{\text{post}}=10$ Hz; green $\bar{\nu}^{\text{post}}=20$ Hz; red $\bar{\nu}^{\text{post}}=30$ Hz). **B.** The threshold θ as a function of $\bar{\nu}^{\text{post}}$ for different choices of the parameter γ , i.e., $\gamma = 0.5$ (purple); $\gamma = 1$ (black); $\gamma = 2$ (orange).

With a parameter γ (set to $\gamma = 1$ for the simulations), the quantity we maximize is therefore

$$\mathcal{L} = I(\mathbf{Y}^K; \mathbf{X}^K) - \gamma D(P(Y^K) || \tilde{P}(Y^K)), \quad [8]$$

where $D(P(Y^K) || \tilde{P}(Y^K)) = \sum_{Y^K} P(Y^K) \log(P(Y^K)/\tilde{P}(Y^K))$ denotes the Kullback-Leibler divergence [24]. The target distribution is that of a neuron with constant instantaneous rate \tilde{g} (set to $\tilde{g} = 30$ Hz throughout the paper except for Fig. 2) modulated by the refractory variable $R(t)$, i.e., that of a renewal process.

The main idea of our approach is as follows. We assume that synaptic efficacies w_j can change within some bounds $0 \leq w_j \leq w_{\text{max}}$ so as to maximize information transmission under the constraint of a fixed target firing rate. To derive the optimal rule of synaptic update, we calculate the gradient of Eq. (8). Applying the chain rule of information theory [24] to both the mutual information I and the Kullback-Leibler divergence D , we can write $\mathcal{L} = \sum_{k=1}^K \Delta \mathcal{L}^k$ where

$$\Delta \mathcal{L}^k = \left\langle \log \frac{P(y^k | Y^{k-1}, X^k)}{P(y^k | Y^{k-1})} - \gamma \log \frac{P(y^k | Y^{k-1})}{\tilde{P}(y^k | Y^{k-1})} \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k}, \quad [9]$$

with $\langle \cdot \rangle_{\mathbf{X}^k, \mathbf{Y}^k} = \sum_{X^k, Y^k} \cdot P(X^k, Y^k)$. Assuming slow changes of synaptic weights, we apply a gradient ascent algorithm to maximize the objective function and change the synaptic efficacy w_j at each time step by $\Delta w_j^k = \alpha (\partial \Delta \mathcal{L}^k / \partial w_j)$ with an appropriate

learning rate α . Evaluation of the gradient (see *Supporting Text*, which is published as supporting information on the PNAS web site) yields

$$\Delta w_j^k = \alpha \langle C_j^k (F^k - \gamma G^k) \rangle_{\mathbf{X}^k, \mathbf{Y}^k}, \quad [10]$$

with three functions C_j, F, G described in the following. First, with ρ' denoting the derivative of ρ with respect to u , the quantity

$$C_j^k = \sum_{l=k-k_a}^k \sum_{n=1}^l \epsilon(t^l - t^n) x_j^n \frac{\rho'^l}{\rho^l} \left[y^l - \frac{1 - y^l}{1 - \rho^l} \rho^l \right] \quad [11]$$

is a measure that counts coincidences between postsynaptic spikes ($y^l = 1$) and the time course of PSPs generated by presynaptic spikes ($x_j^n = 1$) at synapse j , normalized to an expected value $\langle C_j^k \rangle_{\mathbf{Y}^k | \mathbf{X}^k} = 0$. The time span k_a of the coincidence window is given by the width of the autocorrelation of the spike train of the postsynaptic neuron (see *Supporting Text* and Fig. 6 which are published as supporting information on the PNAS web site). The term

$$F^k = \log \frac{P(y^k | Y^{k-1}, X^k)}{P(y^k | Y^{k-1})} = y^k \log \frac{\rho^k}{\bar{\rho}^k} + (1 - y^k) \log \frac{1 - \rho^k}{1 - \bar{\rho}^k} \quad [12]$$

compares the instantaneous firing probability ρ^k at time step k with the average probability $\bar{\rho}^k$ and, analogously, the term $G^k = \log(P(y^k | Y^{k-1}) / \tilde{P}(y^k | Y^{k-1}))$ compares the average probability with the target value $\tilde{\rho}^k = \tilde{g}R(t^k)\Delta t$. We note that both F and G are functions of the postsynaptic variables only. We therefore introduce a postsynaptic factor B^{post} by the definition $B^{\text{post}}(t^k) = [F^k - \gamma G^k] / \Delta t$ and take the limit $\Delta t \rightarrow 0$. Under the assumption of small learning rate α (i.e., $\alpha = 10^{-4}$ in our simulations), the expectations $\langle \rangle_{\mathbf{X}, \mathbf{Y}}$ in Eq. (10) can be approximated by averaging over a single long trial which allows us to define an on-line rule $dw_j(t)/dt = \alpha C_j(t) B^{\text{post}}(t - \delta)$ with a postsynaptic factor

$$B^{\text{post}}(t) = \delta(t - \hat{t} - \delta) \log \left[\frac{g(u(t))}{\bar{g}(t)} \left(\frac{\tilde{g}}{\bar{g}(t)} \right)^\gamma \right] - R(t) [g(u(t)) - (1 + \gamma)\bar{g}(t) + \gamma\tilde{g}] \quad [13]$$

where \hat{t} is the firing time of the last postsynaptic spike. The delay δ in the Dirac- δ function reflects the order of updates in a *single* time step of the numerical implementation, i.e., we first update the membrane potential, then the last firing time \hat{t} , then the factors C_j and B , and finally the synaptic efficacy w_j ; for the mathematical theory we take $\delta \rightarrow 0$. The rate $\bar{g}(t) = \langle g(u(t)) \rangle_{\mathbf{X} | \mathbf{Y}}$ denotes an expectation over the input distribution given the recent firing history of the postsynaptic neuron. For a numerical implementation, it is convenient to estimate the expected rate $\bar{g}(t)$ by a running average with exponential time window (time constant 10s). Similarly, we replace the rectangular coincidence count window in Eq. (11) by an exponential one (time constant 1 s; see *Supporting Text*).

3 Results

We analyzed information transmission for a model neuron which receives input from one hundred presynaptic neurons. A presynaptic spike that arrives at time $t_j^{(f)}$ at synapse j evokes an excitatory postsynaptic potential of time course $\epsilon(t - t_j^{(f)})$. The amplitude w_j of the postsynaptic response is taken as a measure of synaptic efficacy and subject to synaptic dynamics. Firing of the postsynaptic neuron is more likely if the total PSP $u(t) = u_r + \sum_j \sum_f w_j \epsilon(t - t_j^{(f)})$ is large; however, due to refractoriness firing is suppressed after a postsynaptic action potential at time \hat{t} by a factor R which depends on the time since the last postsynaptic spike; see *Methods and Models* for details.

Maximizing the mutual information between several presynaptic spike trains and the output of the postsynaptic neuron can be achieved by a synaptic update rule which depends on the presynaptic spike arrival time $t_j^{(f)}$, the postsynaptic membrane potential u , as well as the last postsynaptic firing time \hat{t} . More precisely, the synaptic update rule can be written as

$$\frac{dw_j}{dt} = \alpha C_j(t) B^{\text{post}}(t - \delta) \quad [14]$$

where C_j is a measure sensitive to correlations between pre- and postsynaptic activity and B^{post} is a variable which characterizes the state of the postsynaptic neuron; (see *Methods and Models*). α is a small learning parameter. We note that in standard formulations of Hebbian learning, changes of synaptic efficacies are driven by correlations between pre- and postsynaptic neurons, similar to the function $C_j(t)$. The above update rule, however, augments these correlations by a further postsynaptic factor B^{post} .

This postsynaptic factor B^{post} depends on the firing time of the postsynaptic neuron; the refractory state of the neuron; its membrane potential u via the instantaneous firing intensity $g(u)$; as well as on its past firing history via $\bar{g}(t)$. The postsynaptic factor can be decomposed into two terms: the first one compares the instantaneous firing intensity $g(u)$ with its running average $\bar{g}(t)$; the second the running average with a target rate \tilde{g} . Thus the first term of B^{post} measures momentary significance of the postsynaptic state whereas its second term accounts for homeostatic processes; see *Methods and Models* for details.

The variable C_j measures correlations between the postsynaptic neuron and its presynaptic input at synapse j

$$\frac{dC_j(t)}{dt} = -\frac{C_j(t-\delta)}{\tau_C} + \sum_f \epsilon(t-t_j^{(f)}) S(t) [\delta(t-\hat{t}-\delta) - g(u(t))R(t)] \quad [15]$$

with time constant $\tau_C=1$ s. Here $g(u(t))R(t)$ is the instantaneous firing rate modulated by the refractory function $R(t)$, and $S(t) = g'(u(t))/g(u(t))$ is the sensitivity (the prime denotes the derivate with respect to u) of the neuron to a change of its membrane potential. The term with the Dirac $\delta(t - \hat{t} - \delta)$ induces a

positive jump of C_j immediately (with short delay δ) after each postsynaptic spike. Between postsynaptic spikes, C_j evolves continuously. Significant changes of C_j are conditioned on the presence of a PSP $\epsilon(t - t_j^{(f)})$ caused by spike arrival at synapse j . In the absence of presynaptic input the correlation estimate decays with time constant τ_C back to zero.

Both the correlation term C_j and the postsynaptic factor B^{post} can be estimated on-line (Fig. 1) and use only information which could, directly or indirectly, be available at the site of the synapse: information about postsynaptic spike firing could be conveyed by backpropagating action potentials; timing of presynaptic spike arrival is transmitted by neurotransmitter receptors; and the total postsynaptic potential can be estimated, although not perfectly, from the local potential at the synapse. The direction of change of a synapse is determined by a subtle interplay between the correlation term C_j and the postsynaptic factor B^{post} which can both be positive or negative.

To elucidate the balance between potentiation and depression of synapses, we first considered a simplified neuron model without refractoriness and firing rate $\nu^{\text{post}} = g_2(u)$. In this special case, the synaptic update rule (10) can be rewritten in the simpler form

$$\frac{d}{dt}w_j = \alpha \nu_j \phi(\nu^{\text{post}}, \theta) \quad [16]$$

where ν_j is the instantaneous firing rate of the presynaptic neuron as estimated from the amplitude of the PSP [if the potential is measured in mV and time in ms, then the proportionality constant a has units $1/(\text{mV ms})$]. $\nu_j = a \sum_f \epsilon(t - t_j^{(f)})$ generated at synapse j ; and $\phi(\nu^{\text{post}}, \theta) = f(\nu^{\text{post}}) \log(\nu^{\text{post}}/\theta)$ is a function which depends on the instantaneous postsynaptic firing rate $\nu^{\text{post}} = g_2(u)$ and a parameter θ . The function f is proportional to the derivative of g_2 , i.e., $f(\nu^{\text{post}}) = g_2'/a$ evaluated at $u = g_2^{-1}(\nu^{\text{post}})$. The parameter θ denotes the transition from a regime of potentiation to that of depression. It depends on the recent firing history of the neuron and is given by

$$\theta(t) = \bar{\nu}^{\text{post}}(t) \left(\frac{\bar{\nu}^{\text{post}}(t)}{\tilde{g}} \right)^\gamma \quad [17]$$

where $\tilde{g} = 20$ Hz denotes a target value of the postsynaptic rate implemented by homeostatic processes [22] and $\bar{\nu}^{\text{post}}(t)$ is a running average of the postsynaptic rate. The function ϕ in Eq. (16), shown in Fig. 2 is characteristic for the BCM learning rule [12]. Our approach by information maximization predicts a specific form of this function which can be plotted either as a function of the postsynaptic firing rate ν^{post} or as a function of the postsynaptic potential $u = g_2^{-1}(\nu^{\text{post}})$, in close agreement with experiments [2, 13]. Moreover, because information maximization was performed under the constraint of a fixed target firing rate, our approach yields automatically a sliding threshold of the form postulated in [12], but on different grounds. Thus for neurons without refractoriness, i.e., a pure rate model, our update rule for synaptic plasticity reduces exactly to the BCM rule.

An application of rate models to stimulation paradigms that vary on a time scale of tens of milliseconds or less, has often been questioned since an interpretation of ‘rate’ is seen as problematic. Our synaptic update rule is based on a *spiking* neuron model that includes refractoriness and which captures well properties of much more detailed neuron models [25]. The learning rule for spiking neurons has a couple of remarkable properties that we explore now.

First, we consider a pattern discrimination task in a rate coding paradigm. Patterns are defined by the firing rate ν^{pre} of 25 presynaptic neurons ($\nu^{\text{pre}} = 2, 13, 25, 40$ Hz for patterns, 1-4, respectively) modeled as independent Poisson spike trains. The remaining 75 synapses received uncorrelated Poisson input at a constant rate of 20 Hz. Each second, a pattern was chosen stochastically and applied during one second. Those synapses that received pattern-dependent input developed strong efficacies close to the maximal value $w_{\text{max}} = 1$ whereas most of the other 75 synapses developed weaker ones; cf. Fig. 3. However, some of the weakly driven synapses also spontaneously increased their efficacy. This was necessary for the neuron to achieve a mean postsynaptic firing rate close to the target firing rate. Despite the fact that the *mean* firing rate approaches on a time scale of tens of seconds the target rate, the spike count in each one-second segment is strongly modulated by the input pattern and can be used to distinguish the patterns (Fig. 3B) with a misclassification error of $< 20\%$. (See *Supporting Text* and Fig. 7, which is published as supporting information on the PNAS web site.)

In a second set of simulation experiments, we studied again rate coding, but considered two groups of input defined by periodic modulation of the presynaptic firing rates. More precisely, 40 neurons received input spike trains generated by an inhomogeneous Poisson process with common rate modulation $\nu(t) = \nu_0 + A \sin(2\pi t/T)$ with amplitude $A = 10$ Hz and period $T = 100$ ms. Another group of 40 neurons received modulated input of the same amplitude and period, however, with a phase shift π . The remaining 20 synapses received Poisson input at a fixed rate $\nu_0 = 20$ Hz. All 100 inputs projected onto 9 postsynaptic neurons. Synaptic weights were initialized randomly (between 0.10 and 0.12) and evolved according to the update rule (14). Out of the 9 postsynaptic neurons, 4 developed strong connections to the first group of correlated inputs, while 5 developed strong connections to the second group of correlated inputs (Fig. 4A). The rate of the output neurons reflects the modulation of their respective inputs (Fig. 4B). To summarize, the synaptic update rule derived from the principle of information maximization drives neurons to spontaneously detect and specialize for groups of coherent inputs. Just as in the standard BCM rule, several output neurons (with different specialization) are needed to account for the different features of the input.

Whereas the two preceding paradigms focused on *rate* modulation, we now show that, even if all presynaptic neurons fire at the same mean rate, the presence of weak spike-spike correlations in the input is sufficient to bias the synaptic selection mechanism; cf. Fig. 5. All synapses received spike input at the same rate of 20 Hz, but the spike trains of 50 synapses showed weak correlations ($c = 0.1$), whereas the remaining 50 synapses received uncorrelated input. Most of the 50 synapses that

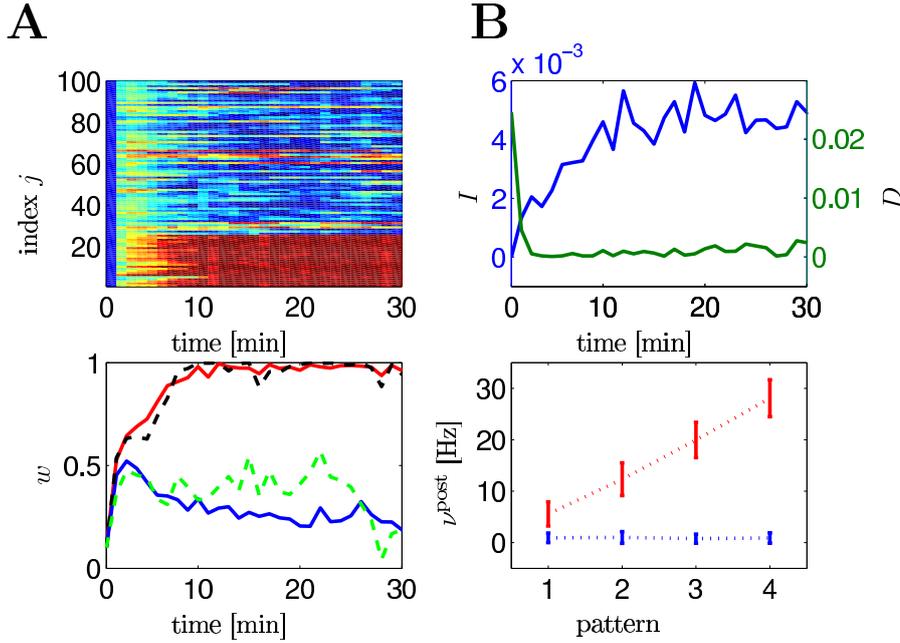


Figure 3: Pattern discrimination. The first 25 synapses $1 \leq j \leq 25$ are stimulated by Poisson input with a rate $\nu^{\text{pre}} = 2, 13, 25, 40$ Hz that changes each second. The remaining 75 synapses receive Poisson input at a constant rate of 20 Hz. **A** *Upper*. Evolution of all synaptic weights as a function of time (red: strong synapses $w_j \approx 1$; blue : depressed synapses $w_j \approx 0$). All synapses are initialized at the same value $w_j = 0.1$. *Lower*. The evolution of the average efficacy of the 25 synapses that receive pattern-dependent input (red line) and that of the 75 other synapses (blue). Typical examples of individual traces (synapse 1: black and synapses 30: green) are given by the dashed lines. **B** *Upper*. Evolution of the average mutual information I per bin (blue line and left scale) and of the average Kullback-Leibler distance D per bin as a function of time. Averages are calculated over segments of 1 min. *Lower* Output rate (spike count during 1 second) as a function of pattern index before (blue bars) and after (red bars) learning.

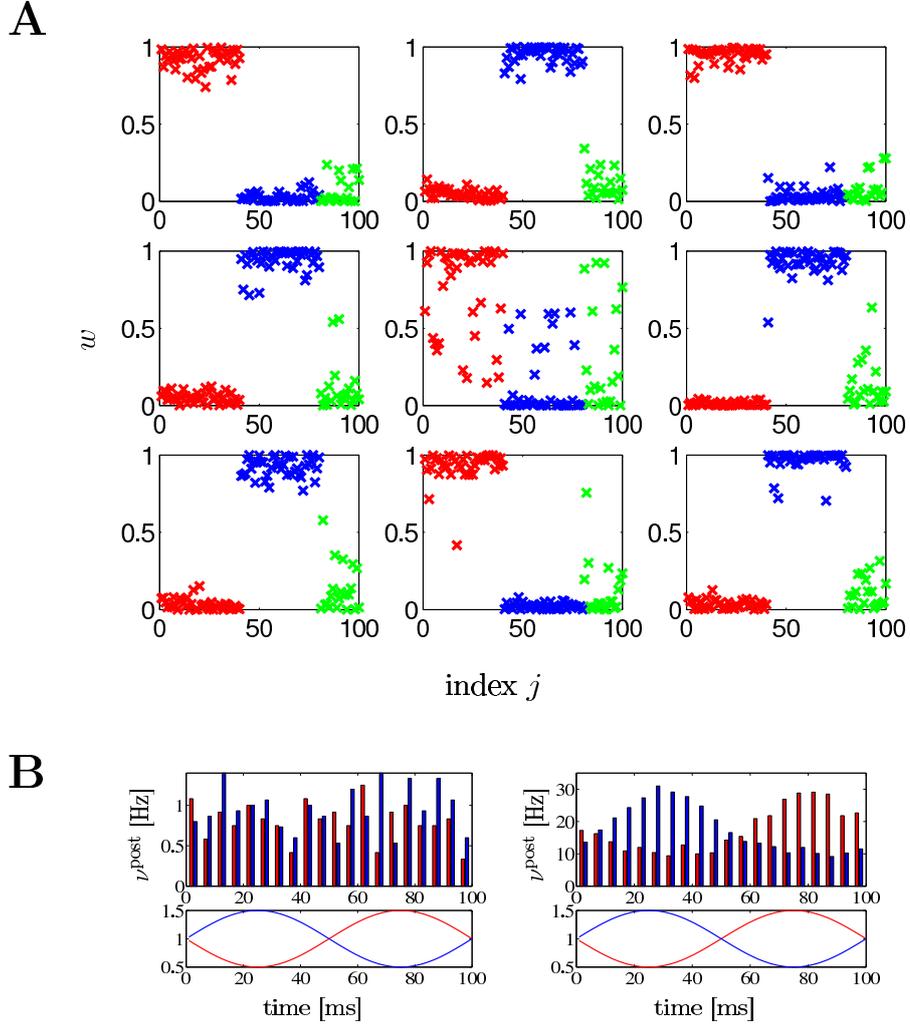


Figure 4: Rate modulation. **A**. Distribution of synaptic efficacies of nine postsynaptic neurons after 60 min of stimulation with identical inputs for all neurons. Synapses $1 \leq j \leq 40$ (red symbols) received Poisson input with common rate modulation; the input at synapses $41 \leq j \leq 80$ (blue symbols) was also rate modulated but phase shifted; and the input at the remaining 20 synapses was uncorrelated (green symbols). Four postsynaptic neurons (number 1, 3, 5, 8) develop a spontaneous specialization for the first group of modulated input (red symbols close to the maximum efficacy of one) and five (2, 4, 6, 7, 9) specialized for the second group. **B**. Modulation of the output rate of the 9 postsynaptic neurons before (left) and after (right) learning. Red/blue bars: neurons responding to the first/second group of input. Red/blue lines: modulation of the input of group 1 and 2.

received a weakly correlated input increased their weights under application of the synaptic update rule (Fig. 5). If at a later stage the group of correlated inputs changes, the newly included synapses will be strengthened as well whereas those that are no longer correlated decay. Moreover, the final distribution of synaptic efficacies persists, even if the input is switched to random spike arrival. Thus the synaptic update rule is sensitive to spike-spike correlations on a millisecond scale, which would be difficult to account for in a pure rate model. A biophysical signature of spike-spike correlations are systematic and large fluctuations of the membrane potential. If several postsynaptic neurons receive the same input, their outputs are again correlated and generate large fluctuations in the membrane potential of a read-out neuron further down the processing chain (Fig. 5B). Thus, information that is potentially encoded in millisecond correlations in the input can be detected, enhanced, transmitted, and read out by other neurons.

4 Discussion

The synaptic update rule discussed in this paper relies on the maximization of the mutual information between an ensemble of presynaptic spike trains and the output of the postsynaptic neuron. As in all optimization approaches, optimization has to be done under some constraint. Since information scales with the postsynaptic firing rate, but high firing rates cannot be sustained by the biophysical machinery of the cell over long times, we imposed that, on average, the postsynaptic firing rate should stay close to a desired firing rate. This idea is consistent with the widespread finding of homeostatic processes that tend to push a neuron always back into its preferred firing state [22]. The implementation of this idea in our formalism gave naturally rise to a control mechanism that corresponds exactly to the sliding threshold in the original BCM model [12]. While derivations of the BCM model from optimality concepts [19, 23] or statistical approaches [26] are not new, our approach gives a new perspective on the concept of a sliding threshold.

Our derivation extends the BCM model, which was originally designed for rate models of neuronal activity to the case of spiking neuron models with refractoriness. Spiking neuron models of the integrate-and-fire type, such as the one presented here, can be used to account for a broad spectrum of neuronal firing behavior, including the role of spike-spike correlations, interspike interval distributions, coefficient of variations, and even timing of single spikes; for a review, see ref. [9]. The essential ingredients of the spiking neuron model considered here were *(i)* PSP generated by presynaptic spike arrival; *(ii)* a heuristic spiking probability that depends on the total postsynaptic potential; and *(iii)* a phenomenological account of absolute or relative refractoriness. The synaptic update rule depends on all three of these quantities. While we do not imply that synaptic potentiation and depression of real neurons are implemented the way it is suggested by our update rule, the rule shows nevertheless some interesting features.

First, in contrast to pure Hebbian correlation driven learning, the update rule uses a correlation term modulated by an additional postsynaptic factor. Thus, presy-

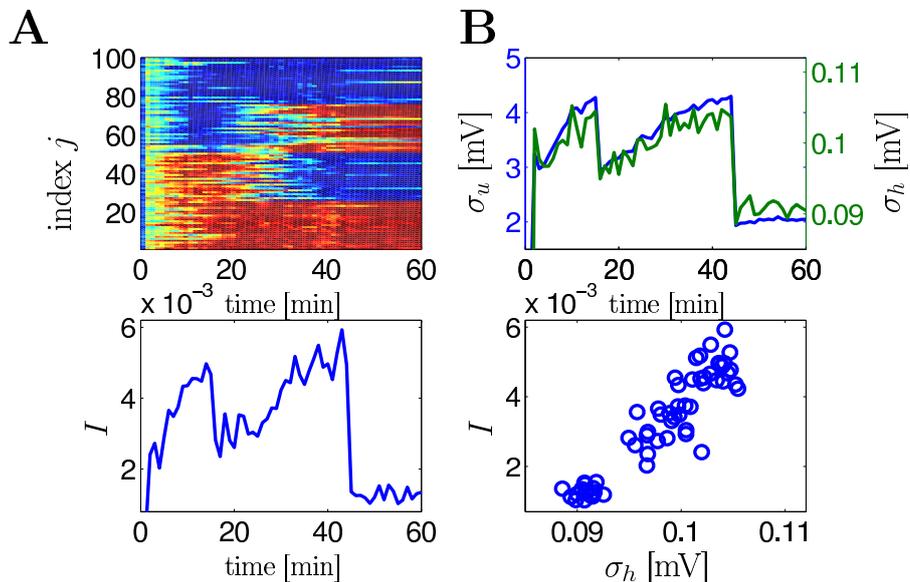


Figure 5: Spike-spike correlations. The $N = 100$ synapses have been separated into four groups of 25 neurons each (group A, $1 \leq j \leq 25$; B, $26 \leq j \leq 50$; C, $51 \leq j \leq 75$; D, $76 \leq j \leq 100$). All synapses were stimulated at the same rate of 20 Hz. But during the first 15 min of simulated time, neurons in groups C and D were uncorrelated whereas the spike trains of the remaining 50 neurons (groups A and B) had correlations of amplitude $c = 0.1$, i.e., ten percent of the spike arrival times were identical between each pair of synapses. After 15 minutes, correlations changed so that group A became correlated with C whereas B and D were uncorrelated. After 45 minutes of simulated time correlations stopped, but stimulation continued at the same rate. **A.** Upper: Evolution of all 100 weights (red: potentiated; blue: depressed). Lower: Average mutual information per bin as a function of time. In the absence of correlations ($t > 45$ min) mutual information is lower than before, but the distribution of synaptic weights remains stable. **B.** 9 postsynaptic neurons $1 \leq i \leq 9$ with membrane potential $u_i(t)$ are stimulated as discussed in A and project to a read-out unit with potential $h(t) = \sum_{i=1}^9 \sum_m \epsilon(t - t_i^m)$ where the sum runs over all output spikes m of all 9 neurons. Mean membrane potentials are \bar{u} and \bar{h} , respectively. The fluctuations $\sigma_u = [\langle (u_i(t) - \bar{u})^2 \rangle]^{0.5}$ of the postsynaptic potentials (blue line, top graph) and those of the readout potentials ($\sigma_h = [\langle (h(t) - \bar{h})^2 \rangle]^{0.5}$, green line) are correlated (Lower) with the mutual information and can serve as neuronal signal.

naptic stimulation is combined with a highly nonlinear function of the postsynaptic state in order to determine the direction and amplitude of synaptic changes. The essence of the BCM rule (presynaptic gating combined with nonlinear postsynaptic term) is hence translated into a spike-based formulation.

Second, the spike-based formulation of a synaptic update rule should allow a connection to spike-timing dependent plasticity [5, 6] and allow its interpretation in terms of optimal information transmission [27–29]. Given the highly nonlinear involvement of postsynaptic spike times and postsynaptic potential in the optimal synaptic update rule, a simple interpretation in terms of pairs or pre- and postsynaptic spikes as in many standard models of synaptic plasticity [30, 31] can only capture a small portion of synaptic plasticity phenomena. The optimal learning rule suggests that nonlinear phenomena [32–35] are potentially highly relevant.

Acknowledgments T.T. is supported by the Japan Society for the Promotion of Science and a Grant-in-Aid for JSPS fellows and J.-P.P. by the Swiss National Science Foundation.

References

- [1] Bliss, T.V.P. & Gardner-Medwin, A.R. (1973) *J. Physiol.* **232**, 357–374.
- [2] Dudek, S. M. & Bear, M. F. (1993) *J. Neurosci.* **13**, 2910–2918.
- [3] Artola, A., Bröcher, S., & Singer, W. (1990) *Nature* **347**, 69–72.
- [4] Levy, W. B. & Stewart, D. (1983) *Neurosci.*, **8**, 791–797.
- [5] Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997) *Science* **275**, 213–215.
- [6] Bi, G.-q. & Poo, M.-m. (2001) *Ann. Rev. Neurosci.* **24**, 139–166.
- [7] Malenka, R. C. & Nicoll, R. A. (1999) *Science* **285**, 1870–1874.
- [8] Dayan, P. & Abbott, L. F. (2001) *Theoretical Neuroscience* (MIT Press, Cambridge).
- [9] Gerstner, W. & Kistler, W. K. (2002) *Spiking Neuron Models* (Cambridge University Press, Cambridge UK).
- [10] Cooper, L.N., Intrator, N., Blais, B.S., & Shouval, H. Z. (2004) *Theory of cortical plasticity*. (World Scientific, Singapore).
- [11] Hebb, D. O. (1949) *The Organization of Behavior* (Wiley, New York).
- [12] Bienenstock, E.L., Cooper, L.N., & Munro, P.W. (1982) *J. Neurosci.* **2**, 32–48, reprinted in Anderson and Rosenfeld, 1990.
- [13] Artola, Alain & Singer, Wolf (1993) *Trends Neurosci.* **16**(11), 480–487.
- [14] Kirkwood, A., Rioult, M.G., & Bear, M.F. (1996) *Nature* **381**, 526–528.
- [15] Barlow, H.B. & Levick, W.R. (1969) *J. Physiology (London)* **200**, 1–24.
- [16] Laughlin, S. (1981) *Z. Naturforschung* **36**, 910–912.
- [17] Fairhall, A. L., Lewen, G.D, Bialek, W., & vanSteveninck, R.R.D. (2001) *Nature* **412**, 787–792.
- [18] Linsker, R. (1989) *Neural Computation* **1**(3), 402–411.
- [19] Nadal, J-P. & Parga, N. (1997) *Neural Computation* **9**, 1421–1456.
- [20] Bell, A.J. & Sejnowski, T.J. (1995) *Neural Computation* **7**, 1129–1159.
- [21] Chechik, G. (2003) *Neural Computation* **15**, 1481–1510.
- [22] Turrigiano, G.G. & Nelson, S.B. (2004) *Nature Reviews Neuroscience* **5**, 97–107.

- [23] Intrator, N. & Cooper, L.N. (1992) *Neural Networks* **5**, 3–17.
- [24] Cover, T.M. & Thomas, J.A. (1991) *Elements of Information Theory* (Wiley, New York).
- [25] Jolivet, R., Lewis, T.J., & Gerstner, W. (2004) *J. Neurophysiol.* **92**, 959–976.
- [26] Beggs, J.M. (2000) *Neural Computation* **13**, 87–111.
- [27] Toyozumi, T., Pfister J.P., Aihara, K., Gerstner W. (2005) in *Advances in Neural Information Processing Systems*, eds. Saul L.K., Weiss Y., Bottou L. (MIT Press, Cambridge, MA), Vol. 17, pp. 1409-1416.
- [28] Bell A.J., Parra L.C. (2005) in *Advances in Neural Information Processing Systems*, eds. Saul L.K., Weiss Y., Bottou L. (MIT Press, Cambridge, MA), Vol. 17, pp. 121-128.
- [29] Bohte S.M., Mozer M.C. (2005) in *Advances in Neural Information Processing Systems*, eds. Saul L.K., Weiss Y., Bottou L. (MIT Press, Cambridge, MA), Vol. 17, pp. 201-208.
- [30] Gerstner, W., Kempter, R., vanHemmen, J. Leo, & Wagner, H. (1996) *Nature* **383**, 76–78.
- [31] Song, S., Miller, K.D., & Abbott, L.F. (2000) *Nature Neuroscience* **3**, 919–926.
- [32] Froemke, R. & Dan, Y. (2002) *Nature* **416**, 433–438.
- [33] Bi, G.-Q. & Wang, H.-X. (2002) *Physiology and Behavior* **77**, 551–555.
- [34] Senn, W. (2002) *Biol. Cybern.* **87**, 344–355.
- [35] Sjöström, P.J., Turrigiano, G.G., & Nelson, S.B. (2001) *Neuron* **32**, 1149–1164.

Supporting Text

Evolution of the Average Synaptic Update Rule

In this appendix we evaluate the derivative of Eq. 9 in the main text, i.e., we need to calculate

$$\frac{\partial}{\partial w_j} \left\langle \log \frac{P(y^k | Y^{k-1}, X^k)}{P(y^k | Y^{k-1})} - \gamma \log \frac{P(y^k | Y^{k-1})}{\tilde{P}(y^k | Y^{k-1})} \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k}. \quad [1]$$

Before we start let us recall some notation. The average of an arbitrary function f_w with arguments x and y is by definition

$$\langle f_w(x, y) \rangle_{\mathbf{x}, \mathbf{y}} = \sum_x \sum_y p_w(x, y) f_w(x, y) \quad [2]$$

where $p_w(x, y)$ denotes the joint probability of the pair (x, y) to occur and the sum runs over all configurations of x and y . The subscript w indicates that both the probability distribution p_w and the function f_w may depend on a parameter w .

By definition, we have $p_w(x, y) = p_w(y|x)p(x)$ where $p(x)$ is a given input distribution and $p_w(y|x)$ the (parameter-dependent) conditional probability of generating an output y given x . Hence Eq. 2 can be transformed into

$$\langle f_w(x, y) \rangle_{\mathbf{x}, \mathbf{y}} = \sum_x p(x) \sum_y p_w(y|x) f_w(x, y) = \left\langle \sum_y p_w(y|x) f_w(x, y) \right\rangle_{\mathbf{x}} \quad [3]$$

If we now take the derivative with respect to the parameter w , the product rule yields two terms

$$\begin{aligned} \frac{\partial}{\partial w} \langle f_w(x, y) \rangle_{\mathbf{x}, \mathbf{y}} &= \left\langle \sum_y p_w(y|x) \frac{\partial}{\partial w} f_w(x, y) \right\rangle_{\mathbf{x}} \\ &+ \left\langle \sum_y p_w(y|x) \left[\frac{\partial}{\partial w} \log p_w(y|x) \right] f_w(x, y) \right\rangle_{\mathbf{x}} \end{aligned} \quad [4]$$

The first term contains the derivative of the function f_w whereas the second term contains the derivative of the conditional probability p_w . We note that Eq. 4 can also be written in the form

$$\frac{\partial}{\partial w} \langle f_w(x, y) \rangle_{\mathbf{x}, \mathbf{y}} = \left\langle \frac{\partial}{\partial w} f_w(x, y) \right\rangle_{\mathbf{x}, \mathbf{y}} + \left\langle \left[\frac{\partial}{\partial w} \log p_w(y|x) \right] f_w(x, y) \right\rangle_{\mathbf{x}, \mathbf{y}}, \quad [5]$$

i.e., as an average over the joint distribution of x and y . This formulation will be useful for the problem at hand.

The gradient in Eq. 1 contains several terms and for the moment we pick only one of these. The others will then be treated analogously. Let us focus on the term $\langle \log P(y^k | Y^{k-1}, X^k) \rangle_{\mathbf{Y}^k, \mathbf{X}^k}$ and apply steps completely analogous to those leading

from Eqs. 2-5.

$$\begin{aligned}
& \frac{\partial}{\partial w_j} \left\langle \log P(y^k | Y^{k-1}, X^k) \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k} \\
&= \left\langle \frac{\partial}{\partial w_j} \log P(y^k | Y^{k-1}, X^k) \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k} \\
&\quad + \left\langle \left[\frac{\partial}{\partial w_j} \log P(Y^k | X^k) \right] \log P(y^k | Y^{k-1}, X^k) \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k}
\end{aligned} \tag{6}$$

We now evaluate the averages using the identity

$\langle \cdot \rangle_{\mathbf{Y}^k, \mathbf{X}^k} = \langle \cdot \rangle_{y^k | Y^{k-1}, X^k} \rangle_{\mathbf{Y}^{k-1}, \mathbf{X}^k}$. We find that the first term on the right-hand side of Eq. 6 vanishes, since

$$\begin{aligned}
& \left\langle \frac{\partial}{\partial w_j} \log P(y^k | Y^{k-1}, X^k) \right\rangle_{y^k | Y^{k-1}, X^k} \\
&= \sum_{y^k \in \{0,1\}} \frac{\partial}{\partial w_j} \left[\log P(y^k | Y^{k-1}, X^k) \right] P(y^k | Y^{k-1}, X^k) \\
&= \frac{\partial}{\partial w_j} \left[\sum_{y^k \in \{0,1\}} P(y^k | Y^{k-1}, X^k) \right] = 0
\end{aligned} \tag{7}$$

because of the normalization of probabilities. The same argument can be repeated to show that $0 = \left\langle \frac{\partial}{\partial w_j} \log P(y^k | Y^{k-1}) \right\rangle_{y^k | Y^{k-1}, X^k}$. The reference distribution $\tilde{P}(y^k | Y^{k-1})$ is by definition independent of w_j .

Hence the only term that gives a non-trivial contribution on the right-hand side of Eq. 6 is the second term. With an analogous argument for the other factors in Eq. 1 we have

$$\begin{aligned}
& \frac{\partial}{\partial w_j} \left\langle \log \frac{P(y^k | Y^{k-1}, X^k)}{P(y^k | Y^{k-1})} - \gamma \log \frac{P(y^k | Y^{k-1})}{\tilde{P}(y^k | Y^{k-1})} \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k} \\
&= \left\langle \left[\frac{\partial \log P(Y^k | X^k)}{\partial w_j} \right] \left(\log \frac{P(y^k | Y^{k-1}, X^k)}{P(y^k | Y^{k-1})} - \gamma \log \frac{P(y^k | Y^{k-1})}{\tilde{P}(y^k | Y^{k-1})} \right) \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k}
\end{aligned} \tag{8}$$

An identification of the factors $C, F,$ and G in the main text is straightforward. From Eq. 4 in the main text we have

$$\log P(y^k | Y^{k-1}, X^k) = y^k \log(\rho^k) + (1 - y^k) \log(1 - \rho^k) \tag{9}$$

Hence we can evaluate the factors

$$\begin{aligned}
F^k &= \log \frac{P(y^k | Y^{k-1}, X^k)}{P(y^k | Y^{k-1})} = y^k \log \frac{\rho^k}{\bar{\rho}^k} + (1 - y^k) \log \frac{1 - \rho^k}{1 - \bar{\rho}^k} \\
G^k &= \log \frac{P(y^k | Y^{k-1})}{\tilde{P}(y^k | Y^{k-1})} = y^k \log \frac{\bar{\rho}^k}{\bar{\rho}} + (1 - y^k) \log \frac{1 - \bar{\rho}^k}{1 - \bar{\rho}}
\end{aligned}$$

Furthermore we can calculate the derivative needed in Eq. 8 using the chain rule from Eq. 6 of the main text, i.e.,

$$P(Y^k | X^k) = \prod_{l=1}^k P(y^l | Y^{l-1}, X^l) \tag{10}$$

which yields

$$\frac{\partial \log P(Y^k|X^k)}{\partial w_j} = \frac{\partial}{\partial w_j} \sum_{l=1}^k \log P(y^l|Y^{l-1}, X^l) \quad [11]$$

$$= \sum_{l=1}^k \left[\frac{y^l}{\rho^l} - \frac{1-y^l}{1-\rho^l} \right] \rho^l \sum_n \epsilon(t^l - t^n) x_j^n \quad [12]$$

We note that in Eq. 8 the factor $\frac{\partial}{\partial w_j} \log P(Y^k|X^k)$ has to be multiplied with F^k or with G^k before taking the average. Multiplication generates terms of the form $\langle y^l y^k \rangle_{\mathbf{Y}^k, \mathbf{X}^k} = \langle \langle y^l y^k \rangle_{\mathbf{Y}^k|X^k} \rangle_{\mathbf{X}^k}$. For any given input X^k , the autocorrelation $\langle y^l y^k \rangle_{\mathbf{Y}^k|X^k}$ with $l < k - k_a$ of the postsynaptic neuron will have a trivial value

$$\langle y^l y^k \rangle_{\mathbf{Y}^k|X^k} = \langle y^l \rangle_{\mathbf{Y}^k|X^k} \langle y^k \rangle_{\mathbf{Y}^k|X^k} \quad \text{for } k-l > k_a \quad [13]$$

where $k_a \Delta t$ is the width of the autocorrelation. As a consequence

$$\left\langle \left[\frac{y^l}{\rho^l} - \frac{1-y^l}{1-\rho^l} \right] (F^k - \gamma G^k) \right\rangle_{\mathbf{Y}^k, \mathbf{X}^k} = 0 \quad \text{for } k-l > k_a \quad [14]$$

Hence, for $k-l > k_a$, we can truncate the sum over l in Eq. 12, i.e., $\sum_{l=1}^k \rightarrow \sum_{l=k-k_a}^k$ which yields exactly the coincidence measure C_j introduced in the main text; cf. Eq. 11 in the main text, and which we repeat here for convenience

$$C_j^k = \sum_{l=k-k_a}^k \left[\frac{y^l}{\rho^l} - \frac{1-y^l}{1-\rho^l} \right] \rho^l \sum_n \epsilon(t^l - t^n) x_j^n \quad [15]$$

From Averages to an Online Rule

The coincidence measure C_j^k counts coincidences in a rectangular time window. If we replace the rectangular time window by an exponential one with time constant τ_C and go to continuous time, the summation $\sum_{l=k-k_a}^k \dots$ in Eq. 15 turns into an integral $\int_{-\infty}^t dt' \exp[-(t-t')/\tau_C] \dots$ which can be transformed into a differential equation

$$\frac{dC_j(t)}{dt} = -\frac{C_j(t-\delta)}{\tau_C} + \sum_f \epsilon(t-t_j^{(f)}) S(t) [\delta(t-\hat{t}-\delta) - g(u(t)) R(t)] ; \quad [16]$$

cf. Eq. 15 in the main text. Based on the considerations in the previous paragraph, the time constant τ_C should best be chosen in the range $k_a \Delta t \leq \tau_C \leq 10 k_a \Delta t$.

Similarly, the average firing rate $\bar{\rho}(t) = \bar{g}(t) R(t)$ can be estimated using a running average

$$\tau_{\bar{g}} \frac{d\bar{g}(t)}{dt} = -\bar{g}(t) + g(u(t)) \quad [17]$$

with time constant $\tau_{\bar{g}}$.

In Fig. 6, we compare the performance of three different update schemes in numerical simulations. In particular, we show that (i) the exact value of the truncation

of the sum in Eq. 15 is not relevant, as long as $k_a \Delta t$ is larger than the width of the autocorrelation; and (ii) that the online rule is a good approximation to the exact solution.

To do so we take the scenario from Fig. 3 of the main text. For each segment of 1 s, we simulate one hundred pairs of input and output spike trains. We evaluate numerically Eq. 8 by averaging over the 100 samples. After each segment of 1 second (=1,000 time steps) we update the weights using a rule without truncation in the sum of Eq. 15. We call this the full batch update; compare. Fig. 6 (*Top*).

Second, we use the definition of C_j^k with the truncated sum and repeat the above steps; Fig. 6 (*Middle*). The truncation is set to $k_a \Delta t = 200\text{ms}$ which is well above the expected width of the autocorrelation function of the postsynaptic neuron. We call this the truncated batch rule.

Third, we use the online rule discussed in the main body of the paper with $\tau_C = 1$ s; Fig. 6 (*Bottom*).

Comparison of top and center graphs of Fig. 6 shows that there is no difference in the evolution of mean synaptic efficacies, i.e., the truncation of the sum is allowed, as expected from the theoretical arguments. A further comparison with Fig. 6 *Bottom* shows that updates based on the online rule add some fluctuations to the results, but its trend captures nicely the evolution of the batch rules.

Supplement to the Pattern Detection Paradigm

In Fig. 3 we presented a pattern detection paradigm where patterns defined by input rates were chosen randomly and applied for one second. After learning, the spike count over one second is sensitive to the index of the pattern. Fig. 7A shows the histogram of spike counts for each pattern. Optimal classification is achieved by choosing for each spike count the pattern which is most likely. With this criterion 81 percent of the patterns will be classified correctly.

The update of synaptic efficacies depends on the choice of the parameter γ in the learning rule. According to the optimality criterion in Eq. 8 of the main text, a high level of γ implies a strong homeostatic control of the firing rate of the postsynaptic neuron whereas a low level of γ induces only a weak homeostatic control. In order to study the role of γ , we repeated the numerical experiments for the above pattern detection paradigm with a value of $\gamma = 100$ instead of our standard value of $\gamma = 1$. Fig. 7B shows that the output firing rate is still modulated by the pattern index, the modulation at $\gamma = 100$ is, however, weaker than that at $\gamma = 1$. As a result, pattern detection is less reliable with 45 percent correct classification only. We note that this is still significantly higher than the chance level of 25 percent.

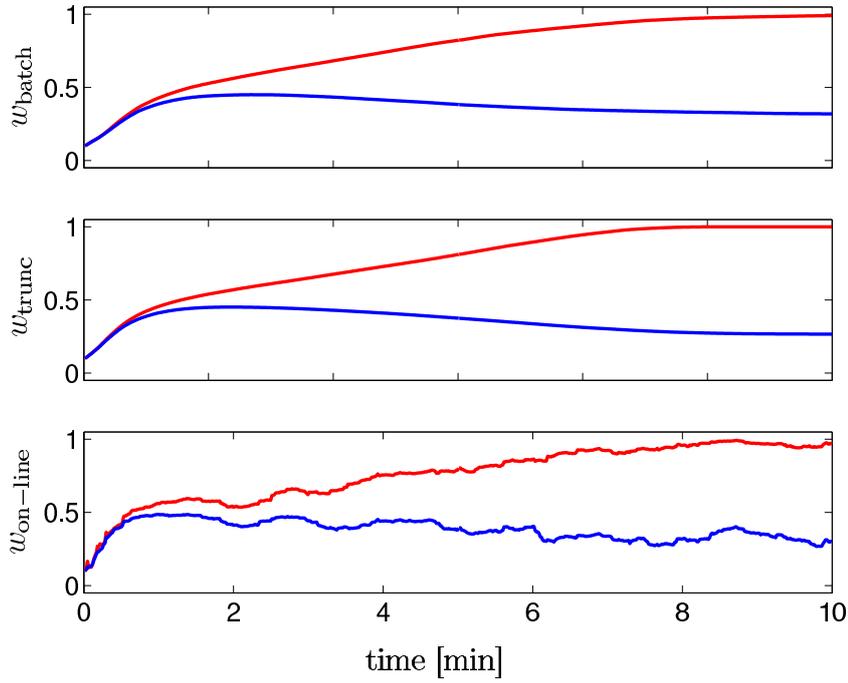


Figure 6: Evolution of the synaptic efficacies for the pattern detection paradigm of Fig. 3 during the first 10 minutes of simulated time. Red: mean synaptic efficacy of the 25 synapses that received pattern-dependent input rates. Blue: mean synaptic efficacy of the remaining 75 synapses. The batch update rule (top), the truncated batch rule (middle) and the online rule (bottom) yield comparable results.

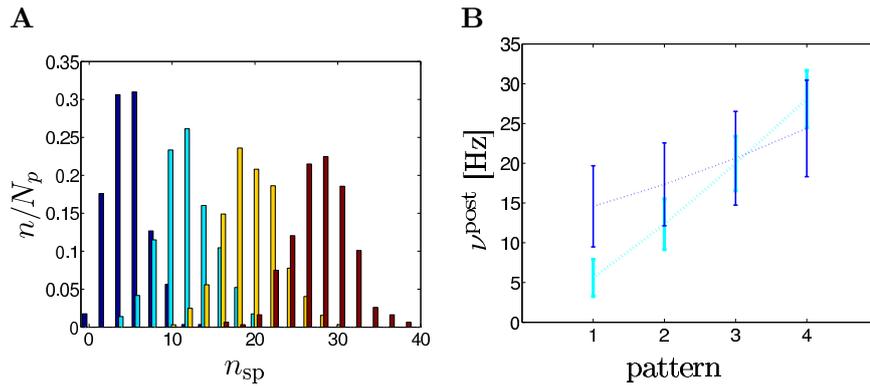


Figure 7: Pattern detection. **A** Histograms of spike counts n_{sp} over one second (horizontal axis, bin size 2) during presentation of pattern 1 (dark blue), pattern 2 (light blue), pattern 3 (yellow), and pattern 4 (red). Vertical scale: number of trials n with a given spike count divided by total number N_p of trials for that pattern. **B** Spike count during one second (mean and variance) for each of the four patterns with a parameter value $\gamma = 1$ (light blue) and $\gamma = 100$ (dark blue). The values for $\gamma = 1$ are redrawn from Fig. 3.

2.5 Paper V

Summary

This paper is the third paper in this thesis (c.f. papers III and IV) which considers an unsupervised scenario. This time the objective function is composed of an information term, an homeostatic term (identical to the one of paper IV) and an extra activity dependent weight decay term.

This last term has been added in order to prevent weights to become arbitrarily large, which is unrealistic. It should be noted that in the previous paper, which did not have the weight decay term, it was necessary to impose a fixed upper bound. Indeed, the homeostatic term prevents the postsynaptic firing to be too large, but does not prevent individual synapses from unrealistic growth. So the weight decay term replaces the strict upper bound.

Another motivation for this quadratic weight penalty term is that the learning rule, when applied to a STDP protocol, elicits a positive offset. So if the relative importance of the extra penalty term is well chosen, it can compensate exactly the positive offset (see Fig. 2d of paper V). Similarly to the previous papers, the shape of the potentiation part of the learning window is determined by the EPSP time course. Depression comes from the EPSP suppression we assumed in this paper. Indeed, if an EPSP is elicited just after the postsynaptic spike, its amplitude is reduced and hence transmits less information.

We also added a weight dependent learning rate which vanishes for small weights. In this way, there is no need to implement a strict lower bound at $w = 0$. This specific implementation of the learning rate brings interesting feature to the learning rule. One of them is the ability to store memories for extremely long time. Indeed, if a neuron specializes to a given input and thereby induces some small synaptic weights, those synapses will hardly change and therefore contribute to the memory maintenance.

Another interesting feature of the derived learning rule is that unspecialized synapse pattern (which corresponds to an unselective neuron) can be stable (see Fig. 3c of paper V). This is in direct contrast with the BCM learning rule. Indeed, the only stable fixed point of the BCM dynamics are the ones where the neuron is selective. This feature of our learning rule directly comes from the quadratic weight penalty term we added in this paper. The advantage of such a feature is that it prevents the neuron to specialize to an input pattern

that has occurred “by chance”.

It should be noted that the stability of the unspecialized synapse pattern (which corresponds to an unimodal weight distribution) can co-exist with the stability of specialized synapses (which corresponds to a bimodal weight distribution). This bistability is highlighted by the presence of an hysteresis behavior of the weights (see Fig 3b of paper V).

Reference

This paper is submitted under the following reference:

Toyoizumi T and Pfister JP and and Aihara K and Gerstner W. Optimality Model of Unsupervised Spike-Timing Dependent Plasticity: synaptic memory and weight distribution. *Accepted to Neural Computation.*

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

T. Toyozumi ^{1,2}, J.-P. Pfister ¹, K. Aihara ^{3,4} and W. Gerstner ¹

(1) School of Computer and Communication Sciences and Brain-Mind Institute,
Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland

(2) Department of Complexity Science and Engineering, Graduate School of
Frontier Sciences, University of Tokyo, Tokyo 153-8505, Japan

(3) Department of Information and Systems, Institute of Industrial Science,
University of Tokyo, 153-8505 Tokyo, Japan

(4) Aihara Complexity Modelling Project, ERATO, JST
3-23-5 Uehara, Shibuya-ku, Tokyo 151-0064, Japan

Manuscript accepted to *Neural Computation*

Summary

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

1 Introduction

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

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

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

The items in the above list are not necessarily exclusive, and the relative importance of a given aspect may vary from one subsystem to the next; for example, synaptic memory maintenance might be more important for a long-term memory system than for primary sensory cortices. Nevertheless, all of the above aspects seem to be important features of synaptic plasticity. However, the development of theoretical learning rules that exhibit all of the above properties has posed problems in the past. For example, traditional learning rules that have been proposed as an explanation of receptive field development (Bienenstock et al. 1982; Miller et al. 1989), exhibit a spontaneous separation of synaptic weights into two groups, even if the input shows no or only weak correlations. This is difficult to reconcile with experimental results in visual cortex of young rats where a unimodal distribution was found (Sjöström et al. 2001). Moreover model neurons that specialize early in development on one subset of features cannot readily re-adapt later on. Other learning rules, however, that exhibit a unimodal distribution of synaptic weights (Gütig et al. 2003) do not lead to a longterm stability of synaptic changes.

In this paper we want to show that all of the above features *(i)* - *(vi)* emerge nat-

urally in a theoretical model where we require only a limited number of objectives that will be formulated as postulates. In particular, we study how the conflicting demands on synaptic memory maintenance, plasticity, and distribution of synaptic synapses could be satisfied by our model. Even though the postulates are rather general and could be adapted to arbitrary neural systems, we had in mind excitatory synapses in neocortex or hippocampus and exclude inhibitory synapses and synapses in specialized systems such as the calyx of held in the auditory pathway. Our arguments are based on three postulates:

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

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

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

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

2 The Model

Neuron Model

We simulated a single stochastic point neuron model with $N = 100$ input synapses. Presynaptic spikes at synapse j are denoted by their arrival time t_j^f and evoke EPSPs with time course $\exp[-(t - t_j^f)/\tau_m]$ for $t \geq t_j^f$ where $\tau_m = 20\text{ms}$ is the membrane time constant. Recent experiments have shown that action potentials propagating back into the dendrite can partially suppress EPSPs measured at

the soma (Froemke et al. 2005). Since our model neuron has no spatial structure, we included EPSP suppression by a phenomenological amplitude factor $a(t_j^f - \hat{t})$ that depends on the time difference between presynaptic spike arrival and the spike trigger time \hat{t} of the last (somatic) action potential of the postsynaptic neuron.

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

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

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

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

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

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

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

Objective Function

Postulates **A** and **B** have been used previously (Toyoizumi et al. 2005) and lead to an optimality criterion $\mathcal{L}' = I - \gamma D$ where I is the mutual information between presynaptic input and postsynaptic output and D a measure of the distance of the mean firing rate of the neuron from its target rate. The parameter γ scales the importance of the information term I (postulate **A**) compared to the homeostatic term D (postulate **B**). It was shown that optimization of \mathcal{L}' by gradient ascent yields a synaptic update rule which shows sensitivity to correlations (see point (i) above), input selectivity (see point (ii) above), and depends on presynaptic firing rates (see point (v) above) (Toyoizumi et al. 2005). However, while the learning rule in (Toyoizumi et al. 2005) showed some dependence on spike timing, it did not (without additional assumptions) have the typical features of STDP as measured in vitro (point (vi) above); and exhibited, like earlier models (Bienenstock et al. 1982),

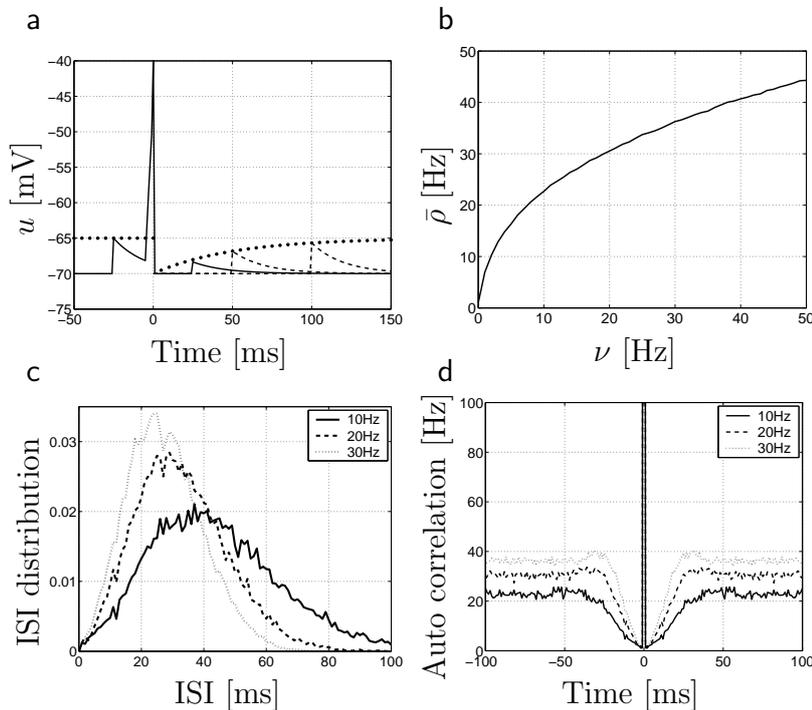


Figure 1: Properties of the stochastically spiking neuron model. (a) Neuron model and EPSP suppression. EPSPs arriving just after a postsynaptic spike at $t = 0$ are attenuated by a factor $(1 - e^{-t/\tau_a})$ and recover exponentially to their maximal amplitude w_j . (b) The output firing rate $\bar{\rho}$ of a neuron that receives stochastic spike input at rate ν at all 100 synapses. Each presynaptic spike evokes an EPSPs with maximal amplitude $w_j = 0.4\text{mV}$. (c) Interspike interval (ISI) distribution with input frequency $\nu = 10\text{Hz}$ (solid line), 20Hz (dashed line), and 30Hz (dotted line) at all 100 synapses. (d) Autocorrelation function of postsynaptic action potentials at an input frequency of 10Hz (solid line), 20Hz (dashed line), and 30Hz (dotted line).

spontaneous synaptic specialization, even for very weak input features, which is in contrast to point (iii) above.

In this earlier theoretical study, synaptic potentiation was artificially stopped at some upper bound w^{\max} (and synaptic depression was stopped at weight $w = 0$), so as to ensure that weights w stayed in a regime $0 \leq w \leq w^{\max}$. In the present paper we take the more realistic assumption that strong weights are more likely to show depression than weaker ones but do not impose a hard upper bound. Similarly, we require that adaptation speed is slowed down for very weak synapses, but do not impose a hard bound at zero weight. We will show that with these assumptions the resulting synaptic update rule shows properties of STDP (see point (vi) above), is suitable for memory retention (see point (iv) above) and leads to synaptic specialization when driven by strong input (see point (iii) above), while keeping the properties (i), (ii) and (v), that were found in (Toyoizumi et al. 2005).

To avoid hard upper bounds for the synapses, we use postulate **C1** and add a

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

$$\mathcal{L} = I - \gamma D - \lambda \Psi \quad (3)$$

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

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

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

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

where angular brackets $\langle \cdot \rangle_{Y,X}$ denote averaging over all combinations of input and output spike trains¹. Here $P(Y|X)$ is the conditional probability density of our stochastic neuron model to generate a specific spike train Y with (one or several) spike times $\{t_{post}^f\}$ during a trial of duration T given 100 known input spike trains X . This conditional probability density is given as a product of the instantaneous probabilities $\rho(t_{post}^f)$ of firing at the postsynaptic spike times $\{t_{post}^f\}$ and the probability of not firing elsewhere, i.e.,

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

Similarly, $P(Y)$ is the probability to generate the very same output spike train Y *not* knowing the input. Here ‘not knowing the input’ implies that we have to average

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

over all possible inputs so as to get the expected instantaneous firing density $\bar{\rho}(t)$ at time t . However, because of the EPSP suppression factor, the expected firing density will also depend on the last output spike before t . We therefore define

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

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

The second term is the homeostatic term

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

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

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

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

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

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

Synaptic update rule

We optimize the synaptic weights by gradient ascent

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

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

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

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

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

$$\frac{\partial \Psi}{\partial w_j} = w_j \langle n_j \rangle_X. \quad (12)$$

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

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

and

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

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

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

and

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

As a first interpretation we may say that c_j represents the causal correlation between input and output spikes (corrected for the expected correlation); and B^{post} is a function of postsynaptic quantities, namely, the output spikes y , current firing rate

ρ via the membrane potential u , average firing rate $\bar{\rho}$ and the target firing rate $\tilde{\rho}$. More precisely, B^{post} compares the actual output with the expected output and, modulated by a factor γ , the expected output with the target.

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

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

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

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

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

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

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

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

See Fig. 2 for an illustration of the dynamics of C_j and B^{post} . The last term has the form of an ‘weight decay’ term common in artificial neural networks (Hertz et al. 1991) and arises from the derivative of the weight-dependent cost term Ψ . The parameter λ is set such that $dw_j/dt = 0$ for large enough $|t^{pre} - t^{post}|$ in the STDP in vitro paradigm. A few steps of calculation (see Appendix A) yield $\lambda = 0.026$. In our simulations, we take $\tau_C = 100\text{ms}$ for the cut-off factor in Eq. (18).

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

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

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

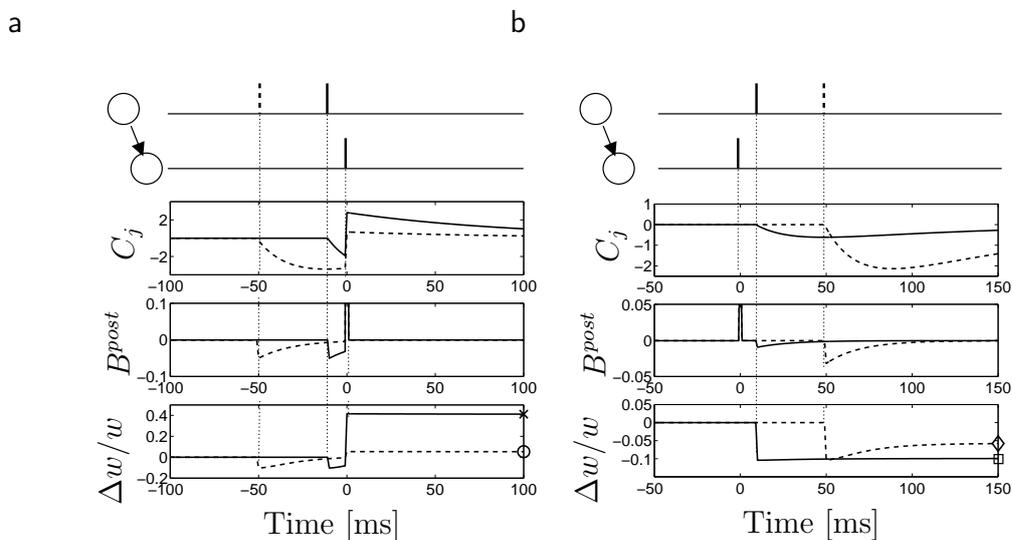


Figure 2: Illustration of the dynamics of B^{post} (Top), C_j (Middle) and $\Delta w/w = (w - w_{init})/w_{init}$ (Bottom). We always start from an initial weight $w_{init} = 4$ mV and induce postsynaptic firing at time $t^{post} = 0$. (a) Pre-before-post timing with $t^{pre} = -10$ ms (solid line) induces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

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

Stimulation paradigm

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

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

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

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

3 Results

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

As in a previous study based on postulates **A** and **B** (Toyoizumi et al. 2005), the optimization of synaptic weights can be understood as a synaptic update rule

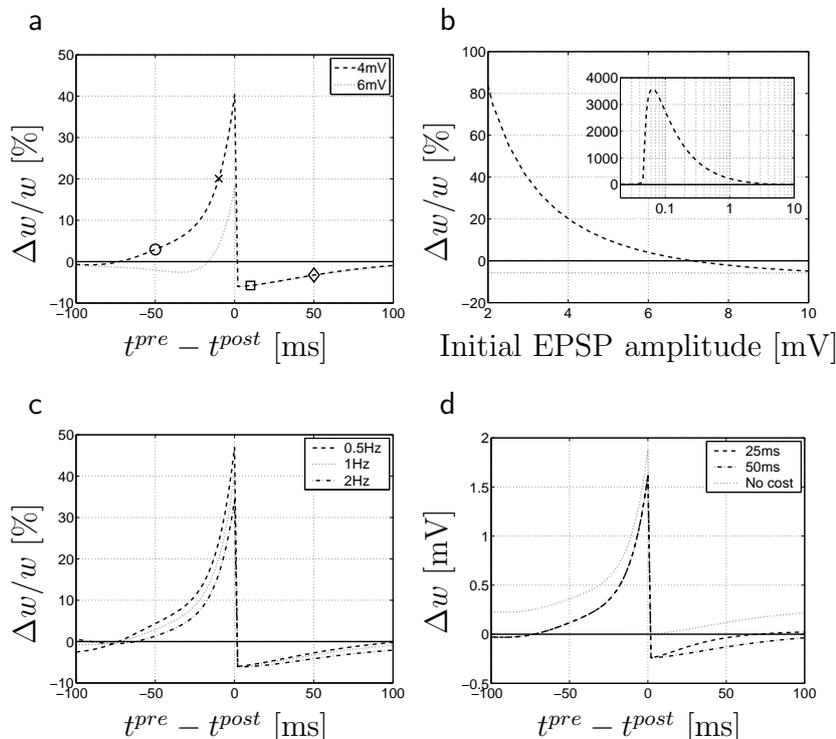


Figure 3: The synaptic update rule of the model shares features with STDP. (a) STDP function (percentage change of EPSP amplitude as a function of $t^{pre} - t^{post}$) determined using 60 pre-and-post spike-pairs injected at 1Hz. The initial EPSP amplitudes are 4mV (*dashed line*) and 6mV (*dotted*).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

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

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

The learning rule exhibits STDP.

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

For a comparison with experimental results, we used sixty pairs of pre- and postsynaptic spikes applied at a frequency of 1Hz and recorded the total change Δw in EPSP amplitude. The experiment is repeated with different spike timings and the result is plotted as a function of spike timing difference $t^{pre} - t^{post}$. As in experiments (Markram et al. 1997; Zhang et al. 1998; Bi and Poo 1998; Bi and Poo 2001; Sjöström et al. 2001), we find that synapses are potentiated if presynaptic spikes occur about 10 ms before a postsynaptic action potential, but are depressed if the timing is reversed. Compared to synapses with amplitudes in the range of 1 or 2 mV, synapses which are exceptionally strong show a reduced effect of potentiation for pre-before-post timing, or even depression (Fig. 3A and B), in agreement with experiments on cultured hippocampal neurons (Bi and Poo 1998). The shape of the STDP function depends only weakly on the stimulation frequency (Fig. 3C), even though a significant reduction of the potentiation amplitude with increasing

frequency can be observed.

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

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

Both unimodal and bimodal synapse distributions are stable

Under random spike arrival with a rate of 10Hz at all 100 synapses, synaptic weights show little variability with a typical EPSP amplitude in the range of 0.4 mV. This unspecific pattern of synapses stays stable even if 20 out of the 100 synapses are subject to a common rate modulation between 1 and 30Hz (Fig. 4A). However, if modulation of presynaptic firing rates becomes strong, the synapses develop rapidly a specific pattern with large values of weights at synapses with rate-modulated spike input and weak weights at those synapses that received input at fixed rates (*synaptic specialization*, see Fig. 4A and B), making the neuron highly selective to input at one group of synapses (*input selectivity*). Thus, our synaptic update rule is capable of selecting strong features in the input, but does also allow a stable unspecific pattern of synapses in case of weak input. This is in contrast to most other Hebbian learning rules, where unspecific patterns of synapses are unstable so that synaptic weights move spontaneously towards their upper or lower bounds (Miller et al. 1989; Miller and MacKay 1994; Gerstner et al. 1996; Kempter et al. 1999; Song et al. 2000).

After induction of synaptic specialization by strong modulation of presynaptic input, we reduced the rate modulation back to the value that previously led to an unspecific pattern of synapses. We found that the strong synapses remained strong and weak synapses remained weak, i.e., synaptic specialization was stable

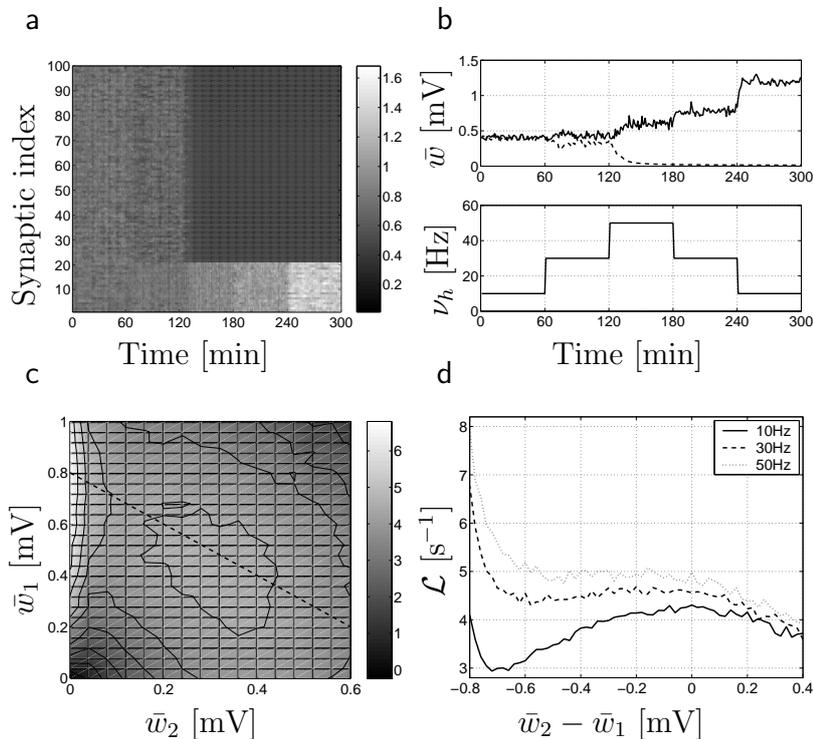


Figure 4: Potentiation and depression depend upon the presynaptic firing rate. 20 synapses of group **1** receive input with common rate modulation while the 80 synapses of group **2** (synapse index 21 - 100) receive Poisson input at a constant rate of 10Hz. The spike arrival rate in group **1** switches stochastically every 200ms between a low rate $\nu_l = 1$ Hz and a high rate ν_h taken as a parameter. (a) Evolution of synaptic weights as a function of time for different amplitudes of rate modulation, that is ν_h changes from 10Hz during the first hour to 30 Hz, then to 50 Hz, 30 Hz, and back to 10 Hz. During the first two hours of stimulation, an unspecific distribution of synapses remains stable even though a slight decrease of weights in group **2** can be observed when the stimulus switches to $\nu_h = 30$ Hz. A specialization of the synaptic pattern with large weights for synapses in group **1** is induced during the third hour of stimulation and remains stable thereafter. (b) Top: Mean synaptic weights (same data as in A) of group **1** (\bar{w}_1 , *solid line*) and group **2** (\bar{w}_2 , *dashed line*). Bottom: The stimulation paradigm, ν_h as a function of time. Note that at $\nu_h = 30$ Hz (2nd and 4th hour of stimulation) both an unspecific pattern of synapses with little difference between \bar{w}_1 and \bar{w}_2 (2nd hour, top) and a highly specialized pattern (4th hour, top, large difference between *solid* and *dashed lines*) are possible. (c) The value of the objective function \mathcal{L} (average value per second of time) in gray code as a function of the mean synaptic weight in group **1** (y-axis, \bar{w}_1) and group **2** (x-axis, \bar{w}_2) during stimulation with $\nu_h = 30$ Hz. Two maxima can be perceived, i.e., a broad maximum for the unspecific synapse pattern ($\bar{w}_2 \approx \bar{w}_1 \approx 0.4$ mV) and a pronounced elongated maximum for the specialized synapses pattern ($\bar{w}_1 \approx 0$; $\bar{w}_2 \approx 0.8$ mV). The dashed line indicates a 1-dimensional section (see d) through the two-dimensional surface. (d) Objective function \mathcal{L} as a function of the difference $\bar{w}_2 - \bar{w}_1$ between the mean synaptic weights in groups 2 and 1 along the line indicated in c. For $\nu_h = 30$ Hz (*dashed*, same data as in c) two maxima are visible, i.e., a broad maximum at $\bar{w}_2 - \bar{w}_1 \approx 0$ and a narrow, but higher maximum corresponding to the specialized synapse pattern to the very left of the graph. For $\nu_h = 50$ Hz (*dotted line*), the broad maximum at $\bar{w}_2 - \bar{w}_1 \approx 0$ disappears and only the specialized synapse pattern remains whereas for $\nu_h = 10$ Hz (*solid line*) the broad maximum of the unspecific synapse pattern dominates.

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

Retention of synaptic memories

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

Synaptic specialization was defined by a bimodality index that compared the distribution of EPSP amplitudes at synapses that received correlated input with those receiving uncorrelated input. For each of the two groups of synapses, we calculated the mean $\bar{w} = \langle w \rangle$ and the variance $\sigma^2 = \langle [w_j - \bar{w}]^2 \rangle$ i.e., \bar{w}_A and σ_A^2 for the group of synapses receiving correlated input and , \bar{w}_B and σ_B^2 for those receiving uncorrelated input. We then approximated the two distributions by Gaussian functions. The bimodality index depends on the overlap between the two Gaussians and is given by $b = 0.5 \left[\operatorname{erf} \left(\frac{\bar{w}_A - \hat{s}}{\sqrt{2}\sigma_A} \right) + \operatorname{erf} \left(\frac{\hat{s} - \bar{w}_B}{\sqrt{2}\sigma_B} \right) \right]$ where $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$ is the error function and \hat{s} is one of the two crossing points of the two Gaussians such that $\bar{w}_B < \hat{s} < \bar{w}_A$.

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

In order to study how synaptic memory retention depended on the induction paradigm, the experiment was then repeated with different values of the correlation index c that characterizes the spike-spike correlations during the induction period. For correlations $c < 0.1$ the two synaptic distributions are not well separated at the end of the induction period (bimodality index < 0.9), but well separated for

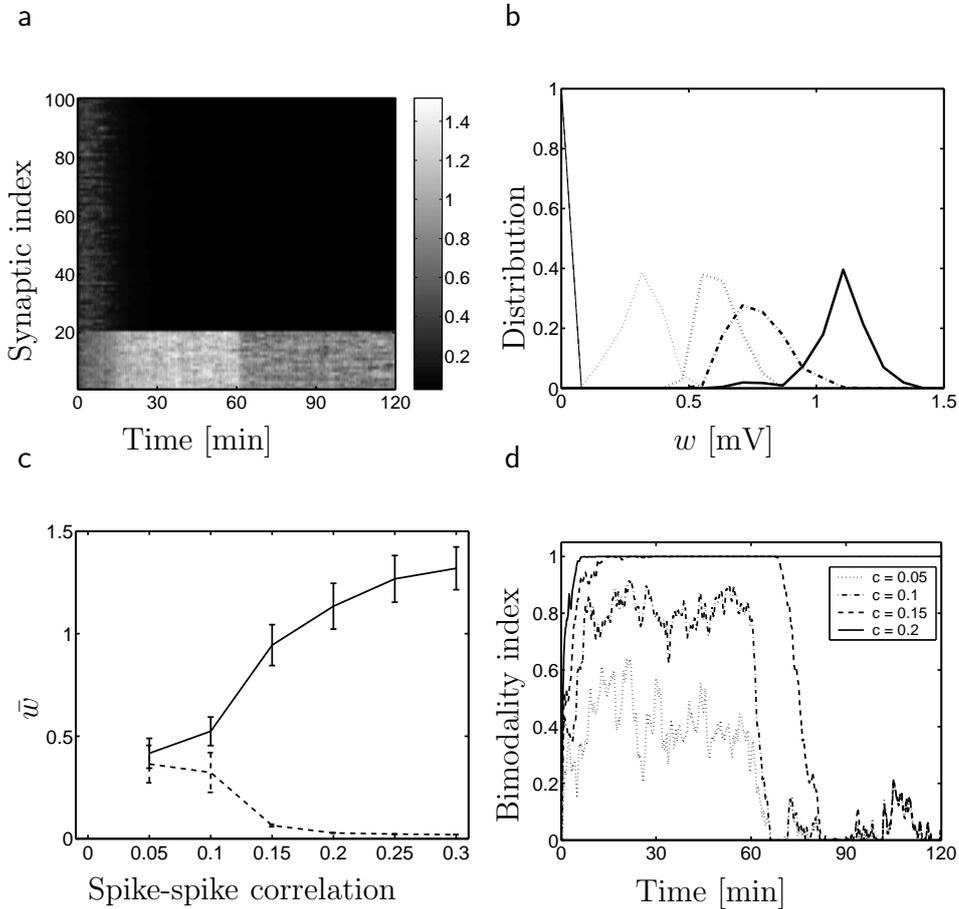


Figure 5: Synaptic memory induced by input with spike-spike correlations. (a) Evolution of 100 synaptic weights (vertical axis) as a function of time. During the first 60 minutes, synapses of group A ($j = 1, \dots, 20$) receive a Poisson stimulation of 10 Hz with correlated spike input ($c = 0.2$) and those of group B ($j = 21, \dots, 100$) receive uncorrelated spike input at 10 Hz. (b) Distribution of the EPSP amplitudes across the 100 synapses after $t = 5$ min. (*dotted line*), $t = 60$ min. (*solid line*) and $t = 120$ min. (*dot-dashed line*). The thick lines denote group A while the blue ones group B. (c) Mean EPSP amplitude of group A (*solid line*) and B (*dashed line*) at $t = 60$ min. for different values of correlation c of the input applied to group A. (d) Bimodality index b of the two groups of weights as a function of time. The memory induction by correlated spike input to group B stops at $t = 60$ min. Memory retention is studied during the following 60 minutes. A bimodality index close to one implies as for the case with $c = 0.2$ implies that synaptic memory is well retained.

hours (only the first hour is plotted in Fig. 5D).

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

Receptive field development

Synaptic plasticity is thought to be involved not only in memory (Hebb 1949), but also in the development of cortical circuits (Hubel and Wiesel 1962; Katz and Shatz 1996; von der Malsburg 1973; Bienenstock et al. 1982; Miller et al. 1989) and, possibly, cortical re-organization (Merzenich et al. 1984; Buonomano and Merzenich 1998). To study how our synaptic update rule would behave during development, we used a standard paradigm of input selectivity (Yeung et al. 2004), which is considered to be a simplified scenario of receptive field development. Our model neuron was stimulated by a Gaussian firing rate profile spanned across the 100 input synapses (Fig. 6A). The center of the Gaussian was shifted every 200 ms to an arbitrarily chosen presynaptic neuron. In order to avoid border effects, neuron number 100 was considered a neighbor of neuron number 1, i.e., we can visualize the presynaptic neurons as being located on a ring.

Nine postsynaptic neurons with slightly different initial values of synaptic weights received identical input from the same set of 100 presynaptic neurons. During one hour of stimulus presentation, six out of the nine neurons developed synaptic specialization leading to input selectivity. The optimal stimulus for these six neurons varies (Fig. 6B and C), so that any Gaussian stimulus at an arbitrary location excites at least one of the postsynaptic neurons. In other words, the six postsynaptic neurons have developed input selectivity with different but partially overlapping receptive fields. The distribution of synaptic weights for the selective neurons is bimodal with a first peak for very weak synapses (EPSP amplitudes less than 0.1 mV) and a second peak around EPSP amplitudes of 0.6 mV (Fig. 6D); the amplitude distribution of the unselective neurons is broader with a single peak at around 0.4mV.

The number of postsynaptic neurons showing synaptic specialization depends on the total stimulation time and the strength of the stimulus. If the stimulation

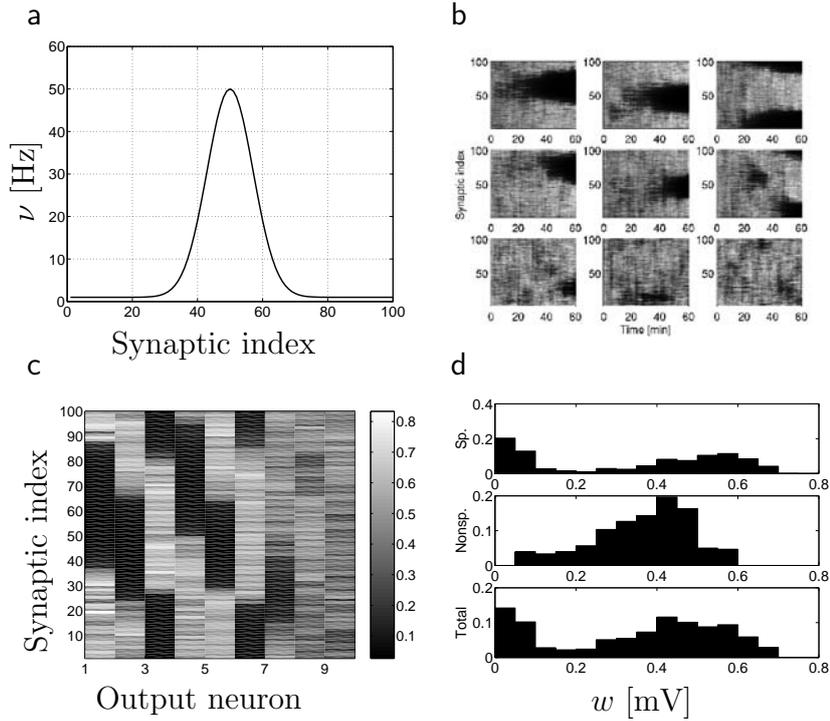


Figure 6: The synaptic update rule leads to input selectivity of the postsynaptic neuron. (a) Gaussian firing rate profile across the 100 presynaptic neurons. The center of the Gaussian is shifted randomly every 200 ms. Presynaptic neurons fire stochastically and send their spikes to nine postsynaptic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

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

4 Discussion

What can we and what can we not expect from optimality models?

Optimality models can be used to clarify concepts, but they are unable to make specific predictions about molecular implementations. In fact, the synaptic update rule derived in this paper shares functional features with STDP and classical LTP, but it is blind with respect to interesting questions such as the role of NMDA, Kainate, endocannabinoid, or CaMKII in the induction and maintenance of potentiation and depression (Bliss and Collingridge 1993; Bortolotto et al. 2003; Frey and Morris 1997; Lisman 2003; Malenka and Nicoll 1993; Sjöström et al. 2004). If molecular mechanisms are in the focus of interest, detailed ‘mechanistic’ models of synaptic plasticity (Senn et al. 2001; Yeung et al. 2004) should be preferred. On the other hand, the mere fact that similar forms of LTP or LTD seem to be implemented across various neural systems by different molecular mechanisms leads us to speculate that common functional roles of synapses are potentially more important for understanding synaptic dynamics than the specific way that these functions are implemented.

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

Correlations.

Hebb postulated an increase in synaptic coupling in case of repeated co-activation of pre- and postsynaptic neurons as a useful concept for memory storage in recurrent

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

Input selectivity.

During cortical development, cortical neurons develop input selectivity typically quantified as the width of receptive fields (Hubel and Wiesel 1962). As illustrated in the scenario of Fig. 6, our synaptic update rule shows input selectivity and stands hence in the research tradition of many other studies (von der Malsburg 1973; Bienenstock et al. 1982; Miller et al. 1989). Input selectivity in our model arises through the combination of the correlation sensitivity of synapses discussed above with the homeostatic term D in Eq. (3) (Toyoizumi et al. 2005). Since the homeostatic term keeps the mean rate of the postsynaptic neuron close to a target rate, it leads effectively to a normalization of the total synaptic input similar to the sliding threshold mechanism in the Bienenstock-Cooper-Munro rule (Bienenstock et al. 1982). Normalization of the total synaptic input through firing rate stabilization has also been seen in previous STDP models (Song et al. 2000; Kempter et al. 1999; Kempter et al. 2001). Its effect is similar to explicit normalization of synaptic weights which is a well-known mechanism to induce input selectivity (von der Malsburg 1973; Miller and MacKay 1994), but our model does not need an explicit normalization step. So far we studied only a single or a small number of independent postsynaptic neurons, but we expect that, as in many other studies, e.g., (Erwin et al. 1995; Song and Abbott 2001; Cooper et al. 2004), our synaptic update rule would yield feature maps if applied to a network of many weakly interacting cortical neurons.

Unimodal versus bimodal synapse distributions.

In several previous models of rate-based or spike-timing based synaptic dynamics, synaptic weights evolved *always* towards a bimodal distribution with some synapses close to zero and others close to maximal weight (Miller et al. 1989; Miller and MacKay 1994; Kempter et al. 1999; Song and Abbott 2001; Toyoizumi et al. 2005). Thus synapses specialize on certain features of the input, even if the input has weak or no correlation at all which seems questionable from a functional point of view and which disagrees with experimentally found distributions of EPSP amplitudes in rat visual cortex (Sjöström et al. 2001). As shown in Figs. 4 — 6, an unspecific

pattern of synapses with a broad distribution of EPSP amplitudes is stable with our synaptic update rule if correlations in the input are weak. Synaptic specialization (bimodal distribution) develops only if synaptic inputs show a high degree of correlations, either on the level of spikes or firing rates. Thus neurons only specialize on highly significant input features, and not in response to noise. A similar behavior was noted in two recent models on spike-timing dependent plasticity (Gütig et al. 2003; Yeung et al. 2004). Going beyond those studies, we also demonstrated stability of the specialized synapse distribution over some time, even if the amount of correlation through rate-modulation is reduced after induction of synaptic specialization stimulation (Fig. 4). Thus, for the same input characteristics our model can show unimodal or bimodal distribution of synaptic weight, depending on the stimulation history. Moreover, we note that in the state of bimodal distribution, the EPSP amplitude at the depressed synapses are so small, that in an experimental setting they could easily remain undetected or classified as ‘silent’ synapse (Kullmann 2003). A large proportion of silent synapses has been previously shown to be consistent with optimal memory storage (Brunel et al. 2004). Furthermore, our results show that in a given population of cells, neurons with specialized synapses can coexist with others that have a broad and unspecific synapse distribution. We speculate that these non-specialized neurons could then be recruited later for new stimuli, as hypothesized in earlier models of neural networks (Grossberg 1987).

Synaptic memory.

In the absence of presynaptic input, synapses in our model do not change significantly. Furthermore, our results show that even in the presence of random pre- and postsynaptic firing activity, synaptic memories can be retained over several hours, even though a slow decay occurs. Essential for long-term memory maintenance under random spike arrival is the reduced adaptation speed for small values of synaptic weights as formulated in postulate **C2**. This postulate is similar in spirit to a recent theory of Fusi et al. (2005), with two important differences. Firstly, we work with a continuum of synaptic states (characterized by the value of w) whereas Fusi et al. assume a cascade of a finite number of discrete internal synaptic states where transitions are uni-directional and characterized by different time constants (Fusi et al. 2005). The scheme of Fusi et al. guarantees slow (i.e., not exponential) decay of memories, whereas in our case decay is always exponential even if the time constant is weight-dependent. Secondly, Fusi et al. use a range of different time constants for both the potentiated and the unpotentiated state, whereas in our model it is sufficient to have a slower time constant for the unpotentiated state only. If the change of unpotentiated synapses is slow compared to homeostatic regulation of mean firing rate, then the maintenance of the strong synapses is given by homeostasis (and also supported by information maximization).

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

1998; Lisman 2003).

Rate dependence.

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

STDP.

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

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

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

role in STDP since it shifts the STDP function to more negative values.

How realistic is a weight-dependent cost term?

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

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

An important aspect of our cost term is that only synapses that have recently been activated are at risk regarding weight decay. We speculate that the weight-dependent cost term could, in a loose sense, be related to the amount of plasticity factors that synapses require and compete for during the first hours of synaptic maintenance (Fonseca et al. 2004). According to the synaptic tagging hypothesis (Frey and Morris 1997), only those synapses that have been activated in the recent past compete for plasticity factors, while unpotentiated synapses do not suffer from decay (Fonseca et al. 2004). We emphasize that such a link of our cost term to the competition for plasticity factors is purely hypothetical. Many relevant details of tagging and competition for synaptic maintenance are omitted in our approach.

Predictions and experimental tests.

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

A second test of our model concerns the pattern of synaptic weights converging on the same postsynaptic neuron. Our results suggest that early in development most neurons would show an unspecific synapse pattern, i.e., a distribution of EPSP

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

Limitations and extensions

We would like to emphasize that properties of our synaptic update rules have so far only been tested for single neurons in an unsupervised learning paradigm. Extensions are possible in several directions. Firstly, instead of single neurons, a large recurrent network could be considered. This could on one side further our understanding of the model properties in the context of cortical map development (Erwin et al. 1995), on the other side scrutinize the properties of the synaptic update rule as a functional memory in recurrent networks (Amit and Fusi 1994). Secondly, instead of unsupervised learning where the synaptic update rule treats all stimuli alike whether they are behaviorally relevant or not, a reward-based learning scheme could be considered (Dayan and Abbott 2001; Seung 2003). Behaviorally relevant situations can be taken into account by optimizing reward instead of information transmission (Schultz et al. 1997).

Acknowledgments

This research is supported by the Japan Society for the Promotion of Science and Grant-in-Aid for JSPS fellows 03J11691 (T.T.), the Swiss National Science Foundation 200020-108097/1 (J.P.P), and Grant-in-Aid for Scientific Research on Priority Areas 17022012 from MEXT of Japan (K.A.).

A Determination of the Parameter λ

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

presynaptic spike at t_j^f , we have

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

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

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

Hence, we find the effect of a presynaptic spike as

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

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

B Weight Dependent Evaluation of the Optimality Criterion

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

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

References

- Abarbanel, H., R. Huerta, and M. Rabinovich (2002). Dynamical model of long-term synaptic plasticity. *Proc. Natl. Academy of Sci. USA* 59(10137-10143).
- Abarbanel, H., S. Talathi, L. Gibb, and M. Rabinovich (2005). Synaptic plasticity with discrete state synapses. *to appear xx(xx)*.
- Abraham, W., B. Logan, J. Greenwood, and M. Dragunow (2002). Induction and experience-dependent consolidation of stable long-term potentiation lasting months in the hippocampus. *J. Neuroscience* 22, 9626 – 9634.
- Amit, D. and S. Fusi (1994). Learning in neural networks with material synapses. *Neural Computation* 6, 957–982.
- Artola, A., S. Bröcher, and W. Singer (1990). Different voltage dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347, 69–72.
- Atick, J. and A. Redlich (1990). Towards a theory of early visual processing. *Neural Computation* 4, 559–572.
- Barlow, H. (1956). Retinal noise and absolute threshold. *J. Opt. Soc. Am.* 46, 634–639.
- Barlow, H. (1980). The absolute efficiency of perceptual decisions. *Phil. Trans. R. Soc. Lond. Ser. B* 290, 71–82.
- Barlow, H. and W. Levick (1969). 3 factors limiting reliable detection of light by retinal ganglion cells of cat. *J. Physiology (London)* 200, 1–24.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenbluth (Ed.), *Sensory Communication*, pp. 217–234. MIT Press.
- Barlow, H. B. (1989). Unsupervised learning. *Neural. Comp.* 1, 295–311.
- Bell, A. and T. Sejnowski (1995). An information maximization approach to blind separation and blind deconvolution. *Neural Computation* 7, 1129–1159.
- Bell, A. J. and L. C. Parra (2005). Maximising sensitivity in a spiking network. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 121–128. Cambridge, MA: MIT Press.
- Bi, G. and M. Poo (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bi, G. and M. Poo (2001). Synaptic modification of correlated activity: Hebb’s postulate revisited. *Ann. Rev. Neurosci.* 24, 139–166.
- Bienenstock, E., L. Cooper, and P. Munroe (1982). Theory of the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* 2, 32–48. reprinted in Anderson and Rosenfeld, 1990.
- Bliss, T. and T. Lomo (1973). Long-lasting potentiation of synaptic transmission in the dendate area of anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 351–356.
- Bliss, T. V. P. and G. L. Collingridge (1993). A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39.
- Bohte, S. M. and M. C. Mozer (2005). Reducing spike train variability: A computational theory of spike-timing dependent plasticity. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 201–208. Cambridge, MA: MIT Press.
- Bortolotto, Z. A., S. Lauri, J. T. R. Isaac, and G. L. Collingridge (2003). Kainate receptors and the induction of mossy fibre long-term potentiation. *Phil. Trans. R. Soc. Lond B: Biological Sciences* 358(657 - 666).
- Britten, K., M. Shadlen, W. Newsome, and J. Movshon (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *J. Neuroscience* 12, 4745–4765.

- Brunel, N., V. Hakim, P. Isope, J.-P. Nadal, and B. Barbour (2004). Optimal information storage and the distribution of synaptic weights: Perceptron versus purkinje cell. *Neuron* *43*, 745–757.
- Buonomano, D. and M. Merzenich (1998). Cortical plasticity: From synapses to maps. *Annual Review of Neuroscience* *21*, 149–186.
- Castellani, G., E. Quinlan, L. Shouval, and H. Cooper (2001). A biophysical model of bidirectional synaptic plasticity: dependence on ampa and nmda receptors. *Proc. Natl. Acad. Sci. USA* *98*, 12772–12777.
- Chechik, G. (2003). Spike-timing-dependent plasticity and relevant mutual information maximization. *Neural Computation* *15*, 1481–1510.
- Cooper, L., N. Intrator, B. Blais, and H. Z. Shouval (2004). *Theory of cortical plasticity*. Singapore: World Scientific.
- Cover, T. and J. Thomas (1991). *Elements of Information Theory*. New York: Wiley.
- D. Liao, N.A. Hessler, R. M. (1995). Activation of postsynaptically silent synapses during pairing-induced ltp in cal region of hippocampal slice. *Nature* *375*, 400–404.
- Dan, Y. and M. Poo (2004). Spike timing-dependent plasticity of neural circuits. *Neuron* *44*, 23–30.
- Dayan, P. and L. F. Abbott (2001). *Theoretical Neuroscience*. Cambridge: MIT Press.
- de Ruyter van Steveninck, R. R. and W. Bialek (1995). Reliability and statistical efficiency of a blowfly movement-sensitive neuron. *Phil. Trans. R. Soc. Lond. Ser. B* *348*, 321–340.
- Dudek, S. M. and M. F. Bear (1992). Homosynaptic long-term depression in area cal of hippocampus and effects of n-methyl-d-aspartate receptor blockade. *Proc. Natl. Acad. Sci. USA* *89*, 4363–4367.
- Erwin, E., K. Obermayer, and K. Schulten (1995). Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comput.* *7*, 425–468.
- Fairhall, A. L., G. Lewen, W. Bialek, and R. van Steveninck (2001). Efficiency and ambiguity in an adaptive neural code. *Nature* *412*, 787–792.
- Fonseca, R., U. Nägerl, R. Morris, and T. Bonhoeffer (2004). Competition for memory: Hippocampal ltp under the regimes of reduced protein synthesis. *Neuron* *44*, 1011–1020.
- Frey, U. and R. Morris (1997). Synaptic tagging and long-term potentiation. *Nature* *385*, 533 – 536.
- Froemke, R. C., M.-M. Poo, and Y. Dan (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* *434*, 221–225.
- Fusi, S. (2002). Hebbian spike-driven synaptic plasticity for learning patterns of mean firing rates. *Biol. Cybern.* *87*, 459–470.
- Fusi, S., M. Annunziato, D. Badoni, A. Salamon, and D.J.Amit (2000). Spike-driven synaptic plasticity: theory, simulation, vlsi implementation. *Neural Computation* *12*, 2227–2258.
- Fusi, S., P. Drew, and L. Abbott (2005). Cascade models of synaptically stored memories. *Neuron* *45*, 599–611.
- Gerstner, W., R. Kempter, J. L. van Hemmen, and H. Wagner (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* *383*, 76–78.
- Gerstner, W. and W. K. Kistler (2002). *Spiking Neuron Models*. Cambridge UK: Cambridge University Press.
- Grossberg, S. (1987). *The Adaptive Brain I*. Elsevier.
- Gütig, R., R. Aharonov, S. Rotter, and H. Sompolinsky (2003). Learning input correlations through non-linear temporally asymmetry hebbian plasticity. *J. Neuroscience* *23*, 3697–3714.
- Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley.

- Hertz, J., A. Krogh, and R. G. Palmer (1991). *Introduction to the Theory of Neural Computation*. Redwood City CA: Addison-Wesley.
- Holmes, W. R. and W. B. Levy (1990). Insights into associative long-term potentiation from computational models of nmda receptor-mediated calcium influx and intracellular calcium concentration changes. *J. Neurophysiol.* *63*, 1148–1168.
- Hubel, D. H. and T. N. Wiesel (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiol. (London)* *160*, 106–154.
- Isaac, J. T., R. A. Nicoll, and R. C. Malenka (1995). Evidence for silent synapses: Implications for the expression of ltp. *Neuron* *15*, 427–434.
- Karmarkar, U. and D. Buonomano (2002). A model of spike-timing dependent plasticity: one or two coincidence detectors. *J. Neurophysiology* *88*, 507–513.
- Katz, L. C. and C. J. Shatz (1996). Synaptic activity and the construction of cortical circuits. *Science* *274*, 1133–1138.
- Kelso, S. R., A. H. Ganong, and T. H. Brown (1986). Hebbian synapses in hippocampus. *Proc. Natl. Acad. Sci. USA* *83*, 5326–5330.
- Kempster, R., W. Gerstner, and J. L. van Hemmen (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* *59*, 4498–4514.
- Kempster, R., W. Gerstner, and J. L. van Hemmen (2001). Intrinsic stabilization of output rates by spike-based hebbian learning. *Neural Computation* *13*, 2709–2741.
- Kempster, R., C. Leibold, H. Wagner, , and J. L. van Hemmen (2001). Formation of temporal-feature maps by axonal propagation of synaptic learning. *Proc. Natl. Acad. of Sci. USA* *98*, 4166–4171.
- Kistler, W. M. and J. L. van Hemmen (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput.* *12*, 385–405.
- Kullmann, D. M. (2003). Silent synapses: what are they telling us about long-term potentiation? *Phil. Trans. R. Soc. Lond B: Biological Sciences* *358*, 727 – 733.
- Laughlin, S. (1981). A simple coding procedure enhances a neurons information capacity. *Z. Naturforschung* *36*, 910–912.
- Laughlin, S., R. R. deRuyter van Steveninck, and J. Anderson (1998). The metabolic cost of neural information. *Nature Neuroscience* *1*(36-41).
- Levy, W. and R. Baxter (2002). Energy-efficient neuronal computation via quantal synaptic failures. *J. Neuroscience* *22*, 4746–4755.
- Linsker, R. (1986). From basic network principles to neural architecture: emergence of spatial-opponent cells. *Proc. Natl. Acad. Sci. USA* *83*, 7508–7512.
- Linsker, R. (1989). How to generate ordered maps by maximizing the mutual information between input and output signals. *Neural Computation* *1*(3), 402–411.
- Lisman, J. (1989). A mechanism for hebb and anti-hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA* *86*, 9574–9578.
- Lisman, J. (2003). Long-term potentiation: outstanding questions and attempted synthesis. *Phil. Trans. R. Soc. Lond B: Biological Sciences* *358*, 829 – 842.
- Malenka, R. C., J. Kauer, R. Zucker, and R. A. Nicoll (1988). Postsynaptic calcium is sufficient for potentiation of hippocampal synaptic transmission. *Science* *242*, 81–84.
- Malenka, R. C., J. A. Kauer, D. J. Perkel, M. D. Mauk, and P. T. Kelly (1989). An essential role for postsynaptic calmodulin and protein kinase activity in long-term potentiation. *Nature* *340*, 554–557.
- Malenka, R. C. and R. A. Nicoll (1993). Nmda-receptor-dependent plasticity: multiple forms and mechanisms. *Trends Neurosci.* *16*, 480–487.
- Markram, H., J. Lübke, M. Frotscher, and B. Sakmann (1997). Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSP. *Science* *275*, 213–215.

- Merzenich, M., R. Nelson, M. Stryker, M. Cynader, A. Schoppmann, and J. Zook (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *J. Comparative Neurology* 224, 591–605.
- Miller, K., J. B. Keller, and M. P. Stryker (1989). Ocular dominance column development: analysis and simulation. *Science* 245, 605–615.
- Miller, K. D. and D. J. C. MacKay (1994). The role of constraints in hebbian learning. *Neural Computation* 6, 100–126.
- Nadal, J.-P. and N. Parga (1997). Redundancy reduction and independent component analysis: Conditions on cumulants and adaptive approaches. *Neural Computation* 9, 1421–1456.
- Oja, E. (1982). A simplified neuron model as a principal component analyzer. *J. Mathematical Biology* 15, 267–273.
- Petersen, C., R. Malenka, R. Nicoll, and J. Hopfield (1998). All-or-none potentiation of ca3-ca1 synapses. *Proc. Natl. Acad. Sci. USA* 95, 4732–4737.
- Rubin, J., D. D. Lee, and H. Sompolinsky (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Physical Review Letters* 86, 364–367.
- Schultz, W., P. Dayan, and R. Montague (1997). A neural substrate for prediction and reward. *Science* 275, 1593–1599.
- Senn, W., M. Tsodyks, and H. Markram (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation* 13, 35–67.
- Seung, H. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* 40, 1063–1073.
- Sjöström, P. and S. Nelson (2002). Spike timing, calcium signals and synaptic plasticity. *Current Opinion in Neurobiology* 12, 305–314.
- Sjöström, P., G. Turrigiano, and S. Nelson (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164.
- Sjöström, P., G. Turrigiano, and S. Nelson (2004). Endocannabinoid-dependent neocortical layer-5 ltd in the absence of postsynaptic spiking. *J. Neurophysiol.* 92, 3338–3343.
- Song, S. and L. Abbott (2001). Column and map development and cortical re-mapping through spike-timing dependent plasticity. *Neuron* 32(339-350).
- Song, S., K. Miller, and L. Abbott (2000). Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience* 3, 919–926.
- Sutton, R. and A. Barto (1998). *Reinforcement learning*. MIT Press, Cambridge.
- Toyoizumi, T., J.-P. Pfister, K. Aihara, and W. Gerstner (2005). Generalized bienenstock-cooper-munro rule for spiking neurons that maximizes information transmission. *Proc. National Academy Sciences (USA)* 102, 5239–5244.
- Toyoizumi, T., J.-P. Pfister, K. Aihara, and W. Gerstner (2005a). Spike-timing dependent plasticity and mutual information maximization for a spiking neuron model. In L. K. Saul, Y. Weiss, and L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17*, pp. 1409–1416. Cambridge, MA: MIT Press.
- Turrigiano, G. and S. Nelson (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience* 5, 97–107.
- van Rossum, M. C. W., G. Q. Bi, and G. G. Turrigiano (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience* 20, 8812–8821.
- von der Malsburg, C. (1973). Self-organization of orientation selective cells in the striate cortex. *Kybernetik* 14, 85–100.
- Yeung, L. C., H. Shouval, B. Blais, and L. Cooper (2004). Synaptic homeostasis and input selectivity follow from a model of calcium dependent plasticity. *Proc. Nat. Acad. Sci. USA*, 101, 14943–14948.

- Zador, A., C. Koch, and T. H. Brown (1990). Biophysical model of a hebbian synapse. *Proc. Natl. Acad. Sci.* *87*, 6718–6722.
- Zhang, L., H. Tao, C. Holt, W.A.Harris, and M.-M. Poo (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature* *395*, 37–44.
- Zhou, Q., H. Tao, and M. Poo (2003). Reversal and stabilization of synaptic modifications in the developping visual system. *Science* *300*, 1953–1957.

2.6 Paper VI

Summary

We have seen in the first five papers different learning rules that maximize a given objective function that should capture the functional role of the synapse. They can be classified as optimal models of STDP (see section 1.2.3). The approach in this paper (and in the next one) is different. The goal is not to determine the functional role of the synapse, but to develop a phenomenological model (see section 1.2.1) which is as simple as possible, but captures most of the experimental results of STDP.

Most of the actual phenomenological models about STDP only consider the effect of spike *pairs* (one pre- and one postsynaptic spike) (Gerstner et al. 1996; Kempter et al. 1999; Song et al. 2000; Gütig et al. 2003). Here, we present a learning rule that also takes into account the effect of *triplets* of spike, i.e. 1 pre and 2 postsynaptic spikes. With this triplet learning rule, it is possible to fit the data of Sjöström, Turrigiano, and Nelson (2001) which could not be fitted by simple pair-based learning rules. Moreover, we show a link between the triplet learning rule and the BCM learning rule.

With this latter property, the triplet learning rule is not only a phenomenological learning rule, but could also be classified as an optimal learning rule in the sense that it maximizes the input selectivity (see definition of selectivity in section 1.2.3).

In this paper, we also discuss the implications of the type of interactions between spikes. Indeed, if a presynaptic spike interacts only with the previous one (see the *Nearest-Past-Spike* interactions on Fig. 1.3 of section 1.2.1) instead of interacting with all the previous ones (*All-to-All*), it induces some non-linearities that can be useful (or harmful) for data fitting. For the experimental data we consider here, both interaction schemes are possible.

Reference

This paper has been presented at the *Arc Lémanique Neuroscience Retreat* (September 30 - October 1, 2005, Les Diablerets, Switzerland) and at the NIPS 2005 conference (December 5 - 10, 2005, Vancouver, Canada) and has been published after a peer-review process under the following reference:

Pfister JP and Gerstner W. Beyond Pair-Based STDP: a Phenomenological Rule for Spike Triplet and Frequency Effects. *Advances in Neural Information Processing Systems 18*, edited by Y. Weiss and B. Schölkopf and J. Platt, MIT Press, Cambridge MA, 1083-1090, 2006.

Beyond Pair-Based STDP: a Phenomenological Rule for Spike Triplet and Frequency Effects

Jean-Pascal Pfister and Wulfram Gerstner

School of Computer and Communication Sciences
and Brain-Mind Institute,

Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne
{jean-pascal.pfister, wulfram.gerstner}@epfl.ch

Abstract

While classical experiments on spike-timing dependent plasticity analyzed synaptic changes as a function of the timing of *pairs* of pre- and postsynaptic spikes, more recent experiments also point to the effect of spike *triplets*. Here we develop a mathematical framework that allows us to characterize timing based learning rules. Moreover, we identify a candidate learning rule with five variables (and 5 free parameters) that captures a variety of experimental data, including the dependence of potentiation and depression upon pre- and postsynaptic firing frequencies. The relation to the Bienenstock-Cooper-Munro rule as well as to some timing-based rules is discussed.

1 Introduction

Most experimental studies of Spike-Timing Dependent Plasticity (STDP) have focused on the timing of spike pairs [1, 2, 3] and so do many theoretical models. The spike-pair based models can be divided into two classes: either all pairs of spikes contribute in a homogeneous fashion [4, 5, 6, 7, 8, 9, 10] (called ‘all-to-all’ interaction in the following) or only pairs of ‘neighboring’ spikes [11, 12, 13] (called ‘nearest-spike’ interaction in the following); cf. [14, 15]. Apart from these phenomenological models, there are also models that are somewhat closer to the biophysics of synaptic changes [16, 17, 18, 19].

Recent experiments have furthered our understanding of timing effects in plasticity and added at least two different aspects: firstly, it has been shown that the mechanism of potentiation in STDP is different from that of depression [20] and secondly, it became clear that not only the timing of pairs, but also of triplets of spikes contributes to the outcome of plasticity experiments [21, 22].

In this paper, we introduce a learning rule that takes these two aspects partially into account in a simple way. Depression is triggered by *pairs* of spikes with *post-before-pre* timing, whereas potentiation is triggered by *triplets* of spikes consisting of 1 pre- and 2 postsynaptic spikes. Moreover, in our model the pair-based depression includes an explicit dependence upon the mean postsynaptic firing rate. We show that such a learning rule accounts for two important stimulation paradigms:

P1 (Relative Spike Timing): *Both the pre- and postsynaptic spike trains consist of a burst*

of N spikes at regular intervals T , but the two spike trains are shifted by a time $\Delta t = t^{\text{post}} - t^{\text{pre}}$.

The total weight change is a function of the relative timing Δt (this gives the standard STDP function), but also a function of the firing frequency $\rho = 1/T$ during the burst; cf. Fig. 1A (data from L5 pyramidal neurons in visual cortex).

P2 (Poisson Firing): *The pre- and postsynaptic spike trains are generated by two independent Poisson processes with rates ρ_x and ρ_y respectively.*

Protocol P2 has less experimental support but it helps to establish a relation to the Bienenstock-Cooper-Munro (BCM) model [23]. To see that relation, it is useful to plot the weight change as a function of the postsynaptic firing rate, i.e., $\Delta w \propto \phi(\rho_y)$ (cf. Fig 1B). Note that the function ϕ has only been measured indirectly in experiments [24, 25].

We emphasize that in the BCM model,

$$\Delta w = \rho_x \phi(\rho_y, \bar{\rho}_y) \quad (1)$$

the function ϕ depends not only on the current firing rate ρ_y , but also on the *mean* firing rate $\bar{\rho}_y$ averaged over the recent past which has the effect that the threshold between depression and potentiation is not fixed but dynamic. More precisely, this threshold θ depends non-linearly on the mean firing rate $\bar{\rho}_y$:

$$\theta = \alpha \bar{\rho}_y^p, \quad p > 1 \quad (2)$$

with parameters α and p . Previous models of STDP have already discussed the relation of STDP to the BCM rule [16, 12, 17, 26], but none of these seems to be completely satisfactory as discussed in Section 4. We will also compare our results to the rule of [21] which was together with the work of [16] amongst the first triplet rules to be proposed.

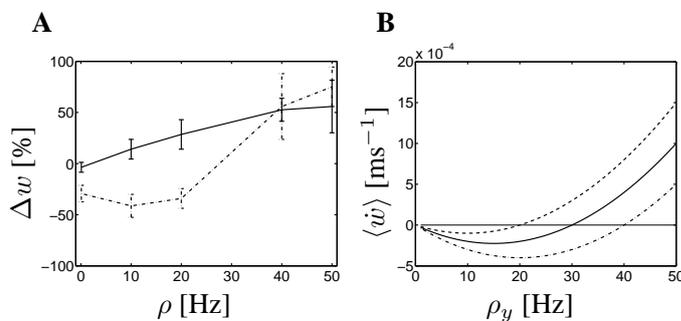


Figure 1: **A.** Weight change in an experiment on cortical synapses using pairing protocol (P1) (solid line: $\Delta t = 10$ ms, dot-dashed line $\Delta t = -10$ ms) as a function of the frequency ρ . Figure redrawn from [11]. **B.** Weight change in protocol P2 according to the BCM rule for $\theta = 20, 30, 40$ Hz.

2 A Framework for STDP

Several learning rules in the modeling literature can be classified according to the two criteria introduced above: (i) all-to-all interaction vs. nearest spike interaction; (ii) pair-based vs. triplet based rules. Point (ii) can be elaborated further in the context of an expansion (pairs, triplets, quadruplets, ... of spikes) that we introduce now.

2.1 Volterra Expansion ('all-to-all')

For the sake of simplicity, we assume that weight changes occur at the moment of presynaptic spike arrival or at the moment of postsynaptic firing. The direction and amplitude

of the weight change depends on the configuration of spikes in the presynaptic spike train $X(t) = \sum_k \delta(t - t_x^k)$ and the postsynaptic spike train $Y(t) = \sum_k \delta(t - t_y^k)$. With some arbitrary functionals $F[X, Y]$ and $G[X, Y]$, we write (see also [8])

$$\dot{w}(t) = X(t)F[X, Y] + Y(t)G[X, Y] \quad (3)$$

Clearly, there can be other neuronal variables that influence the synaptic dynamics. For example, the weight change can depend on the current weight value w [8, 15, 10], the Ca^{2+} concentration [17, 19], the depolarization [25, 27, 28], the mean postsynaptic firing rate $\bar{\rho}_y(t)$ [23],... Here, we will consider only the dependence upon the history of the pre- and postsynaptic firing times and the mean postsynaptic firing rate $\bar{\rho}_y$. Note that even if $\bar{\rho}_y$ depends via a low-pass filter $\tau_\rho \dot{\bar{\rho}}_y = -\bar{\rho}_y + Y(t)$ on the past spike train Y of the postsynaptic neuron, the description of the problem will turn out to be simpler if the mean firing rate is considered as a separate variable. Therefore, let us write the instantaneous weight change as

$$\dot{w}(t) = X(t)F([X, Y], \bar{\rho}_y(t)) + Y(t)G([X, Y], \bar{\rho}_y(t)) \quad (4)$$

The goal is now to determine the simplest functionals F and G that would be consistent with the experimental protocols $P1$ and $P2$ introduced above. Since the functionals are unknown, we perform a Volterra expansion of F and G in the hope that a small number of low-order terms are sufficient to explain a large body of experimental data. The Volterra expansion [29] of the functional G can be written as¹

$$\begin{aligned} G([X, Y]) &= G_1^y + \int_0^\infty G_2^{xy}(s)X(t-s)ds + \int_0^\infty G_2^{yy}(s)Y(t-s)ds \\ &+ \int_0^\infty \int_0^\infty G_3^{xxy}(s, s')X(t-s)X(t-s')ds'ds \\ &+ \int_0^\infty \int_0^\infty \mathbf{G}_3^{xyy}(s, s')X(t-s)Y(t-s')ds'ds \\ &+ \int_0^\infty \int_0^\infty G_3^{yyy}(s, s')Y(t-s)Y(t-s')ds'ds + \dots \end{aligned} \quad (5)$$

Similarly, the expansion of F yields

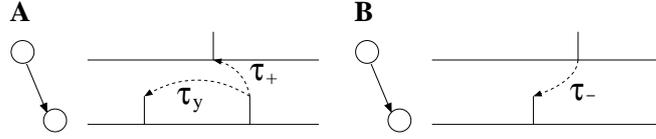
$$F([X, Y]) = F_1^x + \int_0^\infty F_2^{xx}(s)X(t-s)ds + \int_0^\infty \mathbf{F}_2^{xy}(s)Y(t-s)ds + \dots \quad (6)$$

Note that the upper index in functions represents the type of interaction. For example, G_3^{xyy} (in bold face above) refers to a triplet interaction consisting of 1 pre- and 2 postsynaptic spikes. Note that the G_3^{xyy} term could correspond to a *pre-post-post* sequence as well as a *post-pre-post* sequence. Similarly the term F_2^{xy} picks up the changes caused by arrival of a presynaptic spike after postsynaptic spike firing. Several learning rules with all-to-all interaction can be classified in this framework, e.g. [5, 6, 7, 8, 9, 10].

2.2 Our Model

Not all term in the expansion need to be non-zero. In fact, in the results section we will show that a learning rule with $G_3^{xyy}(s, s') \geq 0$ for all $s, s' > 0$ and $F_2^{xy}(s) \leq 0$ for $s > 0$ and all other terms set to zero is sufficient to explain the results from protocols $P1$ and $P2$. Thus, in our learning rule an isolated pair of spikes in configuration *post-before-pre* will lead to depression. An isolated spike pair *pre-before-post*, on the other hand, would not be sufficient to trigger potentiation, whereas a triplet *pre-post-post* or *post-pre-post* will do so (see Fig. 2).

¹For the sake of clarity we have omitted the dependence on $\bar{\rho}_y$.

Figure 2: **A.** Triplet interaction for LTP **B.** Pair interaction for LTD.

To be specific, we consider

$$F_2^{xy}(s) = -A_-(\bar{\rho}_y)e^{-\frac{s}{\tau_-}} \quad \text{and} \quad G_3^{xyy}(s, s') = A_+e^{-\frac{s}{\tau_+}} e^{-\frac{s'}{\tau_y}}. \quad (7)$$

Such an exponential model can be implemented by a mechanistic update involving three variables (the dot denotes a temporal derivative)

$$\begin{aligned} \dot{a} &= -\frac{a}{\tau_+}; & \text{if } t = t_x^k \text{ then } a &\rightarrow a + 1 \\ \dot{b} &= -\frac{b}{\tau_-}; & \text{if } t = t_y^k \text{ then } b &\rightarrow b + 1 \\ \dot{c} &= -\frac{c}{\tau_y}; & \text{if } t = t_y^k \text{ then } c &\rightarrow c + 1 \end{aligned} \quad (8)$$

The weight update is then

$$\dot{w}(t) = -A_-(\bar{\rho}_y)X(t)b(t) + A_+Y(t)a(t)c(t). \quad (9)$$

2.3 Nearest Spike Expansion (truncated model)

Following ideas of [11, 12, 13], the expansion can also be restricted to neighboring spikes only. Let us denote by $f_y(t)$ the firing time of the last postsynaptic spike before time t . Similarly, $f_x(t')$ denotes the timing of the last presynaptic spike preceding t' . With this notation the Volterra expansion of the preceding section can be repeated in a form that only nearest spikes play a role. A classification of the models [11, 12, 13] is hence possible.

We focus immediately on the truncated version of our model

$$\dot{w}(t) = X(t)F_2^{xy}(t - f_y(t), \bar{\rho}_y(t)) + Y(t)G_3^{xyy}(t - f_x(t), t - f_y(t)) \quad (10)$$

The mechanistic model that generates the truncated version of the model is similar to Eq. (8) except that under the appropriate update condition, the variable goes to one, i.e. $a \rightarrow 1$, $b \rightarrow 1$ and $c \rightarrow 1$. The weight update is identical to that of the all-to-all model, Eq. (9).

3 Results

One advantage of our formulation is that we can derive explicit formulas for the total weight changes induced by protocols P1 and P2.

3.1 All-to-all Interaction

If we use protocol P1 with a total of N pre- and postsynaptic spikes at frequency ρ shifted by a time Δt , then the total weight change Δw is for our model with all-to-all interaction

$$\begin{aligned} \Delta w &= A_+ \sum_{k=0}^{N-1} \sum_{k'=1}^{N-1} (N - \max(k, k')) \exp\left(-\frac{k/\rho + \Delta t}{\tau_+}\right) \exp\left(-\frac{k'}{\tau_y \rho}\right) \lambda_k(-\Delta t) \\ &\quad - A_-(\bar{\rho}_y) \sum_{k=0}^{N-1} (N - k) \exp\left(-\frac{k/\rho - \Delta t}{\tau_-}\right) \lambda_k(\Delta t) \end{aligned} \quad (11)$$

where $\lambda_k(\Delta t) = 1 - \delta_{k0}\Theta(\Delta t)$ with Θ the Heaviside step function. The results are plotted in Fig. 3 top-left for $N = 60$ spikes.

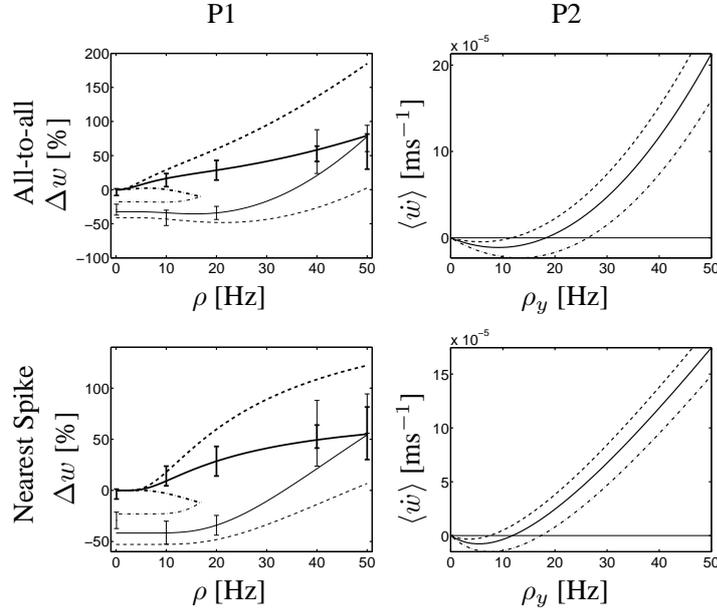


Figure 3: Triplet learning rule. Summary of all results of protocol *P1* (left) and *P2* (right) for an all-to-all (top) and nearest-spike (bottom) interaction scheme. For the left column, the upper thick lines correspond to positive timing ($\Delta t > 0$) while the lower thin lines to negative timing. Dashed line: $\Delta t = \pm 2$ ms, solid line: $\Delta t = \pm 10$ ms and dot-dashed line $\Delta t = \pm 30$ ms. The error bars indicate the experimental data points of Fig. 1A. Right column: dashed-line $\bar{\rho}_y = 8$ Hz, solid line $\bar{\rho}_y = 10$ Hz and dot-dashed line $\bar{\rho}_y = 12$ Hz. Top: $\tau_y = 200$ ms, bottom: $\tau_y = 40$ ms.

The mean firing rate $\bar{\rho}_y$ reflects the firing activity during the recent past (i.e. *before* the start of the experiment) and is assumed as fixed during the experiment. The exact value does not matter. Overall, the frequency dependence of changes Δw is very similar to that observed in experiments. If X and Y are independent Poisson process, the protocol *P2* gives a total weight change that can be calculated using standard arguments [8]

$$\langle \dot{w} \rangle = -A_-(\bar{\rho}_y)\rho_x\rho_y\tau_- + A_+\rho_x\rho_y^2\tau_+\tau_y \quad (12)$$

As before, the mean firing rate $\bar{\rho}_y$ reflects the firing activity during the recent past and is assumed as fixed during the experiment. In order to implement a sliding threshold as in the BCM rule, we take $A_-(\bar{\rho}_y) = \beta_- \bar{\rho}_y^2 / \rho_0^2$ where we set $\rho_0 = 10$ Hz. This yields a frequency dependent threshold $\theta(\bar{\rho}_y) = \beta_- \tau_- \bar{\rho}_y^2 / (A_+ \tau_+ \tau_y \rho_0^2)$. As can be seen in Fig. 3 top-right our model exhibits all essential features of a BCM rule.

3.2 Nearest Spike Interaction

We now apply protocols *P1* and *P2* to our truncated rule, i.e. restricted to the *nearest-spike* interaction; cf. Eq. (10) where the expression of F_2^{xy} and G_3^{xyy} are taken from Eq. (7). The weight change Δw for the protocol *P1* can be calculated explicitly and is plotted in Fig. 3 bottom-left. For protocol *P2* (see Fig. 3 bottom-right) we find

$$\langle \dot{w} \rangle = \rho_x \left(-\frac{A_-(\bar{\rho}_y)\rho_y}{\rho_y + \alpha_-} + \frac{A_+}{\rho_x + \alpha_+} \frac{\rho_y^2}{\rho_y + \alpha_y} \right) \quad (13)$$

where $\alpha_y = \tau_y^{-1}$. If we assume that $\rho_x \ll \alpha_x$, Eq. (13) is a BCM learning rule.

In summary, both versions of our learning rule (all-to-all or nearest-spike) yield a frequency dependence that is consistent with experimental results under protocol P1 and with the BCM rule tested under protocol P2. We note that our learning rule contains only two terms, i.e., a triplet term (1 pre and 2 post) for potentiation and a *post-pre* pair term for depression. The dynamics is formulated using five variables ($a, b, c, \bar{\rho}_y, w$) and five parameters ($\tau_+, \tau_-, \tau_y, A_+, \beta_-$). $\tau_+ = 16.8$ ms and $\tau_- = 33.7$ ms are taken from [14]. A_+ and β_- are chosen such that the weight changes for $\Delta t = \pm 10$ ms and $\rho = 20$ Hz fit the experimental data [11].

4 Discussion - Comparison with Other Rules

While we started out developing a general framework, we focused in the end on a simple model with only five parameters - why, then, this model and not some other combination of terms? To answer this question we apply our approach to a couple of other models, i.e., pair-based models (all-to-all or nearest spike), triplet-based models, and others.

4.1 STDP Models Based on Spike Pairs

Pair-based models with all-to-all interaction [4, 5, 6, 7, 8, 9, 10] yield under Poisson stimulation (protocol P2) a total weight change that is linear in presynaptic and postsynaptic frequencies. Thus, as a function of postsynaptic frequency we always find a straight line with a slope that depends on the integral of the STDP function [5, 7]. Thus pair-based models with all-to-all interaction need to be excluded in view of BCM features of plasticity [25, 24].

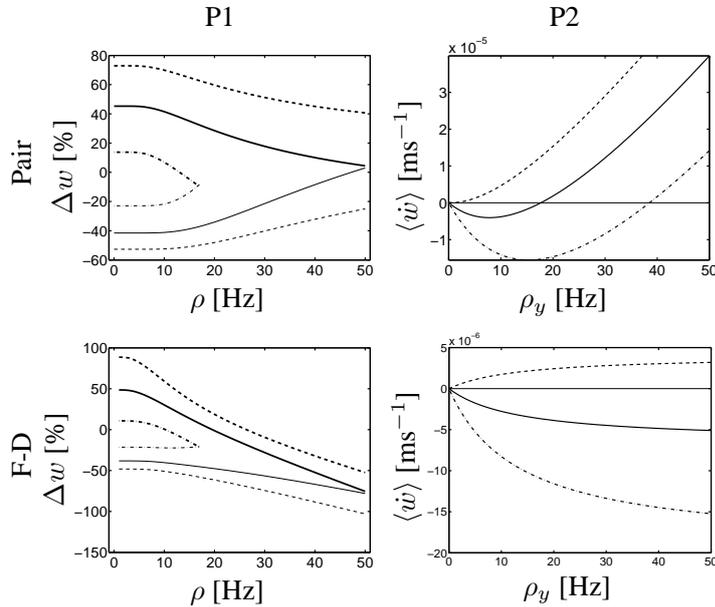


Figure 4: Pair learning rule in a nearest spike interaction scheme (top) and Froemke-Dan rule (bottom). For the left column, the higher thick lines correspond to positive timing ($\Delta t > 0$) while the lower thin lines to negative timing. Dashed line: $\Delta t = \pm 2$ ms, solid line: $\Delta t = \pm 10$ ms and dot-dashed line $\Delta t = \pm 30$ ms. Right column: dashed-line $\bar{\rho}_y = 8$ Hz, solid line $\bar{\rho}_y = 10$ Hz and dot-dashed line $\bar{\rho}_y = 12$ Hz. The parameters of the F-D model are taken from [21]. The dependence upon $\bar{\rho}_y$ has been added to the original F-D rule ($A_- \rightarrow \beta_- \bar{\rho}_y^2 / \rho_0^2$).

A pair-based model with nearest-spike interaction, however, can give a non-linear dependence upon the postsynaptic frequency under protocol P2 with fixed threshold between

depression and potentiation [12]. We can go beyond the results of [12] by adding a suitable dependence of the parameter A_- upon $\bar{\rho}_y$ which yields a sliding threshold; cf. Fig. 4 top right.

But even a pair rule restricted to nearest-spike interaction is unable to account for the results of protocol P1. An important feature of the experimental results with protocol P1 is that potentiation only occurs above a minimal firing frequency of the postsynaptic neuron (cf. Fig. 1A) whereas pair-based rules *always* exhibit potentiation with pre-before-post timing even in the limit of low frequencies; cf. Fig. 4 top left. The intuitive reason is that at low frequency the total weight change is proportional to the number of *pre-post* pairings and this argument can be directly transformed into a mathematical proof (details omitted). Thus, pair-based rules of potentiation (all-to-all or nearest spike) cannot account for results of protocol P1 and must be excluded.

4.2 Comparison with Triplet-Based Learning Rules

The model of Senn et al. [16] can well account of the results under protocol P1. A classification of this rule within our framework reveals that the update algorithm generates pair terms of the form *pre-post* and *post-pre*, as well as triplet terms of the form *pre-post-post* and *post-pre-pre*. As explained in the previous paragraph, a pair term *pre-post* generated potentiation even at very low frequencies which is not realistic. In order to avoid this effect in their model, Senn et al. included additional threshold values which increased the number of parameters in their model to 9 [16] while the number of variables is 5 as in our model. Moreover, the mapping of the model of Senn et al. to the BCM rule is not ideal, since the sliding threshold is different for each individual synapse [16].

An explicit triplet rule has been proposed by Froemke and Dan [21]. In our framework, the rule can be classified as a combination of triplet terms for potentiation and depression. Following the same line or argument as in the preceding sections we can calculate the total weight change for protocols P1 and P2. The result is shown in Fig. 4 bottom. We can clearly see that the pairing experiment P1 yields a behavior opposite to the one found experimentally and the BCM behavior is not at all reproduced in protocol P2.

4.3 Summary

We consider our model as a minimal model to account for results of protocol P1 and P2, but, of course, several factors are not captured by the model. First, our model has no dependence upon the current weight value, but, in principle, this could be included along the lines of [10]. Second, the model has no explicit dependence upon the membrane potential or calcium concentration, but the postsynaptic neuron enters only via its firing activity. Third, and most importantly, there are other experimental paradigms that have to be taken care of.

In a recent series of experiments Bi and colleagues [22] have systematically studied the effect of symmetric spike triplets (*pre-post-pre* or *post-pre-post*) and spike quadruplets (e.g., *pre-post-post-pre*) in hippocampal cultures. While the model presented in this paper is intended to model the synaptic dynamic for L5 pyramidal neurons in the visual cortex [11], it is possible to consider a similar model for the hippocampus containing two extra terms (a pair term for potentiation and and triplet term for depression).

References

- [1] Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. *Science* **275**, 213–215 (1997).
- [2] Zhang, L., Tao, H., Holt, C., W.A.Harris, and Poo, M.-M. *Nature* **395**, 37–44 (1998).

- [3] Bi, G. and Poo, M. *Ann. Rev. Neurosci.* **24**, 139–166 (2001).
- [4] Gerstner, W., Kempter, R., van Hemmen, J. L., and Wagner, H. *Nature* **383**, 76–78 (1996).
- [5] Kempter, R., Gerstner, W., and van Hemmen, J. L. *Phys. Rev. E* **59**, 4498–4514 (1999).
- [6] Roberts, P. *J. Computational Neuroscience* **7**, 235–246 (1999).
- [7] Song, S., Miller, K., and Abbott, L. *Nature Neuroscience* **3**, 919–926 (2000).
- [8] Kistler, W. M. and van Hemmen, J. L. *Neural Comput.* **12**, 385–405 (2000).
- [9] Rubin, J., Lee, D. D., and Sompolinsky, H. *Physical Review Letters* **86**, 364–367 (2001).
- [10] Gütig, R., Aharonov, R., Rotter, S., and Sompolinsky, H. *J. Neuroscience* **23**, 3697–3714 (2003).
- [11] Sjöström, P., Turrigiano, G., and Nelson, S. *Neuron* **32**, 1149–1164 (2001).
- [12] Izhikevich, E. and Desai, N. *Neural Computation* **15**, 1511–1523 (2003).
- [13] Burkitt, A. N., Meffin, M. H., and Grayden, D. *Neural Computation* **16**, 885–940 (2004).
- [14] Bi, G.-Q. *Biological Cybernetics* **319-332** (2002).
- [15] van Rossum, M. C. W., Bi, G. Q., and Turrigiano, G. G. *J. Neuroscience* **20**, 8812–8821 (2000).
- [16] Senn, W., Tsodyks, M., and Markram, H. *Neural Computation* **13**, 35–67 (2001).
- [17] Shouval, H. Z., Bear, M. F., and Cooper, L. N. *Proc. Natl. Acad. Sci. USA* **99**, 10831–10836 (2002).
- [18] Abarbanel, H., Huerta, R., and Rabinovich, M. *Proc. Natl. Academy of Sci. USA* **59**, 10137–10143 (2002).
- [19] Karmarkar, U., Najarian, M., and Buonomano, D. *Biol. Cybernetics* **87**, 373–382 (2002).
- [20] Sjöström, P., Turrigiano, G., and Nelson, S. *Neuron* **39**, 641–654 (2003).
- [21] Froemke, R. and Dan, Y. *Nature* **416**, 433–438 (2002).
- [22] Wang, H. X., Gerkin, R. C., Nauen, D. W., and Bi, G. Q. *Nature Neuroscience* **8**, 187–193 (2005).
- [23] Bienenstock, E., Cooper, L., and Munro, P. *Journal of Neuroscience* **2**, 32–48 (1982). reprinted in Anderson and Rosenfeld, 1990.
- [24] Kirkwood, A., Rioult, M. G., and Bear, M. F. *Nature* **381**, 526–528 (1996).
- [25] Artola, A. and Singer, W. *Trends Neurosci.* **16**(11), 480–487 (1993).
- [26] Toyozumi, T., Pfister, J.-P., Aihara, K., and Gerstner, W. In *Advances in Neural Information Processing Systems 17*, Saul, L. K., Weiss, Y., and Bottou, L., editors, 1409–1416. MIT Press, Cambridge, MA (2005).
- [27] Fusi, S., Annunziato, M., Badoni, D., Salamon, A., and D.J.Amit. *Neural Computation* **12**, 2227–2258 (2000).
- [28] Toyozumi, T., Pfister, J.-P., Aihara, K., and Gerstner, W. *Proc. National Academy Sciences (USA)* **102**, 5239–5244 (2005).
- [29] Volterra, V. *Theory of Functionals and of Integral and Integro-Differential Equations*. Dover, New York, (1930).

2.7 Paper VII

Summary

This last paper can be seen as the extension of paper VI. Because our triplet rule fitted surprisingly well the frequency data of Sjöström et al. (2001), we decided to further test the idea of triplet learning rule by applying it to different preparations and different brain regions.

We therefore considered the triplet and quadruplet experiments of Wang et al. (2005) and fitted our triplet models. We systematically studied the pair and triplet rules for both an All-to-All and a Nearest-Past-Spike interaction schemes.

The main message of this paper is that triplet learning rules succeed (and pair-based models fail) to reproduce frequency effects in pairing protocols as well as triplet and quadruplet experiments. Moreover, it is possible to reduce the triplet model into a *minimal* triplet model (5 parameters) which contains only one extra parameter, when compared to pair-based models (4 parameters), and fits experimental data almost as well as the *full* triplet model (8 parameters).

Similarly to paper VI, we show the mapping to the BCM learning rule. Moreover, if Poisson statistics is assumed for the pre- and postsynaptic neuron and under some conditions, we show that the averaged triplet learning follows the gradient of an objective function L . In this way, the phenomenological triplet learning rule can be seen as well as an optimal learning rule.

Reference

This paper has been presented at the *Cosyne conference* (March 5 - 11, 2006, Salt-Lake City, USA). This paper is submitted under the following reference:

Pfister JP and Gerstner W. Why Triplets of Spike are Necessary in Models of Spike-Timing-Dependent Plasticity. *Submitted to The Journal of Neuroscience.*

Why Triplets of Spike are Necessary in Models of Spike-Timing-Dependent Plasticity

Jean-Pascal Pfister, Wulfram Gerstner
Laboratory of Computational Neuroscience,
School of Computer and Communication Sciences
and Brain-Mind Institute,
Ecole Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne
{jean-pascal.pfister, wulfram.gerstner}@epfl.ch

Submitted to *The Journal of Neuroscience*

Abstract

Classical experiments on spike-timing dependent plasticity (STDP) use a protocol based on *pairs* of pre- and postsynaptic spikes repeated at a given frequency in order to induce synaptic potentiation or depression. Therefore standard STDP models have expressed the weight change as a function of pairs of pre- and postsynaptic spike. Unfortunately, those paired-based STDP models cannot account for the dependence upon the repetition frequency of the pairs of spike. Moreover, those STDP models cannot reproduce recent triplet and quadruplet experiments. Here we examine a *triplet* rule, i.e. a rule which considers sets of three spikes (2 pre and 1 post or 1 pre and 2 post) and compare it to classical pair-based STDP learning rules. With such a triplet rule, it is possible to fit experimental data from visual cortical slices as well as from hippocampal cultures. Moreover, when assuming stochastic spike trains, the triplet learning rule can be mapped to a Bienenstock-Cooper-Munro learning rule.

1 Introduction

During the last decade an increasing number of experiments have shown that synaptic strength changes as a function of the precise timing of the pre- and postsynaptic neurons. In the early experiments (Markram, Lübke, Frotscher, and Sakmann 1997; Zhang, Tao, Holt, Harris, and Poo 1998; Bi and Poo 1998; Bi and Poo 2001), potentiation has been elicited by a sequence of N pairs of “pre then post” spikes whereas depression occurred when the timing was reversed i.e. when each postsynaptic spike precedes a presynaptic one. At this point, it was natural to characterize synaptic plasticity as a function of the time difference $\Delta t = t^{\text{post}} - t^{\text{pre}}$ within pairs of spikes. However performing experiments with pairs of spike does not mean that pairs of spikes are the elementary building block. There is no *a priori* reason to think that pairs of spikes are more relevant than three spikes (*triplets*), four spikes (*quadruplets*) or even more. It is clear that a lot of other neuronal variables such as calcium concentration (Malenka, Kauer, Zucker, and Nicoll 1988; Lisman 1989; Lisman and Zhabotinsky 2001; Shouval, Bear, and Cooper 2002) or postsynaptic membrane potential (Sjöström, Turrigiano, and Nelson 2001; Rao and Sejnowski 2001; Lisman and Spruston 2005) play an important role in triggering potentiation or depression. The point of this study is to see how far we can explain experiments that only use spike timing as a parameter with models that only use spike timing.

Recent experiments (Bi and Wang 2002; Froemke and Dan 2002; Wang, Gerkin, Nauen, and Bi 2005; Froemke, Tsay, Raad, Long, and Dan 2006) have studied the detailed role of spike timing by triggering synaptic plasticity with spike triplets (1 presynaptic spike combined with 2 postsynaptic ones or 1 postsynaptic spike with 2 presynaptic spikes). The results of those experiments indicate that a pair rule is not sufficient to explain synaptic changes triggered by triplets or quadruplets of spikes.

In the first part of this study, we will review some experimental protocols performed in visual cortex and hippocampal culture and show why the classical pair-based STDP models fail to reproduce those experimental data. The visual cortex data set (Sjöström, Turrigiano, and Nelson 2001) used in this paper consists of a standard pairing protocol where the frequency of the pairing has been changed. We also considered a hippocampal culture data set (Wang, Gerkin, Nauen, and Bi 2005) which consists of pair, triplet and quadruplet protocols. Since both data sets disagree on some specific protocols (at low frequency of the pairing protocol, no potentiation is elicited in Sjöström’s data while a large amount of potentiation is present in Wang’s data) and since the preparations are different, we fitted our models with different parameters for each data set.

In the second part of this study, we will show that if we assume that synaptic plasticity is governed by a suitable combination of pairs and triplets of spikes, the results from the above mentioned protocols can be surprisingly well reproduced. Moreover, we show that our triplet learning rule elicits input selectivity analogous to that of the Bienenstock-Cooper-Monro (BCM) theory (Bienenstock, Cooper, and Munro 1982).

Claiming that triplet of spikes are more relevant than pairs of spike is not enough

to construct a model of synaptic plasticity. It is also necessary to determine how those pairs or triplets of spikes integrate. For both the pair-based models and the triplet-based models, we will consider the case where a presynaptic spike interacts with all previous postsynaptic ones or vice versa (we will call this, the *All-to-All* interaction) (Gerstner, Kempter, van Hemmen, and Wagner 1996; Kempter, Gerstner, and van Hemmen 1999; Song, Miller, and Abbott 2000; Kistler and van Hemmen 2000) and the case where only neighboring spikes are taken into account (*Nearest-Spike* interaction) (van Rossum, Bi, and Turrigiano 2000; Bi 2002; Izhikevich and Desai 2003; Burkitt, Meffin, and Grayden 2004; Pfister and Gerstner 2006).

In this study, we will show that a *Nearest-Spike* interaction scheme does not rescue classical STDP pair-based models. However, if we go from pair-based STDP rules to triplet models, both Nearest-Spike and the All-to-All interactions are possible. We will have a slight preference for the All-to-All interaction because it allows a better fit to quadruplets experiments.

2 Material and Methods

2.1 Synaptic Learning Rule

In this paper, we study a new triplet-based model of STDP and compare it to classical pair-based STDP models. Traditional *mechanistic* models of STDP involve a small number of variables that are updated by pre- and postsynaptic firing events (Kistler and van Hemmen 2000; Karmarkar and Buonomano 2002; Abarbanel, Huerta, and Rabinovich 2002; Gerstner and Kistler 2002). The new triplet rule is formulated in such a framework.

In order to introduce the variables used in our model, let us consider the process of synaptic transmission. Whenever a presynaptic spike arrives at an excitatory synapse, glutamate is released into the synaptic cleft and binds to glutamate receptors. Let r_1 denote the amount of glutamate bound to a postsynaptic receptor. The variable r_1 increases whenever there is a presynaptic spike and decreases back to zero otherwise with a time constant of τ_+ . This can be written as

$$\frac{dr_1(t)}{dt} = -\frac{r_1(t)}{\tau_+} \quad \text{if } t = t^{\text{pre}} \text{ then } r_1 \rightarrow r_1 + 1 \quad (1)$$

Here, t^{pre} denotes the moment of spike arrival at the presynaptic terminal. The units of r_1 are chosen such that glutamate binding increases by one unit upon spike arrival. We emphasize that r_1 is an abstract variable. Instead of glutamate binding, it could describe equally well some other quantity that increases upon presynaptic spike arrival. We call r_1 a detector of presynaptic events.

Instead of having only one process triggered by a presynaptic spike, it is possible to consider several different quantities which increase in the presence of a presynaptic spike. In our model, we will consider two different detectors of presynaptic events, namely r_1 and r_2 . The dynamics of r_2 is analogous to that of r_1 except that its time

constant τ_x is larger than τ_+ . Similarly, we assume that each postsynaptic spike t^{post} induces an increase of two different quantities that we will denote o_1 and o_2 . Potential interpretations of o_1 and o_2 are given below. In the absence of postsynaptic spike, these postsynaptic detectors decrease their value with a time constant τ_- and τ_y , respectively. Formally, this gives

$$\begin{aligned} \frac{dr_2(t)}{dt} &= -\frac{r_2(t)}{\tau_x} && \text{if } t = t^{\text{pre}} \text{ then } r_2 \rightarrow r_2 + 1 \\ \frac{do_1(t)}{dt} &= -\frac{o_1(t)}{\tau_-} && \text{if } t = t^{\text{post}} \text{ then } o_1 \rightarrow o_1 + 1 \\ \frac{do_2(t)}{dt} &= -\frac{o_2(t)}{\tau_y} && \text{if } t = t^{\text{post}} \text{ then } o_2 \rightarrow o_2 + 1 \end{aligned} \quad (2)$$

We do not want to identify the variables r_1 , r_2 , o_1 and o_2 with specific biophysical quantities. Candidates of detectors of presynaptic events are, e.g., the amount of glutamate bound (Karmarkar and Buonomano 2002) or the number of NMDA receptors in an activated state (Senn, Markram, and Tsodyks 2001). Postsynaptic detectors o_1 and o_2 could represent the influx of calcium concentration through voltage-gated Ca^{2+} channels and NMDA channels (Karmarkar and Buonomano 2002) or the number of secondary messengers in a deactivated state of the NMDA receptor (Senn et al. 2001), or the voltage trace of a back-propagating action potential (Shouval, Bear, and Cooper 2002).

Since our present model is formulated as a mechanistic model, it is possible to define changes of synaptic efficacies for our triplet learning rule with All-to-All interactions as a function of those four detectors without making any assumption on the biophysical quantities they represent. We assume that the weight decreases upon presynaptic spike arrival by an amount that is proportional to the value of the postsynaptic variable o_1 , but depends also on the value of the second presynaptic detector r_2 . Hence presynaptic spike arrival at time t^{pre} triggers a change given by

$$w(t) \rightarrow w(t) - o_1(t) [A_2^- + A_3^- r_2(t - \epsilon)] \quad \text{if } t = t^{\text{pre}} \quad (3)$$

Similarly, a postsynaptic spike at time t^{post} triggers a change that depends on the presynaptic variable r_1 and the second postsynaptic variable o_2 :

$$w(t) \rightarrow w(t) + r_1(t) [A_2^+ + A_3^+ o_2(t - \epsilon)] \quad \text{if } t = t^{\text{post}} \quad (4)$$

Here A_2^+ and A_2^- denote the amplitude of the weight change whenever there is a pre-post pair or a post-pre pair. Similarly A_3^+ and A_3^- denote respectively the amplitude of the *triplet* term for potentiation and depression (see Fig. 1A). All those 4 amplitude parameters are assumed to be greater or equal to zero. ϵ is a small positive constant in order to make sure that the weight is updated *before* the

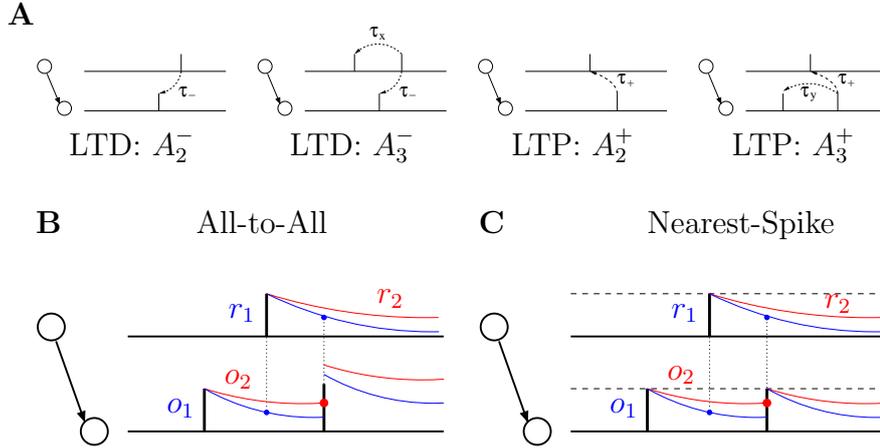


Figure 1: Schematic description of the triplet learning rules. **A**. Schematic description of the terms contributing to long-term depression (LTD) controlled by A_2^- and A_3^- and the two long-term potentiation (LTP) terms controlled by A_2^+ and A_3^+ . A presynaptic spike after a postsynaptic one (post \rightarrow pre) induces LTD if the temporal difference is not much larger than τ_- (pair term, A_2^-). The presence of a previous presynaptic spike gives a further contribution (2-pre-1-post triplet term, A_3^-) if the interval between the two presynaptic spikes is not much larger than τ_x . Similarly, the triplet term for LTP depends on 1 presynaptic spike, but 2 postsynaptic spikes. The presynaptic spike must occur *before* the second postsynaptic one with a temporal difference not much larger τ_+ . **B**. Time course of detectors of pre- and postsynaptic events r_1 , r_2 , o_1 and o_2 . The presynaptic variables r_1 and r_2 are increased by a fixed amount upon arrival of a presynaptic spike. Analogously, postsynaptic variables are updated upon presynaptic firing. With All-to-All interactions, each postsynaptic spike interacts with all previous postsynaptic spikes and vice versa, i.e. the internal variables r_1 , r_2 , o_1 and o_2 accumulate over several postsynaptic spike timings. The red and blue dots denote the values of those internal variables “read” by the triplet model whenever a spike occurs; e.g. the value of the postsynaptic variable o_1 is “read out” at the moment of presynaptic spike arrival leading to synaptic depression proportional to the momentary value of o_1 (blue dot). Similarly, the value of the presynaptic variable r_1 and the postsynaptic variable o_2 are “read out” at the moment of the second postsynaptic spike and determine the amplitude of synaptic potentiation. **C**. Same as in **B**, but with Nearest-Spike interactions: the extension of the spike interaction is restricted to the last spike; no accumulation occurs.

detectors r_2 or o_2 . In other words, r_2 is zero unless a *previous* presynaptic spike has led to an increase of r_2 . This ensures the detection of spike triplets.

Fig. 1B illustrates how 1-pre-2-post triplets are detected by the learning rule. At the time of a postsynaptic spike, the learning rule “reads” the value of the second postsynaptic variable o_2 just before the spike (see red dot at time $t^{\text{post}} - \epsilon$ in Fig. 1B) as well as the value of the presynaptic detector r_1 (see blue dot at time t^{post} in Fig. 1B) and increases the weight by an amount $A_2^+ r_1(t^{\text{post}}) + A_3^+ r_1(t^{\text{post}}) o_2(t^{\text{post}} - \epsilon)$.

Note that if we set $A_3^+ = 0$ and $A_3^- = 0$, the model becomes a classical pair-based STDP model (Gerstner, Kempter, van Hemmen, and Wagner 1996; Kempter, Gerstner, and van Hemmen 1999; Song, Miller, and Abbott 2000; Kistler and van Hemmen 2000). This pair-based STDP model was used for the results of Fig. 2. It should be further noted that the two extra triplet terms vanish if a single spike pair is presented or if spike pairings are repeated at low frequency. This means that in the limit of low frequency, the classical pair-based learning rule is identical to this triplet learning rule.

The triplet learning rule of Eqs. (3) and (4) can also describe a Nearest-Spike interaction scheme if we redefine the update rule of pre- and postsynaptic detectors. Instead of simply low-pass filtering the spike trains, i.e. adding the effects of all spikes, the detector variables saturate at 1, i.e. $0 \leq r_1, r_2, o_1, o_2 \leq 1$. This is achieved by updating the variables *to* the value of 1 instead of updating *by* at step of 1. In this way, the synapse *forgets* all other previous spikes and keeps only the memory of the last one (see Fig. 1C).

In this paper we consider first a *full triplet model* which takes into account all the four terms of Eqs. (3) and (4). Then, we will see that only some of the terms are really necessary. This is why we define two different minimal models. The first one is intended to fit the visual cortex data and disregards two terms i.e. $A_2^+ = 0$ and $A_3^- = 0$. For the hippocampal culture data set we consider a slightly different minimal model which disregards only one term, i.e. $A_3^- = 0$.

In principle, the amplitude parameters A_2^+ , A_2^- , A_3^+ and A_3^- could change on a slow time scale. For example, similar to the threshold in the BCM rule, those parameters could, because of homeostatic processes, depend on the mean postsynaptic firing rate $\bar{\rho}_y$ averaged over a time scale of 10 minutes or more.

2.2 Protocols

In order to compare our model to experimental data, we followed three different experimental protocols (see Fig. 2) in which the synaptic weight changes as a function of the pre- and postsynaptic spike statistics. The fourth protocol is of a more theoretical value in the sense that it can be compared to the BCM learning rule which has interesting computational properties.

Pairing Protocol. (see Fig. 2A). This is the classical STDP protocol (Markram, Lübke, Frotscher, and Sakmann 1997; Zhang, Tao, Holt, Harris, and Poo 1998; Bi and Poo 1998; Bi and Poo 2001; Sjöström, Turrigiano, and Nelson 2001; Froemke and Dan 2002). $N = 60$ pairs of pre- and postsynaptic spikes shifted by Δt are

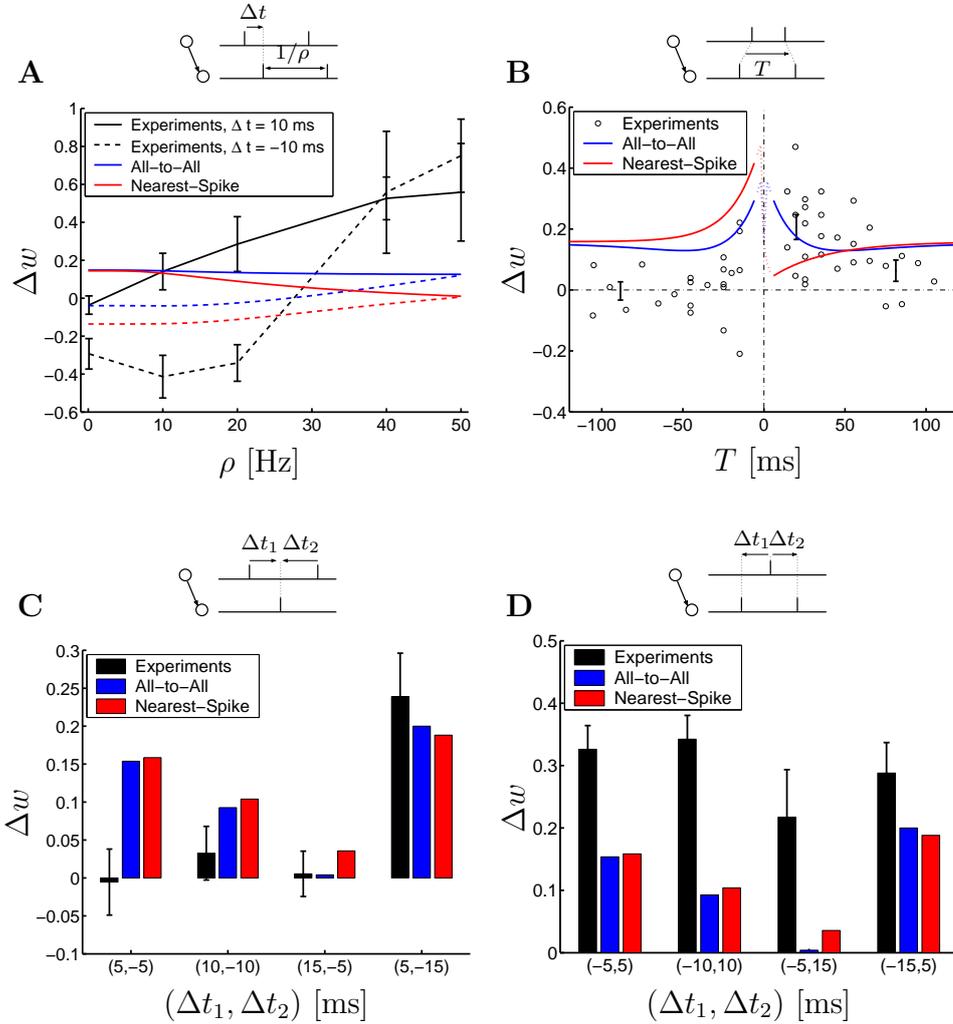


Figure 2: Failure of pair-based STDP learning rules. In all four subgraphs, black lines or symbols denote experimental data, blue lines correspond to the All-to-All pair model and the red lines to the Nearest-Spike pair model (see text for details). **A**. Weight change in a pairing protocol as a function of the frequency ρ (*solid lines*: $\Delta t = +10$ ms, *dashed lines*: $\Delta t = -10$ ms). Black lines and data points (with errors) redrawn from Sjöström (2001). The experimental data are neither reproduced at high nor at low values of the repetition frequency ρ . **B**. Quadruplet protocol. Black circles are redrawn from Wang et al. (2005). **C**. Triplet protocol for the *pre-post-pre* case and **D** for the *post-pre-post* case. Black dots in **B** and black bars (and standard errors) in **C** and **D** from Wang et al. (2005). The asymmetry of the experimental results (no potentiation for $(\Delta t_1, \Delta t_2) = (5 \text{ ms}, -5 \text{ ms})$ in **C** but strong potentiation for $(-5 \text{ ms}, 5 \text{ ms})$ in **D**) is not captured by the pair-based models.

elicited at regular intervals of $1/\rho$. The interest of the study of Sjöström, Turrigiano, and Nelson (2001) is that the authors analyzed, in this pairing protocol, the weight change as a function of the frequency ρ for a fixed Δt . Changing the frequency ρ is a good way to check the validity of a model, especially at high frequency where many spikes are potentially in the temporal range of interaction.

It should be noted that the amount of potentiation for a pre-post ($\Delta t = 10$ ms) pair reported in (Wang, Gerkin, Nauen, and Bi 2005) is significantly lower than the one originally measured (Bi and Poo 1998) under the same conditions. As mentioned in (Wang et al. 2005), this can be accounted for by the difference in initial synaptic strength which was higher in (Bi and Poo 1998). In order to test our model on a consistent set of data, we took the measurements of Wang et al. (c.f. sup. Fig. 1 of Wang et al. 2005), i.e. $\Delta w \simeq 0.25 \pm 0.05$ for $\Delta t = +10$ ms and $\Delta w \simeq -0.17 \pm 0.05$ for $\Delta t = -10$ ms. Data from Wang et al. (2005) including error bars are redrawn in Fig. 3. Since the potentiation and depression time constant are not present in (Wang et al. 2005), we took $\tau_+ = 16.8$ ms and $\tau_- = 33.7$ from (Bi and Poo 2001).

Triplet Protocol. The first triplet protocol (see Fig. 2C) consists of $N = 60$ sets of three spikes repeated at a given frequency $\rho = 1$ Hz. Each triplet consists of two presynaptic spikes and one postsynaptic spike characterized by $\Delta t_1 = t_1^{\text{post}} - t_1^{\text{pre}}$ and $\Delta t_2 = t_2^{\text{post}} - t_2^{\text{pre}}$ where t_1^{pre} and t_2^{pre} are the first and second presynaptic spike of the triplet.

The second triplet protocol (see Fig. 2D) also consists of $N = 60$ triplets. The only difference is that each triplet consists of 1 pre and 2 postsynaptic spikes. In this case $\Delta t_1 = t_1^{\text{post}} - t^{\text{pre}}$ and $\Delta t_2 = t_2^{\text{post}} - t^{\text{pre}}$ where t_1^{post} and t_2^{post} are respectively the first and second postsynaptic spike of the triplet.

Experiments with such a triplet protocol have been performed by Froemke and Dan (2002) in L2/3 pyramidal neurons of the rat visual cortex and by Wang et al. (2005) in hippocampal cultures. In order to have a consistent and broad data set, i.e. pair, triplet and quadruplet experiments, we decided, in the present study, to focus only on the data of Wang et al. (2005) (since we did not find enough quantitative information about quadruplets in (Froemke and Dan 2002)).

Quadruplet Protocol. (see Fig. 2B) This protocol consists of $N = 60$ quadruplets at frequency $\rho = 1$ Hz. It was used by Wang et al. (2005) and is characterized as follows: a post-pre pair with a delay of $\Delta t_1 = t_1^{\text{post}} - t_1^{\text{pre}} < 0$ is followed after a time T by a pre-post pair with a delay of $\Delta t_2 = t_2^{\text{post}} - t_2^{\text{pre}} > 0$. When T is negative, the opposite happens. A pre-post pair ($\Delta t_2 = t_2^{\text{post}} - t_2^{\text{pre}} > 0$) is followed by a post-pre pair ($\Delta t_1 = t_1^{\text{post}} - t_1^{\text{pre}} < 0$). Formally, T is defined by $T = (t_2^{\text{pre}} + t_2^{\text{post}})/2 - (t_1^{\text{pre}} + t_1^{\text{post}})/2$. Throughout this paper, we took $\Delta t = -\Delta t_1 = \Delta t_2 = 5$ ms.

Poisson Protocol. The pre- and postsynaptic spike trains are Poisson spike trains with firing rate ρ_x and ρ_y respectively. The interest of such a protocol is that it is possible to establish a link with the BCM learning rule (Bienenstock, Cooper, and Munro 1982) which has attractive theoretical properties. Indeed, this learning rule was originally used to explain the emergence of orientation selectivity in the visual cortex. Even if this protocol has less experimental support than the other protocols,

some aspects of it have been indirectly measured in the visual cortex (Kirkwood, Rioult, and Bear 1996) and in hippocampal slice (Dudek and Bear 1992; Artola, Bröcher, and Singer 1990).

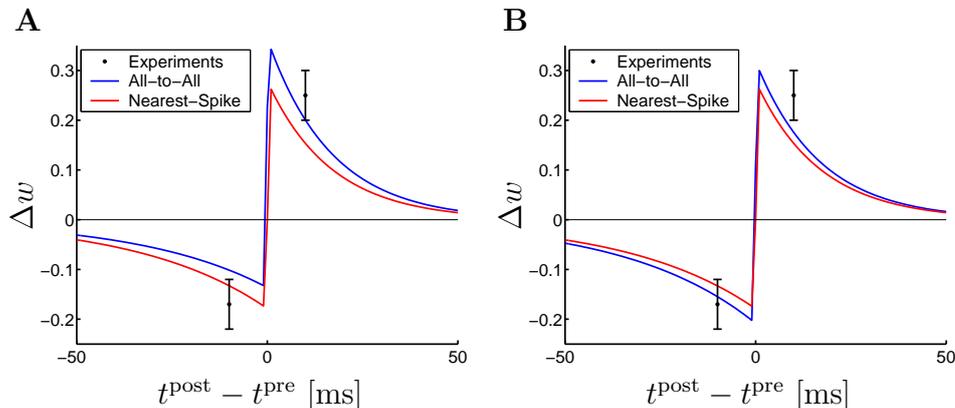


Figure 3: The triplet learning rule can reproduce the STDP learning window. Weight change induced by a repetition of 60 pairs of pre- and postsynaptic spike with a delay of Δt at a repetition frequency of 1 Hz. **A**. Weight change as a function of the time difference between post- and presynaptic spike timing for the full triplet model and **B** for the minimal triplet model. The parameters taken for the triplet models are those which correspond to the hippocampal culture data. See table 3. Experimental data points and standard errors redrawn from Wang et al. (2005).

2.3 Data Fitting

In order to fit the amplitude parameters A_2^+ , A_2^- , A_3^+ , A_3^- and the time constants τ_x and τ_y ($\tau_+ = 16.8$ ms and $\tau_- = 33.7$ ms are kept fixed), we calculated the total weight change Δw_i^{mod} for a given pairing or triplet protocol and compared it to the experimental value Δw_i^{exp} . For the optimization of the parameters, we performed a minimization of the normalized mean-square error E defined by

$$E = \frac{1}{P} \sum_{i=1}^P \left(\frac{\Delta w_i^{\text{exp}} - \Delta w_i^{\text{mod}}}{\sigma_i} \right)^2 \quad (5)$$

where Δw_i^{exp} and σ_i are the experimental mean weight change and standard error of the mean weight change for a given data point i . P is the number of data points within a data set; $P = 10$ for the visual cortex data set (see table 1) and $P = 13$ for the hippocampal culture data set (see table 2). Δw_i^{mod} is the weight change for a given model (pair or triplet model).

	$\Delta t = 10$ ms	$\Delta t = -10$ ms
$\rho = 0.1$ Hz	-0.04 ± 0.05	-0.29 ± 0.08
$\rho = 10$ Hz	0.14 ± 0.1	-0.41 ± 0.11
$\rho = 20$ Hz	0.29 ± 0.14	-0.34 ± 0.1
$\rho = 40$ Hz	0.53 ± 0.11	0.56 ± 0.32
$\rho = 50$ Hz	0.56 ± 0.26	0.75 ± 0.19

Table 1: Experimental weight change Δw as a function of the delay $\Delta t = t^{\text{post}} - t^{\text{pre}}$ induced by a pairing protocol in the visual cortex. Those values are used for the fitting of the pair-based and triplet-based models of visual cortical neurons. Data from Sjöström (2001).

pairing			quadruplet		
Δw	Δt [ms]		Δw	T [ms]	Δt [ms]
0.25 ± 0.05	10		-0.003 ± 0.03	-88.5	5
-0.17 ± 0.05	-10		0.06 ± 0.04	83.7	5
			0.21 ± 0.04	20	5

triplet (2-pre-1-post)			triplet (1-pre-2-post)		
Δw	Δt_1 [ms]	Δt_2 [ms]	Δw	Δt_1 [ms]	Δt_2 [ms]
-0.01 ± 0.04	5	-5	0.33 ± 0.04	-5	5
0.03 ± 0.04	10	-10	0.34 ± 0.04	-10	10
0.01 ± 0.03	15	-5	0.22 ± 0.08	-5	15
0.24 ± 0.06	5	-15	0.29 ± 0.05	-15	5

Table 2: Experimental weight change Δw as a function of the relative spike timing Δt , Δt_1 , Δt_2 and T induced by pairing, triplet and quadruplet protocols in hippocampal cultures. Those data are used for the fitting of the pair-based and triplet-based models of hippocampal culture neurons. Data from Wang et al. (2005).

2.4 Numerical Procedures

The weight change Δw^{mod} for a given model and a given protocol can be either simulated numerically with Eqs. (1) - (4) or calculated analytically. See appendix A for an example of analytical calculation of the weight change of the triplet model applied to the pairing protocol with Nearest-Spike interactions.

In the present paper, the weight changes predicted by all different models (pair-based models, minimal and full triplet-based models with both Nearest-Spike and All-to-All interactions) have been calculated analytically and then evaluated numerically with `Matlab` on a Sun machine. The normalized mean-square error E of Eq. (5) has been minimized with the `Matlab` built-in function `lsqnonlin` which uses a reflective Newton method.

3 Results

3.1 Standard Pair-Based STDP Models Fail to Reproduce Frequency Effects

In a first series of experiments, we applied a classical pair-based STDP learning rule (c.f. Eqs. (3) and (4) with $A_3^- = 0$ and $A_3^+ = 0$) to the pairing protocol with 60 pairs of pre and postsynaptic spikes (see Material and Methods section). Obviously, the weight change predicted by the model depends on the precise choice of the parameters A_2^+ , A_2^- , τ_+ and τ_- . We therefore set those parameters in such a way that the normalized mean square error E across all experimental protocols is minimal (see Eq. (5)). We find that even with the best set of parameters, the classical STDP model fails, for both the All-to-All interaction and the Nearest-Spike interaction, to reproduce the experimental data (see Figure 2A). This is due to the following reasons.

First, as pointed out by Sjöström et al. (2001), a surprising aspect of their finding is that at low repetition frequency ρ there is no potentiation. This can not be captured by standard pair-based STDP models, since for any choice of the parameter $A_2^+ > 0$, the pair-based model induces LTP if a presynaptic spike precedes a postsynaptic one by a few milliseconds.

Secondly, as we can see in Fig. 2, for $\Delta t > 0$ potentiation increases when frequency increases. This behavior can also not be reproduced by classical STDP models. Indeed, in pair-based STDP models, as soon as the frequency increases, the pre-post pairs approach each other and generate an interaction between the postsynaptic spike of one pair and the presynaptic spike of the next pair. The effect of these post-pre pairs should increase with frequency and therefore depress the synapse, which is what is seen in experiments. Therefore classical pair-based model fail to reproduce the pairing experiment of Sjöström et al. (2001).

It should be noted that the absence of potentiation at low frequency is in direct conflict with the results of (Bi and Poo 1998; Zhang, Tao, Holt, Harris, and Poo 1998; Froemke and Dan 2002) where there is a reasonable potentiation at low frequency.

Since the preparation of Sjöström et al. (2001) is different from the one of Bi and Poo (1998) and Wang et al. (2005) and the results in conflict, it seems natural to use different parameters in our model for each data set.

3.2 Standard Pair-Based STDP Models Fail to Reproduce Triplet and Quadruplet Experiments

A second evidence of the limits of pair-based STDP learning rules is the following. In triplet experiments (see Fig. 2C,D), there is a clear asymmetry between a pre-post-pre and a post-pre-post experiment. For example, 60 repetitions of a pre-post-pre triplet with relative timing $(\Delta t_1, \Delta t_2) = (5 \text{ ms}, -5 \text{ ms})$ yields no weight change whereas the same number of repetitions of a post-pre-post triplet with $(\Delta t_1, \Delta t_2) = (-5 \text{ ms}, 5 \text{ ms})$ yields a weight change of approximately 30%. However, any pair-based model would predict the same result for pre-post-pre and post-pre-post experiments since the same pairs occur. Therefore triplet results cannot be explained by a sum of a pre-post potentiation term and a post-pre depression term (see Fig. 2C,D).

Finally, the asymmetry present in the quadruplets experiments (see fig. 2B) also causes some problems for pair-based STDP models. A quadruplet consists of a pre-post-post-pre sequence or a post-pre-pre-post sequence and $|T|$ denotes the interval between the first and last pair of spikes within the quadruplet (see Material and Methods for more details). In a pair-based model with All-to-All interactions and for a given interval $|T|$ between the pairs, the weight changes for post-pre-pre-post and pre-post-post-pre are strictly identical due to the symmetry of the protocol and the symmetry of the All-to-All interaction. The weight change predicted by a pair-model can therefore not explain the asymmetry seen in the data. With Nearest-Spike interactions, the situation gets even worse: pre-post-post-pre quadruplets consist of two pre-post pairs and one post-pre term whereas for the post-pre-pre-post case, the opposite occurs: two post-pre pairs and only one pre-post pair, therefore the Nearest-Spike interaction scheme leads to an asymmetry that is opposite to the one found in experiments (see Fig. 2B).

3.3 Triplet Rule

So far, we have seen that standard pair-based STDP models fail to reproduce frequency effects of the pairing protocol as well as triplet and quadruplet experiments. This is mainly due to the fact that pair-based models are intrinsically symmetric, in the sense that they predict the same weight change for a pre-post pair followed by a post-pre pair with the same delay Δt as for the inverted order, i.e. a post-pre pair followed by a pre-post pair (with the same delay Δt). However, there is no *a priori* reason to think that a pre-post-pre and a post-pre-post triplet should give the same result since they will activate different pre- and postsynaptic pathways.

We therefore included extra terms in the learning rule in order to break the symmetry induced by pair-based models. Specifically, we added a triplet depression term (i.e. a 2-pre-1-post term) as well as a triplet potentiation term (i.e. a 1-pre-

2-post term); see *Material and Methods* for more details. We call this model a *full triplet model* since it includes both pair terms and triplet terms. The full triplet model is described by 8 parameters: 4 amplitude parameters A_2^- , A_2^+ , A_3^- and A_3^+ and 4 time constants: τ_+ , τ_- , τ_x and τ_y . Note that pair-based models are described by 4 parameters (A_2^- , A_2^+ , τ_+ , τ_-).

In analogy to our approach in the previous subsection, we applied our triplet model to the protocols described in *Material and Methods*. More precisely, we calculated analytically for each protocol the weight change predicted by our triplet learning rule (see appendix A for an example of explicit expression of the weight change). As before, we want our triplet learning rule to fit as best as possible the experimental data of Sjöström et al. (2001) or Wang et al. (2005). We therefore minimized the normalized mean square error across all data points of a given data set (see table 1 or 2) by adjusting the 8 parameters mentioned above. The resulting parameters are summarized in table 3.

As a first test for the triplet learning rule, we checked if it can reproduce the biphasic learning window observed by Bi and Poo (1998). Our triplet learning rule succeeds to reproduce the classical STDP learning window (see Fig. 3) since the triplet terms specific to our model play a minor role at a fixed low frequency.

3.4 Triplet Learning Rules can Reproduce Frequency Effects

In this section, we study the pairing protocol used by Sjöström et al. (2001) in visual cortex, i.e., we apply 60 pairs of pre- and postsynaptic spikes at a given frequency ρ . As we can see in Fig. 4A, our full triplet learning rule succeeds to reproduce frequency effects of the pairing protocol. Indeed, the two main problems of the pair-based STDP models explained in section 3.1 for the pairing protocol are solved by the triplet model for the following reasons. First, the absence of potentiation at low frequency is achieved by setting A_2^- to a low value; second, the increase of potentiation with frequency is implemented via the triplet potentiation term controlled by A_3^+ which has a stronger effect than the triplet depression term A_3^- ; see table 3A. Thus our model can explain results at different frequencies without an explicit “potentiation wins” mechanism suggested previously (Sjöström et al. 2001; Wang et al. 2005).

Since some of the optimized parameters of the triplet learning rule have values close to zero, we concluded that the terms controlled by these parameters can be neglected. This allowed us to define a *minimal triplet model* with less parameters. The first parameter we can easily drop is the amplitude A_2^+ of the pair potentiation term since it is extremely small in both the All-to-All and Nearest-Spike interaction scheme (see table 3A). The second parameter we neglect is A_3^- . This is possible for the following reason. In the All-to-All interaction scheme we have $A_3^- \ll A_2^-$, therefore the effect of the triplet depression term is negligible compared to the depression induced by spike pairs.

Results with the minimal triplet model show good agreement with experimental

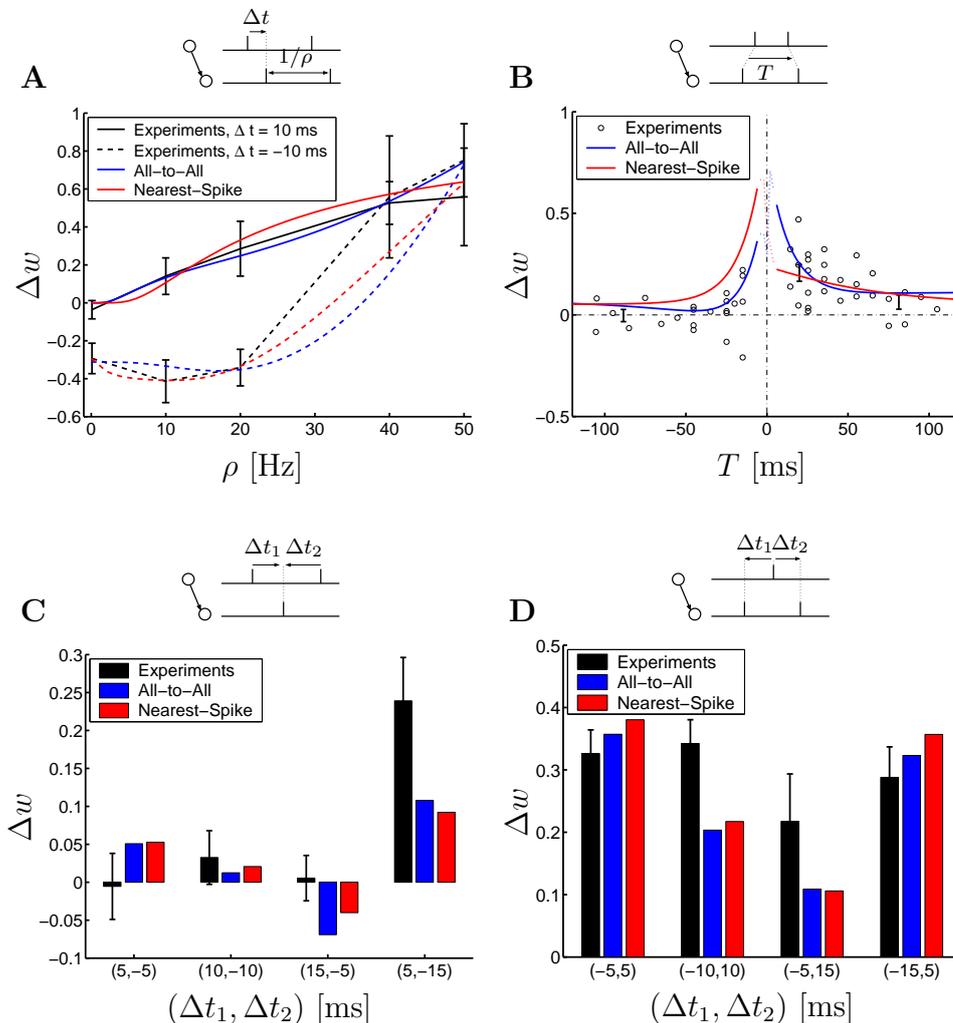


Figure 4: The full triplet learning rule succeeds to reproduce the pairing experiment and most of the triplet and quadruplet experiments. In all four subgraphs, black lines or circle denote experimental data, blue lines correspond to the All-to-All pair model and the red lines to the Nearest-Spike pair model. **A**. Weight change in a pairing protocol as a function of the frequency ρ (*solid lines*: $\Delta t = +10$ ms, *dashed lines*: $\Delta t = -10$ ms). Black lines and data points (with errors) redrawn from Sjöström (2001). **B**. Quadruplet protocol. Black circles are redrawn from Wang et al. (2005). **C**. Triplet protocol for the *pre-post-pre* case and **D** for the *post-pre-post* case. Black dots in **B** and black bars (and standard errors) in **C** and **D** from Wang et al. (2005). The triplet-based models succeed to reproduce the asymmetry in triplets protocols (no potentiation for $(\Delta t_1, \Delta t_2) = (5 \text{ ms}, -5 \text{ ms})$ in **C** and strong potentiation for $(-5 \text{ ms}, 5 \text{ ms})$ in **D**: for those triplets the model results (with All-to-All interactions) are within 1.1σ (standard error of experimental data) whereas the results of the pair-based models are off by more than 4σ

data (see Fig. 5A). Hence the minimal model with 5 parameters can explain the visual cortex data that the classical pair-based STDP model with 4 parameters fails to explain.

A. Visual cortex data set

Model		A_2^+	A_3^+	A_2^-	A_3^-	τ_x [ms]	τ_y [ms]	E
All-to-All	full	$5 \cdot 10^{-10}$	$6.2 \cdot 10^{-3}$	$7 \cdot 10^{-3}$	$2.3 \cdot 10^{-4}$	(101)	125	0.33
	min.	0	$6.5 \cdot 10^{-3}$	$7.1 \cdot 10^{-3}$	0	-	114	0.34
Nearest-Spike	full	$8.8 \cdot 10^{-11}$	$5.3 \cdot 10^{-2}$	$6.6 \cdot 10^{-3}$	$3.1 \cdot 10^{-3}$	714	40	0.22
	min.	0	$5 \cdot 10^{-2}$	$8 \cdot 10^{-3}$	0	-	40	0.34

B. Hippocampal culture data set

Model		A_2^+	A_3^+	A_2^-	A_3^-	τ_x [ms]	τ_y [ms]	E
All-to-All	full	$6.1 \cdot 10^{-3}$	$6.7 \cdot 10^{-3}$	$1.6 \cdot 10^{-3}$	$1.4 \cdot 10^{-3}$	946	27	2.9
	min.	$5.3 \cdot 10^{-3}$	$8 \cdot 10^{-3}$	$3.5 \cdot 10^{-3}$	0	-	40	3.4
Nearest-Spike	full	$4.6 \cdot 10^{-3}$	$9.1 \cdot 10^{-3}$	$3 \cdot 10^{-3}$	$7.5 \cdot 10^{-9}$	(575)	47	2.9
	min.	$4.6 \cdot 10^{-3}$	$9.1 \cdot 10^{-3}$	$3 \cdot 10^{-3}$	0	-	48	2.9

Table 3: List of parameters used to model for **A** the visual cortex data set and **B** the hippocampal culture data set. In this table the terms “full” and “min.” stand for full triplet model and minimal triplet model respectively. The additional parameters $\tau_+ = 16.8$ ms and $\tau_- = 33.7$ ms are taken from Bi and Poo (2001) and kept fixed for all models and data sets. In some cases, parenthesis are added in the τ_x column in order to indicate that the error function is insensitive to the exact value of τ_x in those cases. The last column corresponds to the fitting error given by Eq. (5) and plotted in Fig. 6.

3.5 Triplet Learning Rules can Reproduce Triplet and Quadruplet Experiments

By following the same procedure as the one described in the previous paragraph, we applied our full triplet learning rule to the second set of data, i.e. the hippocampal culture data set (Bi and Poo 1998; Bi and Poo 2001; Wang et al. 2005). The parameters resulting from the minimization of the normalized mean square error across the pair, triplet and quadruplet data are summarized in table 3.

Our triplet learning rule does not only reproduce the classical STDP learning window (see Fig. 3), but it also captures the results of most of the triplet and the quadruplet experiments. See Fig. 4B-D. For example, the asymmetry between the pre-post-pre ($(\Delta t_1, \Delta t_2) = (5 \text{ ms}, -5 \text{ ms})$) and the post-pre-post ($(\Delta t_1, \Delta t_2) = (-5 \text{ ms}, 5 \text{ ms})$) triplets can be well captured by our model. For those two specific

triplet protocols, the predicted weight change of the full triplet learning rule with All-to-All interactions is within 1.1σ (standard error on the mean) off the experimental mean weight change whereas the pair-based learning predictions are off by more than 4σ . We should, however, note that even if our triplet learning rule captures most of the triplet experiments, the fit is not perfect. For example, the pre-post-pre (with $(\Delta t_1, \Delta t_2) = (5 \text{ ms}, -15 \text{ ms})$) triplet experiment is not well reproduced by our triplet learning rule (see Fig. 4C and 5C).

With arguments similar to those applied in section 3.4, it is possible to reduce the complexity of the model. Specifically, we have set $A_3^- = 0$ as before. However, in contrast to the above minimal model for visual cortex data, the pair-term controlled by A_2^+ is kept as part of the model, since it is necessary to explain the potentiation at 1 Hz repetition frequency. The resulting weight change of the minimal model applied to the triplet and quadruplet experiments is depicted in Fig. 5B-D. We emphasize that the minimal model for the hippocampal data is different from the one used for the visual cortex data.

In order to compare the pair models and the minimal and full triplet models, we plotted the fitting error given by Eq. (5) as a function of the number of parameters in the model (see Fig. 6). The best types of model are those who can predict the experimental data as well as possible while being as simple as possible (i.e. having as few parameters as possible). In this sense, the minimal models are the best since they perform almost as well as the full triplet models while having only one extra parameter compared to standard pair-based models (two extra parameters for the hippocampal culture data set).

Finally, for future test of the triplet models, we propose two new protocols that have not yet been used experimentally. The first protocol consists of pre-post-pre triplets with relative timing $(\Delta t_1, \Delta t_2) = (5 \text{ ms}, -5 \text{ ms})$ and the second protocol consists of post-pre-post triplets with relative timing $(\Delta t_1, \Delta t_2) = (-5 \text{ ms}, 5 \text{ ms})$. Triplets are repeated 60 times at different frequencies ρ . Fig. 6C and D depict the weight change predicted by the minimal triplet models (with All-to-All and Nearest-Spike interactions) for the two triplet protocols. The models predict a frequency dependence with a positive slope. However, the overall level of potentiation predicted by the All-to-All model is clearly different from that of the Nearest-Spike interaction model. Thus the above experimental protocol would allow to test the triplet models and distinguish between its two variants.

3.6 Triplet Learning Rule Can Be Mapped to the BCM Learning Rule

Functional consequences of our new triplet model can be studied in two different ways, i.e. analytically or by numerical simulations. We use a combination of the two and proceeded as follows. First, we show analytically a close analogy ("mapping") between our triplet model and the traditional BCM theory. As a result of this mapping, we may conclude that, under random spike arrival with rate ρ_x , our triplet model behaves as a BCM model and inherits all its functional properties. In

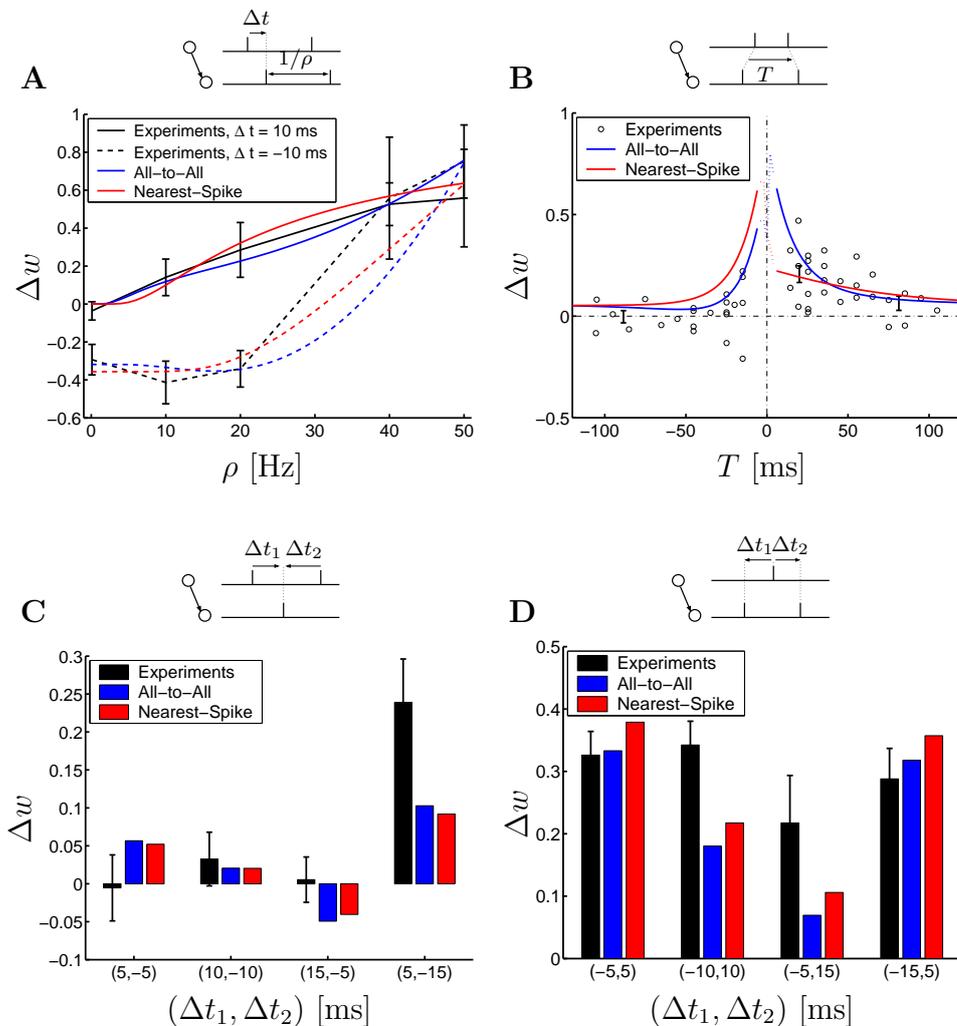


Figure 5: Minimal triplet learning rules are almost as good as full triplet learning rules. In all four subgraphs, black line or circle denote experimental data, blue lines correspond to the All-to-All pair model and the red lines to the Nearest-Spike pair model. **A.** Weight change in a pairing protocol as a function of the frequency ρ (*solid lines*: $\Delta t = 10$ ms, *dashed lines*: $\Delta t = -10$ ms). Black lines and data points (with errors) redrawn from Sjöström (2001). **B.** Quadruplet protocol. Black circles are redrawn from Wang et al. (2005). **C.** Triplet protocol for the *pre-post-pre* case and **D** for the *post-pre-post* case. Black dots in **B** and black bars (and standard errors) in **C** and **D** from Wang et al. (2005).

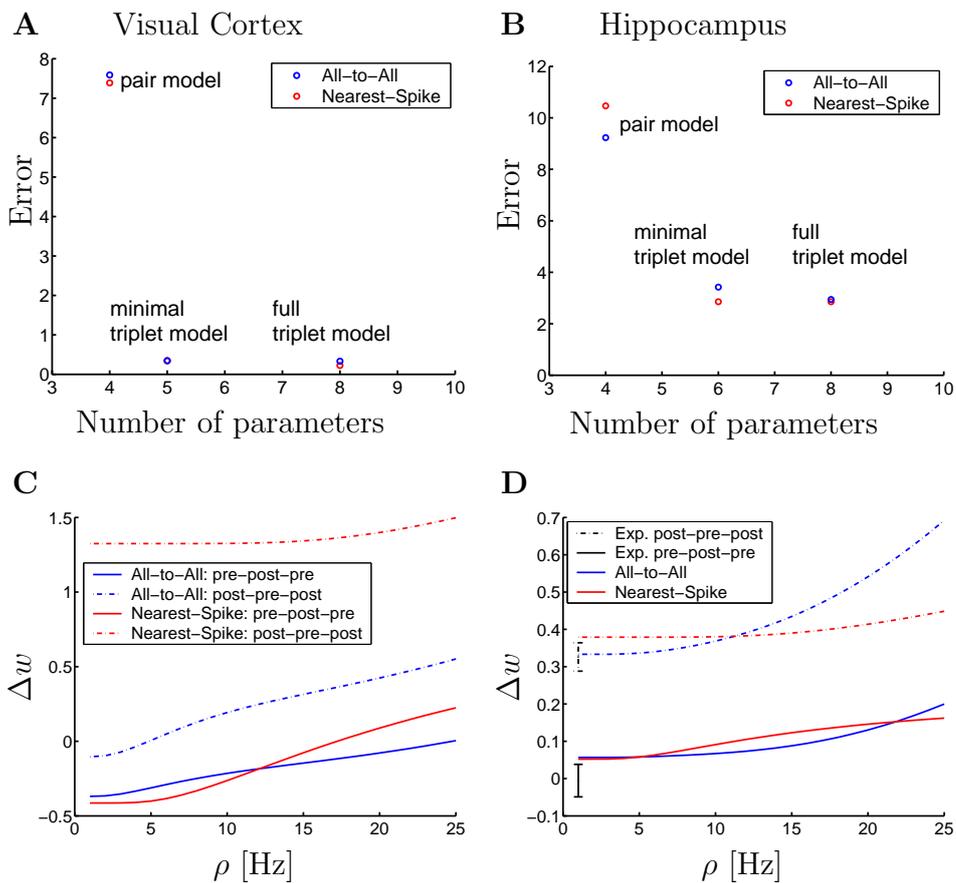


Figure 6: **A** and **B**: Comparison between the pair and triplet models. **C** and **D** Predictions of the triplet models. **A**. Fitting error (c.f. Eq. (5)) for the visual cortex data set of Sjöström et al. (2001) as a function of the number of parameters in the model. The minimal model has only one extra parameter compared to a pair-based model but performs more than 20 times better. **B**. Fitting error for the hippocampal data set of Wang et al. (2005). **C** Predicted weight change (*visual cortex*) of the triplet protocol (*solid lines*: pre-post-pre with $(\Delta t_1, \Delta t_2) = (+5, -5)$ ms, *dot-dashed lines*: post-pre-post with $(\Delta t_1, \Delta t_2) = (-5, +5)$ ms) with All-to-All interactions (*blue lines*) and with Nearest-Spike interactions (*red lines*). **D**. Same as in **C** but for the hippocampal culture data set. Black bars correspond to the experimental results also present in subplots C and D of Figs. 2, 4 and 5.

particular, we expect that our triplet model exhibits synaptic competition leading to input selectivity as required for receptive field development. In a second step, we have tested this prediction of input selectivity by numerical simulation.

First, we show that unlike standard pair-based STDP learning rules, our triplet learning rule can be mapped to the BCM learning rule. If we assume that the pre- and postsynaptic spike trains have Poisson statistics with ρ_x and ρ_y respectively as firing rate, the expected weight change can be calculated analytically. Intuitively, we may expect that a triplet term with one pre and two postsynaptic spikes leads to a weight change that is proportional to the postsynaptic rate and the *square* of the presynaptic rate. An analogous argument holds for the other terms. Indeed, a detailed calculation for the All-to-All triplet learning rule based on Eqs. (1)-(4) yields an expected weight change:

$$\left\langle \frac{dw}{dt} \right\rangle = -A_2^- \tau_- \rho_x \rho_y - A_3^- \tau_- \tau_x \rho_x^2 \rho_y + A_2^+ \tau_+ \rho_x \rho_y + A_3^+ \tau_+ \tau_y \rho_x \rho_y^2 \quad (6)$$

Fig. 7A depicts the expected weight change of Eq. (6) as a function of the postsynaptic frequency ρ_y . The above weight dynamics can be written as a BCM learning rule. Indeed, the BCM theory requires first that the weight change can be written as $dw/dt = \phi(\rho_y, \theta) \rho_x$ where ϕ is such that $\phi(\rho_y < \theta, \theta) < 0$, $\phi(\rho_y > \theta, \theta) > 0$ and $\phi(0, \theta) = 0$. Our Eq. (6) can satisfy this condition if $A_3^- = 0$ as is the case for our minimal triplet models. The second requirement is that the threshold θ between potentiation and depression is proportional to the expectation of the p^{th} power of the postsynaptic firing rate, i.e. $\theta = \alpha \bar{\rho}_y^p$ where $p > 1$ (Bienenstock, Cooper, and Munro 1982; Intrator and Cooper 1992). This second requirement can be fulfilled if the parameters A_2^- and A_2^+ depend on the mean firing rate $\bar{\rho}_y$ (or powers thereof) of the postsynaptic neuron. Specifically, we set $A_2^- \rightarrow A_2^- \bar{\rho}_y^p / \rho_0^p$ as well as $A_2^+ \rightarrow A_2^+ \bar{\rho}_y^p / \rho_0^p$. By doing so the threshold becomes $\theta = \bar{\rho}_y^p (A_2^- \tau_- - A_2^+ \tau_+) / (\rho_0^p A_3^+ \tau_+ \tau_y)$.

Strictly speaking, $\bar{\rho}_y^p$ corresponds to the expectation over the input statistics of the p^{th} power of the postsynaptic firing rate. Practically, this quantity can be evaluated online by low-pass filtering ρ_y^p with a time constant of the order of 10 minutes or more. With this range of time scale, $\bar{\rho}_y^p$ can be considered as constant (i.e. $\bar{\rho}_y^p \simeq \rho_0^p$) over the duration of the pairing, triplet and quadruplet protocols we used in this paper. As an aside, we note that with Nearest-Spike interactions, our triplet learning rule can almost (but not strictly) be mapped to a BCM learning rule.

Since the triplet rule shares properties with BCM theory, we expect that it generates input selectivity if a neuron receives a large number of inputs. Development of input selectivity is thought to be an important property to account for receptive fields development.

For a numerical illustration of the input selectivity property of the triplet learning rule, we simulated the following scenario. We assume that our model neuron receives 100 afferents ($1 \leq i \leq 100$) which are stimulated with Gaussian profiles $\nu_i = 1 \text{ Hz} + 50 \text{ Hz} \exp(-(i - \mu)/2 \cdot 10^2)$, $i = 1, \dots, 100$ whose center μ are shifted

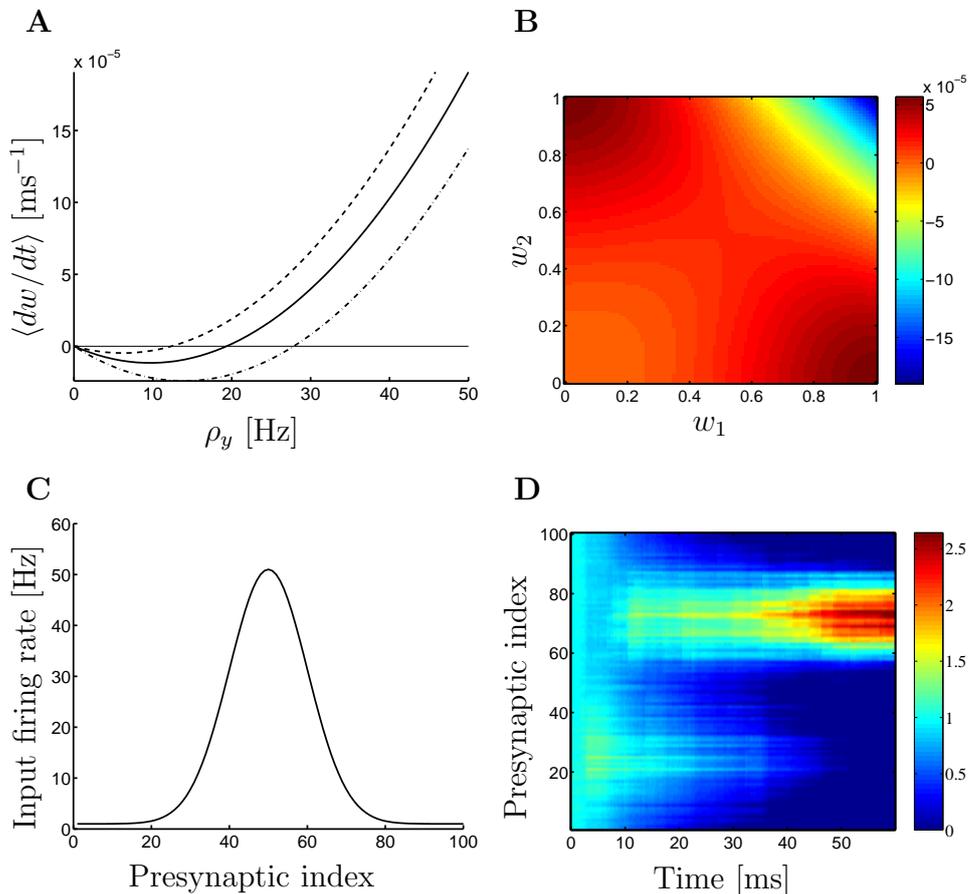


Figure 7: The triplet learning rule can be mapped to a BCM learning rule. **A.** Instantaneous weight change as a function of the postsynaptic frequency for a minimal triplet model. (c.f. Eq. (6) with $A_3^- = 0$). The pre- and postsynaptic spike trains are Poisson spike trains. The *dashed line* corresponds to $\lambda = \rho_y^p / \rho_0^p = 0.64$, *solid line*: $\lambda = 1$ and *dot-dashed line*: $\lambda = 1.44$. **B.** Energy landscape produced by the minimal triplet learning rule (with $p = 2$ and $\rho_0 = 10$ Hz) in a two-input environment: $\rho_x^1 = (10 \text{ Hz}, 0)^T$ and $\rho_x^2 = (0, 10 \text{ Hz})^T$. The presence of two specialized (and stable) fixed point as well as two unspecialized (and unstable) fixed points is an essential feature of the BCM learning rule. **C.** Gaussian stimulation profile across 100 presynaptic neurons. The center of the Gaussian is shifted randomly every 200 ms to one of 10 random positions. Periodic boundary conditions are assumed. **D.** Evolution of the 100 weights as a function of time under the stimulation described in C. After one minute of stimulation, the postsynaptic neuron becomes sensitive to a stimulation centered around the 70th presynaptic neuron. The parameters taken in the minimal model are those which correspond to the visual cortex; c.f. table 3.

randomly every 200 ms (see Fig. 7C) over 10 possible positions. Presynaptic spikes are generated at time t_i^f with a rate ν_i . Each presynaptic spike generates an exponential postsynaptic potential with decay time constant $\tau = 10$ ms, so that the total potential is $u = \sum_i \sum_{t_i^f < t} w_i \epsilon_0 \exp(-(t - t_i^f)/\tau)$ where $\epsilon_0 = 1$ mV. The postsynaptic firing rate increases with the membrane potential according to $\nu^{\text{post}} = 1 \text{ Hz} + g \cdot u$ where $g = 10 \text{ Hz/mV}$. The neuron is stimulated over 60 seconds while synapses change according to our triplet learning rule. As we can see in Fig. 7D, the neuron becomes automatically specialized to one of the 10 input patterns, i.e. the one with $\mu = 70$. In other words, learning leads to input selectivity, a necessary property for receptive field development.

It is interesting to note that the dynamics of Eq. (6) can be seen as a gradient ascent of an objective function L , i.e. $\Delta w \propto \partial L / \partial w$. Let $p = 2$ and $\beta = A_3^+ \tau_+ \tau_y$. L can be written as $L = \frac{\beta}{3} \bar{\rho}_y^3 - \frac{\beta}{\alpha} \theta^2$. If the model neuron has only two input afferents and hence only two synapses, this objective function L (or energy landscape) (see Fig. 7B) elicits two selective points which correspond to the two maxima of the function L . The first maximum is at $w_1 = 1$ and $w_2 = 0$ and the second at $w_1 = 0$ and $w_2 = 1$. The pattern of synaptic weight corresponds therefore to input selectivity, i.e., the neuron is sensitive to only one of two inputs. Thus the objective function can be used for a mathematical demonstration of input selectivity. Note that the objective function exists only if we assume that ϕ is a function of $\bar{\rho}_y^p$ and not $\bar{\rho}_y^p$ (Cooper, Intrator, Blais, and Shouval 2004).

4 Discussion and Conclusion

In this paper we first showed the limitations of the standard pair-based STDP models in terms of predicting the outcome of several spike-timing based protocols. We then showed that a triplet learning rule is more suitable to reproduce those experimental protocols, namely, the frequency dependence of the pairing protocol as well as the triplet and quadruplet protocols. Finally, we showed the link between our triplet learning rule and the BCM learning rule. We find noteworthy and somewhat unexpected that our detailed modeling of frequency dependence of pair-based protocols and asymmetries in triplet protocols should lead under the assumption of Poisson spike trains to a known theoretical rule with well characterized features.

Throughout this paper, we compared the All-to-All interactions versus the Nearest-Spike interactions for pair-based models as well as for triplet-based models. Even though Nearest-Spike interactions induce some potentially interesting non-linearities in pair-based models (van Rossum, Bi, and Turrigiano 2000; Izhikevich and Desai 2003; Burkitt, Meffin, and Grayden 2004) (especially in the Poisson protocol), it is not possible to make a strict mapping of Nearest-Spike interactions models to the BCM rule and more importantly pair-based models with Nearest-Spike interactions fail to reproduce the correct frequency dependence in the pairing protocol as well as triplet and quadruplet experiments.

Limitations

Even if our triplet model can capture most of the triplet and quadruplet experiments, it is necessary to keep in mind the kind of experiments this model cannot reproduce. Since our model predicts weight changes as a function of the spike timing only, it fails to make any kind of inference for experiments which trigger explicitly other biophysical parameters such as Ca^{2+} concentration or postsynaptic membrane potential. We nevertheless think that this approach is interesting since in one way or another those biophysical parameters depend on the timings of the pre- and postsynaptic spikes. For example, the calcium concentration depends on the timing of the postsynaptic spike (via the back-propagating action potential) and the presynaptic spike (via voltage-gated calcium channels and NMDA channels).

Recent experiments (Froemke, Poo, and Dan 2005) show that the shape of the depression part of the learning window depends on the position of the synapse on the dendritic tree. Although our model does not include such geometrical properties, it is possible to account for the position of the synapse by changing explicitly the time constant τ_- (characterizing the LTD part of the learning window) as a function of the distance between the synapse and the soma.

Even in the context of typical STDP experiments, some aspects are not covered by our model. In most STDP experiments, plasticity is induced after a repetition of a fixed number of pairs of pre- and postsynaptic spikes. Clearly the amount of plasticity depends on the number of pairs. In fact, the amount of potentiation increases with the number of pairs of pre- and postsynaptic spikes and saturates at a given value (Senn, Markram, and Tsodyks 2001; Froemke, Tsay, Raad, Long, and Dan 2006). This saturation is not taken into account in our present model since the weight dependence is not explicitly mentioned. The dependence upon the weights can easily be added in the triplet models (the parameters A_2^+ , A_2^- , A_3^+ and A_3^- could also depend on w). Even if we have some indications (Bi and Poo 2001; Wang, Gerkin, Nauen, and Bi 2005) of how synapses change as a function of w , more experimental data are clearly needed in order to determine the correct weight dependence. It should be noted that if we add the dependence upon the weight, there would not be an unambiguous mapping to the BCM theory.

In our present study with the minimal model, we needed to fit 5 parameters for the first preparation (slice of visual cortex) or 6 parameters for a different preparation (hippocampal culture), so we had in total 11 parameters to tune. It could be interesting to define a single model (with less than 11 parameters) where some parameters, but not all, are shared between the two preparations.

Alternative Interpretations of the Experimental Data

The goal of this paper is to go as far as possible in the prediction of the weight change with only spike timing and no other neuronal variables or mechanism. It is interesting to note that our triplet learning rule can reproduce both data that has been explained by a postsynaptic potential effect (Sjöström, Turrigiano, and Nelson 2001) or by a suppression effect (Wang et al. 2005). In Sjöström's experiment, the increased potentiation at high frequency is explained by the increased membrane potential due to the accumulation of presynaptic inputs whereas in our

model the increased potentiation is due to the increase of the postsynaptic variable o_2 . Combined with a suitable neuron model, an increased frequency would of course yield a higher potential on average. Wang et al. (2005) interpreted their triplet and quadruplet experiment as a result of a suppression mechanism, i.e. if a pre-post pair is followed by a post-pre pair, the latter depression term suppresses the first potentiation term (and not the other way round). This phenomenon is captured in our framework by the extra potentiation due to the 1-pre-2-post triplet term.

Expansion Perspective

It is possible to see pair terms and triplet terms of Eqs. (3) and (4) in a more general framework (Gerstner and Kistler 2002). From the point of view of pure spike-timing dependent plasticity, we can say that the instantaneous weight change is an unknown functional of the presynaptic spike train X and the postsynaptic spike train Y , i.e. $dw/dt = H[X, Y]$. In this framework, a pair-based STDP learning rule corresponds to the Volterra expansion of H to the second order. In this article, we pushed the expansion to the third-order and showed that the prediction power increased a lot with only one or two extra parameters.

It is interesting to note that quadruplets data can be fitted with a triplet rule. This suggests that third-order terms (triplet terms) are good enough and therefore there is no need to take into account higher-order terms. It is of course possible that new experiments will show the limitations of a triplet model and force us to consider higher-order terms. Clearly, the relevance of such an approach depends on how far we have to go in the expansion.

The learning rule as it is now does not depend directly on the membrane potential and therefore cannot reproduce the experiments of Sjöström, Turrigiano, and Nelson (2001) in which the membrane potential is controlled by the experimentalist. We should however note that in the framework of the Volterra expansion of the unknown functional H , it is possible to assume that H depends explicitly on the membrane potential (Gerstner and Kistler 2002) and therefore capture the voltage dependence experiments of Sjöström, Turrigiano, and Nelson (2001).

In this paper, we have considered only one pre- and one postsynaptic neuron, in order to be simple and to make the synaptic change a function of locally available quantities. There is now partial evidence (see (Bi and Poo 2001) for a review) that synaptic plasticity does not only depend on the state of the pre- and postsynaptic neuron. Neighboring synapses can have a direct influence on a given synapse leading to heterosynaptic plasticity. This could be included in this framework by making the functional H depend on all presynaptic firing spike trains, i.e. $H[X_1, \dots, X_M, Y]$.

Comparison to Other Models

It is interesting to note that our triplet learning rule has some similarities with the Senn-Markram-Tsodyks (SMT) model (Senn, Markram, and Tsodyks 2001). In their model, they need a first presynaptic spike in order activate a fraction of NMDA channels, then a postsynaptic spike to set some secondary messengers in an up-state and finally a last postsynaptic spike in order to trigger synaptic potentiation. Thus

their rule essentially consists of a triplet pre-post-post term for potentiation and a triplet post-pre-pre term for depression. Even if their model makes implicitly use of triplet terms, it should be noted that the order of spikes is different compared to our present triplet model. In our present triplet model we need one pre and two postsynaptic spikes regardless of the order, i.e. it encompasses pre-post-post as well as post-pre-post triplets of spikes. This difference is of particular importance if we want to fit the post-pre-post triplet experiments or quadruplet experiments of Wang et al. (2005). For those protocols the SMT model cannot reproduce the data.

Another model which takes into account a multi-spike interaction is the Froemke-Dan learning rule (Froemke and Dan 2002). Their model (which is in fact a quadruplet model) predicts a synaptic behavior which is in direct contrast with the synaptic dynamics given by Eqs. (3) and (4). In their model, if a postsynaptic spike precedes a pre-post pair of spikes, the effective potentiation will decrease as soon as the two postsynaptic spikes get closer to each other, whereas in our model, the opposite occurs. This is the main reason why the Froemke-Dan model, under a Poisson assumption for the pre- and postsynaptic spike trains, predicts an increasing depression as the postsynaptic firing rate increases, as reported by (Izhikevich and Desai 2003) which seems unplausible in view of the results in Fig. 8C of Sjöström et al. (2001). It should be noted that the *revised suppression model* of Froemke, Tsay, Raad, Long, and Dan (2006) gets around this problem by setting two different saturation values: one for depression and one potentiation.

Up to this point we have not made any assumption about the cellular processes that are described by our triplet learning rule. It is known that the amount of potentiation or depression expressed by a synapse depends critically on the concentration of calcium. Moreover we know that there is a supralinear summation of calcium on the postsynaptic site when an EPSP precedes an AP (Waters, Larkum, Sakmann, and Helmchen 2003). From this point of view, two closely spaced postsynaptic spikes can increase the level of calcium and therefore increase potentiation. This would correspond to the 1-pre-2-post triplet term of our formalism.

In this study, we showed that a minimal triplet model can capture most, but not all, aspects of the pairing, triplet and quadruplet experiments. A natural extension of this study could be to include explicitly the dependence upon biophysical quantities such as the Ca^{2+} concentration, the postsynaptic membrane potential or other neuronal quantities. A more appealing approach would be to consider existing detailed biophysical models of synaptic plasticity and try to reduce them to a triplet model and therefore identify the underlying biological quantities of our triplet model.

A Analytical Calculations for Nearest-Spike Interactions

The predicted weight change of a given model for a given protocol can be either simulated numerically or calculated analytically. As an example of analytical calcu-

lation, we are showing here the explicit expression of the weight change predicted by a triplet learning rule with Nearest-Spike interactions for a pairing protocol. The explicit expression of the weight change is simply obtained by counting the number pairs and triplets with fixed intervals. For example, in a pairing protocol with N pairs of pre- and postsynaptic spikes repeated at a frequency ρ with a positive delay Δt , there are N pre-post pairs separated by Δt and $N - 1$ post-pre pairs separated by $\rho^{-1} - \Delta t$ where ρ^{-1} is the interval between pairs. The weight change Δw gives

$$\begin{aligned} \Delta w &= (N - 1)F_2(\rho^{-1} - \Delta t) + (N - 1)F_3(\rho^{-1}, \rho^{-1} - \Delta t) \\ &+ NG_2(\Delta t) + (N - 1)G_3(\Delta t, \rho^{-1}) \end{aligned} \quad (7)$$

For negative Δt , an analogous argument applies which gives

$$\begin{aligned} \Delta w &= NF_2(-\Delta t) + (N - 1)F_3(\rho^{-1}, -\Delta t) \\ &+ (N - 1)G_2(\rho^{-1} + \Delta t) + (N - 1)G_3(\rho^{-1} + \Delta t, \rho^{-1}) \end{aligned} \quad (8)$$

where $F_2(s) = -A_2^- \exp(-s/\tau_-)$ and $G_2(s) = A_2^+ \exp(-s/\tau_+)$ denote the pair-based depression kernel and pair-based potentiation kernel respectively, $F_3(s, s') = -A_3^- \exp(-s/\tau_x) \exp(-s'/\tau_-)$ and $G_3(s, s') = A_3^+ \exp(-s/\tau_+) \exp(-s'/\tau_y)$ denote the triplet-based depression kernel and triplet-based potentiation kernel respectively.

The weight change for other combinations of model and protocol can be calculated in the same way by counting the number of pairs and triplets with given intervals.

B Poisson Statistics with Nearest-Spike Interactions

In the Nearest-Spike interaction, we have almost a BCM behavior. If we assume Poisson statistics, i.e. an interval distribution $P_x(s) = \rho_x \exp(-\rho_x s)$ for the presynaptic spikes and $P_y(s) = \rho_y \exp(-\rho_y s)$ for the postsynaptic spikes, we can calculate the expectation of the pair terms and the triplet terms of Eqs. (3) and (4), where ρ_x and ρ_y correspond to the pre- and postsynaptic firing rate respectively. For example, for each postsynaptic spike, the triplet potentiation term gives:

$$A_3^+ \int_0^\infty P_y(s) e^{-s/\tau_y} ds \int_0^\infty P_x(s) e^{-s/\tau_+} ds = \frac{A_3^+ \rho_x \rho_y}{(\rho_y + \alpha_y)(\rho_x + \alpha_x)} \quad (9)$$

where $\alpha_+ = \tau_+^{-1}$ and $\alpha_y = \tau_y^{-1}$. All other terms can be calculated in the same way. This gives

$$\begin{aligned}
\left\langle \frac{dw}{dt} \right\rangle &= -A_2^- \frac{\rho_x \rho_y}{\rho_y + \alpha_-} - A_3^- \frac{\rho_x^2 \rho_y}{(\rho_x + \alpha_x)(\rho_y + \alpha_-)} \\
&+ A_2^+ \frac{\rho_x \rho_y}{\rho_x + \alpha_+} + A_3^+ \frac{\rho_x \rho_y^2}{(\rho_y + \alpha_y)(\rho_x + \alpha_+)}
\end{aligned} \tag{10}$$

where $\alpha_- = \tau_-^{-1}$ and $\alpha_x = \tau_x^{-1}$. If we do the same assumption as we did with the All-to-All interaction scheme, i.e. $A_3^- = 0$ and we further assume that $\rho_x \ll \alpha_+$ (in our case $\rho_x \ll 60Hz$), we can satisfy the first constraint of a BCM learning rule, i.e. $dw/dt = \phi(\rho_y, \theta)\rho_x$, but we cannot satisfy the second one. It is impossible to redefine the amplitude parameters in such a way that the threshold becomes proportional to $\bar{\rho}_y^p$. Note that this does not mean that the learning rule Eq. (10) cannot elicit the properties of a BCM learning rule, such as input selectivity. In fact, this learning rule has still the input selectivity property, but only for a limited range of input frequencies.

References

- Abarbanel, H., R. Huerta, and M. Rabinovich (2002). Dynamical model of long-term synaptic plasticity. *Proc. Natl. Academy of Sci. USA* 59, 10137–10143.
- Artola, A., S. Bröcher, and W. Singer (1990). Different voltage dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347, 69–72.
- Bi, G. and M. Poo (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Bi, G. and M. Poo (2001). Synaptic modification of correlated activity: Hebb's postulate revisited. *Ann. Rev. Neurosci.* 24, 139–166.
- Bi, G. and H. Wang (2002). Temporal asymmetry in spike timing-dependent synaptic plasticity. *Physiology and Behavior* 77, 551–555.
- Bi, G.-Q. (2002). Spatiotemporal specificity of synaptic plasticity: cellular rules and mechanisms. *Biological Cybernetics* 319-332.
- Bienenstock, E., L. Cooper, and P. Munro (1982). Theory of the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* 2, 32–48. reprinted in Anderson and Rosenfeld, 1990.
- Burkitt, A. N., M. H. Meffin, and D. Grayden (2004). Spike-timing-dependent plasticity: The relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Computation* 16, 885–940.

- Cooper, L. N., N. Intrator, B. S. Blais, and H. Z. Shouval (2004). *Theory of Cortical Plasticity*. Singapore: World Scientific.
- Dudek, S. M. and M. F. Bear (1992). Homosynaptic long-term depression in area ca1 of hippocampus and effects of n-methyl-d-aspartate receptor blockade. *Proc. Natl. Acad. Sci. USA* 89, 4363–4367.
- Froemke, R. and Y. Dan (2002). Spike-timing dependent plasticity induced by natural spike trains. *Nature* 416, 433–438.
- Froemke, R., MM. Poo, and Y. Dan (2005). Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434, 221–225.
- Froemke, R., I. Tsay, M. Raad, J. Long, and Y. Dan (2006). Contribution of individual spikes in burst-induced long-term synaptic modification. *Journal of Neurophysiology* 95, 1620–1629.
- Gerstner, W., R. Kempter, J. L. van Hemmen, and H. Wagner (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383, 76–78.
- Gerstner, W. and W. K. Kistler (2002). *Spiking Neuron Models*. Cambridge University Press.
- Intrator, N. and L. N. Cooper (1992). Objective function formulation of the bcm theory of visual cortical plasticity. *Neural Networks* 5, 3–17.
- Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks* 14, 1569–1572.
- Izhikevich, E. and N. Desai (2003). Relating stdp to bcm. *Neural Computation* 15, 1511–1523.
- Karmarkar, U. and D. Buonomano (2002). A model of spike-timing dependent plasticity: one or two coincidence detectors. *J. Neurophysiology* 88, 507–513.
- Kempter, R., W. Gerstner, and J. L. van Hemmen (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514.
- Kirkwood, A., M. G. Rioult, and M. F. Bear (1996). Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* 381, 526–528.
- Kistler, W. M. and J. L. van Hemmen (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput.* 12, 385–405.
- Lisman, J. (1989). A mechanism for hebb and anti-hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA* 86, 9574–9578.
- Lisman, J. and N. Spruston (2005). Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nature Neuroscience* 8, 839–841.
- Lisman, J. and A. Zhabotinsky (2001). A model of synaptic memory: A CaMKII/PP1 switch that potentiates transmission by organizing an AMPA receptor anchoring assembly. *Neuron* 31, 191–201.

- Malenka, R. C., J. Kauer, R. Zucker, and R. A. Nicoll (1988). Postsynaptic calcium is sufficient for potentiation of hippocampal synaptic transmission. *Science* 242, 81–84.
- Markram, H., J. Lübke, M. Frotscher, and B. Sakmann (1997). Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSP. *Science* 275, 213–215.
- Pfister, J.-P. and W. Gerstner (2006). Beyond pair-based stdp: a phenomenological rule for spike triplet and frequency effects. In Y. Weiss, B. Schölkopf, and J. Platt (Eds.), *Advances in Neural Information Processing Systems 18*, pp. 1083–1090. Cambridge, MA: MIT Press.
- Rao, R. P. N. and T. J. Sejnowski (2001). Spike-timing dependent Hebbian plasticity as temporal difference learning. *Neural Computation* 13, 2221–2237.
- Senn, W., H. Markram, and M. Tsodyks (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation* 13, 35–67.
- Shouval, H. Z., M. F. Bear, and L. N. Cooper (2002). A unified model of nmda receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA* 99, 10831–10836.
- Sjöström, P., G. Turrigiano, and S. Nelson (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164.
- Song, S., K. Miller, and L. Abbott (2000). Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience* 3, 919–926.
- van Rossum, M. C. W., G. Q. Bi, and G. G. Turrigiano (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience* 20, 8812–8821.
- Wang, H. X., R. C. Gerkin, D. W. Nauen, and G. Q. Bi (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nature Neuroscience* 8, 187–193.
- Waters, J., M. Larkum, B. Sakmann, and F. Helmchen (2003). Supralinear ca^{2+} influx into dendritic tufts of layer 2/3 neocortical pyramidal neurons in vitro and in vivo. *The Journal of Neuroscience* 23, 8558–8567.
- Zhang, L., H. Tao, C. Holt, W. Harris, and M.-M. Poo (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395, 37–44.

3.1 Summary of the Results

Synapses are dynamic and change as function of locally available quantities. In this thesis, we developed several models on synaptic plasticity that mainly depend on the timings of pre- and postsynaptic spikes as locally available quantities. Are there fundamental principles that help us to understand why synaptic strength changes as a function of the time difference between the pre- and postsynaptic spike in a way depicted by Fig. 1.1? What is the synaptic functional role which governs its dynamics? This question was addressed by the first five papers. The two first papers dealt with supervised scenarios, while the papers III - V studied unsupervised scenarios.

The second question we asked in this thesis was the following. What is the simplest description of STDP experiments? Is this description consistent with experiments that consists of pairs of spikes, triplets or quadruplets of spikes? Is the description still valid if the repetition frequency of the pairs increases? Those questions were addressed in papers VI and VII.

In the next sections, we will review the main results of the optimal models (papers I-V) and the phenomenological models (papers VI-VII). We will then discuss the limitations of our approaches and finish with some propositions of future developments.

STDP is the Best Rule that Achieves Temporal Precision

In the supervised scenarios of the two first papers, we have shown that the best learning rule that maximizes the probability of firing at precise times, while not firing (or not firing more than background activity) at other times is similar to the STDP learning window. In the last scenario of paper II, we have shown that STDP makes the postsynaptic neurons, the best detector neurons. Indeed, we have shown that the learning rule that maximizes the probability of firing at given times (and not at the other times) when a specific input pattern is presented and minimizes the probability of firing when other inputs are presented, is similar to the STDP learning rule if an extra locality constraint is assumed.

As mentioned in the discussion part of paper II, the main interest of such an approach is the link to reinforcement learning. If $\langle R \rangle_{\mathbf{x},y} = \sum_{\mathbf{x},y} R(\mathbf{x},y)P(y|\mathbf{x})P(\mathbf{x})$ denotes the expected reward that should be maximized through learning, where $P(y|\mathbf{x})$ is the probability of having a postsynaptic spike train y for a given input \mathbf{x} , it is straightforward to calculate the gradient of $\langle R \rangle_{\mathbf{x},y}$ with the formalism of paper I. A similar formalism has already been proposed by Xie and Seung (2004) in the context of reinforcement learning, but they did not calculate the corresponding optimal learning window. See also appendix B.1 for a detailed description of the relation to the tempotron (Gütig and Sompolinsky 2006).

Infomax + Homeostatic Constraint Leads to BCM

We have shown in paper IV, that the maximization of information between input and output under an homeostatic constraint leads to the BCM learning rule. Even if the relation between infomax and BCM has already be made (Nadal and Parga 1997), our approach sheds some light on the problem. Indeed, in our framework, the homeostatic constraint is the essential element which drives the non-linear dynamics of the BCM threshold.

Moreover, the learning rule of paper IV has interesting properties that cannot be captured by the BCM learning rule. Indeed, the BCM learning rule is rate-based learning rule and can not elicit properties such as sensitivity to spike-spike correlations which is captured by the learning rule of paper IV.

Infomax + Constraint Leads to STDP

In the unsupervised scenarios considered in papers III and V, we have shown that the infomax principle (Linsker 1988) leads to the STDP learning rule. More precisely, in paper III, we have shown that in the limit of small fluctuations of the membrane potential, the infomax principle leads to a weight change which is composed of a positive bias and a term which depends on the time difference between the pre- and postsynaptic spike. This last term can be compared to the STDP learning window.

In paper V, extra constraints are added in the objective function. We have shown in this paper that the maximization of an information term and a quadratic weight decay term is sufficient to elicit a learning rule which is consistent with STDP. The weight decay was introduced to remove the positive bias induced by the information term. Note that in this paper we also assumed a homeostatic constraint, but this does not contribute to the derived learning window.

The weight decay term has also an effect on the stability of the fixed points. In addition to the specialized weight patterns (bimodal weight distribution) which are also stable in the BCM formalism, the weight change of V has also a fixed point for non-specialized weight patterns (unimodal weight distribution). This bistability of unimodal and bimodal weight distribution is a specificity of the learning rule of V. The stability of the unimodal weight distribution is useful to prevent the neuron to be selective to an input that came “by chance”.

Triplets are Necessary in Models of STDP

In papers VI and VII, we have shown that the standard pair-based learning rule (which only depends on the time difference between the pre- and postsynaptic spike) fail to reproduce experiments during which the repetition frequency between the pairs is changed. Those pair-based rules also fail to capture triplet and quadruplet experiments.

We have shown that we can construct a triplet model with All-to-All interactions that consists of a pre-post term, a post-pre term and a 1-pre-2-post term. This triplet model can reproduce surprisingly well the frequency dependence experiment and can capture most of the triplet and quadruplet experiments.

Triplet Learning Rule can be Mapped to BCM

Moreover, this triplet learning rule can be mapped to the BCM learning rule if Poisson statistics is assumed for the pre- and postsynaptic neuron. The BCM threshold can be obtained if we assume a specific dependence of the amplitude parameters A_2^- , A_2^+ (c.f. paper VII) on the mean postsynaptic firing rate.

τ_{LTP} Comes from τ_{EPSP} , no Unique Interpretation of τ_{LTD}

In all optimal models of papers I-V, the weight change is derived from a different objective function. This means that the functional role and of course the exact expression of the learning rule differs from one case to the other. It is however interesting to note that in all those five papers, the shape of potentiation part of the learning window is essentially given by the time course of the EPSP. Indeed, the EPSP's implement the causal relation between the pre- and postsynaptic events.

The situation is different for the acausal part of the learning window. Indeed, in each scenario the interpretation of the depression is different. In paper I and in scenarios *A* and *B* of paper II, synapses that transmit a presynaptic spike arriving after the postsynaptic one should be depressed because of the increased probability of emitting a subsequent spike (due to the DAP or the depolarizing teaching signal). So depression occurs in order to prevent the neuron from bursting. In this case, the time constant of the LTD part of the learning window τ_{LTD} is linked to the membrane time constant and eventually to the duration of the teaching signal if it lasts for a long time.

In scenario *C* of paper II, synapses are mainly depressed in order to prevent the neuron to fire when an undesirable input pattern is presented. In this scenario, τ_{LTD} comes from the locality constraint.

In paper III the depression part of the weight change (due to the correlation of pre- and postsynaptic spike) comes essentially from the refractoriness of the neuron. Intuitively, the synapse transmits less information when the neuron is in a refractory state than when the presynaptic spike precedes the postsynaptic one. In paper V, the interpretation is similar. In this paper, the amplitude of an EPSP is suppressed and hence transmits less information if it follows a postsynaptic spike. τ_{LTD} is therefore related to the EPSP suppression time constant.

In all those papers, we it was relatively straightforward to define an objective function that elicits potentiation in the causal part of the learning window and a global offset (the simplest example is illustrated by scenario A of paper II, without teaching signal and with no spike-afterpotential). It was however much more difficult to find scenarios in which depression should be stronger than a given baseline if the presynaptic spike occurs just after the postsynaptic spike. As proposed by Song et al. (2000), it is still possible that a learning window with potentiation in the pre-before-post region and a global offset is sufficient to fulfill most of the functional role of STDP. From this point of view, depression is concentrated in the post-before-pre region not for any fundamental reason but for ease of implementation in the biological substrate.

3.2 Limitations

Panglossian Paradigm

Everything is for the best in the best of all possible worlds.

Voltaire 1759 in *Candide*

Optimism. This is what the philosophical tale *Candide* is about. Candide, the main character of the story, is educated by a tutor, Dr. Pangloss who teaches him a very simplistic and optimistic theory: “Everything is for the best in the best of all possible worlds”. This approach which states that “everything is for the best” or is in some sense optimal has been termed as *Panglossian paradigm* or *adaptationism* (Gould and Lewontin 1979).

Since most of this thesis deal with optimal models, let us discuss now the limitations of those models. There is first a methodological concern we have to address: it is simple to define an objective function for which a specific feature could be optimal. It is certainly more difficult to find an objective function that can capture many different features. This point is somehow related to the overfitting problem. If there are too many different possibilities to explain a given phenomenon, in our case STDP, it is necessary to add constraints, i.e. consider other phenomena that have also to be explained by the same rule.

A common mistake (see the criticism of Gould and Lewontin (1979)) is to state that “because of evolution or natural selection, synaptic dynamic (or any kind of biological mech-

anism) are in some sense optimal”. This hazardous statement relies on many implicit assumptions. First, it assumes that the organisms can be considered as an assemblage of features (or *traits*) that can be optimized independently. Second, it considers natural selection as the only (or the major) actor that forges the properties of the organisms and thirdly it fails to distinguish the current utility of a feature from its historical reasons for origin.

Finally, a last remark about the optimal models that use the mutual information as an objective function. Information theory has been developed to describe key features of man-made communication systems and has demonstrated real efficiency in this domain, but we should be cautious when applying it to living systems.

As a summary, it is really important to be careful when interpreting the results of those optimal models. Even if, by definition, they do not have a great power to predict the outcome of new experiments, they can still be of great interest by giving some general principles or general guidelines which can help “not to be lost in a mass of irrelevant details” (Barlow 1961).

3.3 Propositions for Future Developments

In paper V, we derived a learning rule which has a bistability property. Both the unimodal and the bimodal weight distribution are stable. An interesting line of research would be to look for a minimal model that can reproduce this bistability in order to isolate the essential elements that cause this bistability. The next step would be then to add biological relevance of such a model and hence determine what are the biological parameters that need to be changed in order to get this bistability. Can the neuron dynamically switch on and off this bistability property? What are the relevant parameters?

The triplet learning rule developed in papers VI and VII, does not have any explicit dependence upon the weights. In an approach similar to the one of Gütig et al. (2003, Meffin et al. (2006), it would be interesting to study weight dependence of the amplitude coefficients ($A_2^-, A_3^-, A_2^+, A_3^+$) with a Fokker-Planck equation (van Rossum et al. 2000; Rubin et al. 2001). What are the conditions that lead to a unimodal/bimodal distribution of the weights? Is there a specific choice of the weight dependence that could lead to a bistability between unimodal and bimodal weight distribution in a way which similar to paper V? All this study can be done for the All-to-All interaction scheme; does it make any difference on

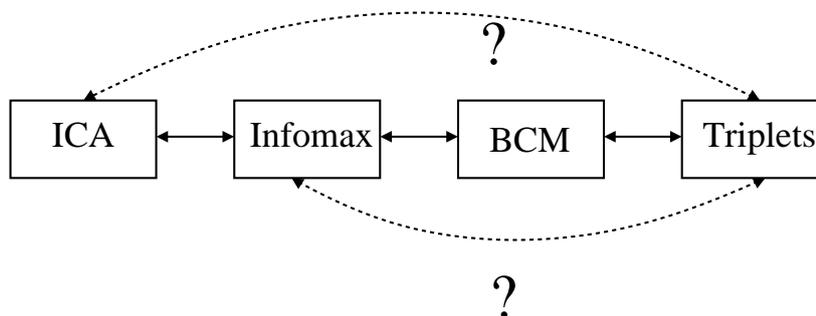


Figure 3.1: Schematic view of a possible line of research. Solid arrows denote links between theories that can be found in the literature. Dashed arrows indicate links that remain to be investigated.

the weight distribution if we assume instead a Nearest-Past-Spike interaction scheme?

A second line of research for those triplet models would be to show the computational properties of the triplet learning rule (see Fig. 3.1). Is there any relation with independent component analysis (ICA). What is the relation between the triplet learning rule and the learning rule that maximizes the mutual information? We have shown in papers VI and VII, that it is possible to do a mapping from the triplet model to the BCM model. In paper IV, we have shown a link between the BCM learning rule and the weight dynamics which results from the mutual information maximization, so there is probably a way to show the link between infomax and the triplets. It is also known that the ICA algorithm can be obtained by maximization of the mutual information (Nadal and Parga 1994; Bell and Sejnowski 1995), so there might be a link between the triplets and ICA.

The research for this plausible link between the triplet learning rule and ICA is also motivated by the fact that pair-based learning rules lead to principal component analysis (PCA) (Oja 1982). So perhaps, the transition from “pair” to “triplet” learning rule could cause a transition from PCA to ICA, which can be seen as a non-linear generalization of the PCA (Oja et al. 1997).

Finally, it could be interesting to further test the triplet learning rules on different preparations. Up to now, we used the data from Sjöström et al. (2001) performed in L5 pyramidal cells of the visual cortex and data from Wang et al. (2005) obtained from measurements in hippocampal cultures. A lot of recordings have also been done in L2/3 pyramidal neurons of

visual cortex (Froemke and Dan 2002; Froemke et al. 2006). The triplet experiments in those last papers give results which are different from the ones in the hippocampal culture. Can our triplet model still reproduce the experimental results? If we relax the constraint that all amplitude parameters ($A_2^-, A_3^-, A_2^+, A_3^+$) should be bigger than zero, the triplet learning rule might be consistent with those experimental data and therefore give a further indication that triplets of spike are relevant in the modeling of STDP.

Lausanne, July 3, 2006

A.1 All-to-All Interaction Scheme

In section 1.2.1, we wrote the instantaneous weight change \dot{w} as

$$\dot{w}(t) = H_0 + X(t)F[X, Y] + Y(t)G[X, Y] \quad (\text{A.1.1})$$

where $F[X, Y]$ and $G[X, Y]$ are unknown functionals of the presynaptic spike train X and the postsynaptic spike train Y . We wrote in Eqs. (1.2.3) and (1.2.4) the first terms of a Volterra expansion of those functionals. As already mentioned in the introduction, this framework corresponds to the All-to-All interaction scheme (c.f. Fig 1.3).

In order to write the general form of a Nearest-Past-Spike interaction scheme, it is necessary to first write in a compact form the Eqs. (1.2.2), (1.2.3) and (1.2.4). Let $x^l = xx \dots x$ denote a sequence of l presynaptic spike trains X and y^l a sequence of l postsynaptic spike trains Y . In this way we can define $F_k^{x^l y^{k-l}}$ as the kernel governing the interaction of the order k with l presynaptic spikes and $k - l$ postsynaptic spikes. For example the kernel $G_3^{xyy} \equiv G_3^{x^1 y^2}$ in Eq. (1.2.4) governs the interactions between 1 pre and 2 postsynaptic spikes. Therefore, the weight change of Eq. (1.2.2) can be rewritten as

$$\begin{aligned}
\dot{w}(t) = & X(t) \sum_{k=1}^{\infty} \sum_{l=1}^k \int_{\Omega^{k-1}} F_k^{x^l y^{k-l}}(\mathbf{s}^{k-1}) \prod_{i=1}^{l-1} X(t-s_i) \prod_{j=l}^{k-1} Y(t-s_j) d\mathbf{s}^{k-1} \\
& + Y(t) \sum_{k=1}^{\infty} \sum_{l=0}^{k-1} \int_{\Omega^{k-1}} G_k^{x^l y^{k-l}}(\mathbf{s}^{k-1}) \prod_{i=1}^l X(t-s_i) \prod_{j=l+1}^{k-1} Y(t-s_j) d\mathbf{s}^{k-1} \quad (\text{A.1.2})
\end{aligned}$$

where $\Omega =]0, \infty]$ is the domain of integration and $\mathbf{s}^{k-1} = (s_1, \dots, s_{k-1})$ is a $k-1$ dimensional vector. In Eq. (A.1.2), the index k denote the order of interactions. For example, $k=3$ corresponds to spike triplets.

A.2 Nearest-Past-Spike Interaction Scheme

Instead of considering all possible interactions, i.e. all the pair interactions for the second order, all the triplet interactions for the third order, \dots it is possible to consider only the last spike before the actual one or the second last spike, \dots . This is what we will call the *Nearest-Past-Spike* expansion. Let $f_x(t)$ denote the last presynaptic spike before t . Similarly, let $f_y(t)$ denote the last postsynaptic spike before t . Let us further denote $f_{x^l}(t) = f_x(f_x(\dots f_x(t)\dots))$ the l^{th} presynaptic spike before t and $f_{y^l}(t)$ the l^{th} postsynaptic spike before t . Finally, let $\mathbf{f}_x^l = (f_x(t), \dots, f_{x^l}(t))$ be the vector containing the timings of the l previous spikes. With this notation, we can replace the domain of integration Ω^{k-1} on the first line of Eq. (A.1.2) by

$$\begin{aligned}
\Omega^{l-1, k-l} = &]0, f_x(t)] \times]f_x(t), f_{x^2}(t)] \times \dots \times]f_{x^{l-2}}(t), f_{x^{l-1}}(t)] \\
& \times]0, f_y(t)] \times \dots \times]f_{y^{k-l-1}}, f_{y^{k-l}}] \quad (\text{A.2.1})
\end{aligned}$$

and on the second line of Eq. (A.1.2), Ω^{k-1} is replaced by $\Omega^{l, k-l-1}$. The instantaneous weight change yields

$$\begin{aligned}
\dot{w}(t) = & X(t) \sum_{k=1}^{\infty} \sum_{l=1}^k F_k^{x^l y^{l-k}}(\mathbf{t}^{l-1} - \mathbf{f}_x^{l-1}, \mathbf{t}^{k-l} - \mathbf{f}_y^{k-l}) \\
& + Y(t) \sum_{k=1}^{\infty} \sum_{l=0}^{k-1} G_k^{x^l y^{l-k}}(\mathbf{t}^l - \mathbf{f}_x^l, \mathbf{t}^{k-l-1} - \mathbf{f}_y^{k-l-1}) \quad (\text{A.2.2})
\end{aligned}$$

where $\mathbf{t}^l = (t, \dots, t)$ is a time vector with l identical components. As an illustration, the term $G_4^{xyyy}(t - f_x(t), t - f_y(t), t - f_y(f_y(t)))$ corresponds to the quadruplet term which depends on the last presynaptic spike, the last postsynaptic spike and the second last postsynaptic spike.

B.1 Notes on the Tempotron

Recently Gütig and Sompolinsky (2006) proposed a learning rule which is derived in a similar context as the one in paper II. This learning rule, termed as the *tempotron* is intended to be the generalization of the perceptron learning rule (Minsky and Papert 1969). Indeed, this rule is a spike-based and not a rate-based learning rule. The idea is the following. The goal of the tempotron is to classify input patterns into two classes: the class C^\oplus and the class C^\ominus . More precisely, after learning the neuron should elicit at least one spike if a pattern $\mathbf{x} \in C^\oplus$ is presented and no spike if a pattern $\mathbf{x} \in C^\ominus$ is presented.

If the neuron makes a classification error, i.e. generating spikes when $\mathbf{x} \in C^\ominus$ is presented and generating no spike in the presence of $\mathbf{x} \in C^\oplus$, the weight w_j of synapse j is changed according to the following rule:

$$\Delta w_j = \alpha^\pm \sum_{t_j < t_{\max}} \epsilon(t_{\max} - t_j) \quad (\text{B.1.1})$$

where $\epsilon(s)$ is the EPSP kernel, t_j are the presynaptic spike timings of synapse j and $\alpha^+ = \alpha$ is a small positive constant if an input pattern $\mathbf{x} \in C^\oplus$ is presented and no postsynaptic spike elicited. Conversely $\alpha^- = -\alpha$ is a small negative constant if an input pattern $\mathbf{x} \in C^\ominus$ is presented and at least one postsynaptic spike is generated. If the neuron does not make any error, the weight w_j does not change.

A more biological plausible version of the tempotron learning rule (c.f. Eq. (B.1.1)) has been proposed by Gütig and Sompolinsky (2006) in order to avoid the calculation of t_{\max} :

$$\Delta w_j = \alpha^\pm \sum_{t_j} \int_0^T f(u(s)) \epsilon(s - t_j) ds \quad (\text{B.1.2})$$

where $f(u)$ is an increasing function of the membrane potential. Typically $f(u) = u$ or $f(u) = \exp(u)$.

Here we will show that this learning rule is a very specific case of the reinforcement learning rule proposed in paper II (c.f. Eq. (31)) as well as the one proposed by (Xie and Seung 2004). Let $\langle R \rangle$ denote the expected reward:

$$\langle R \rangle = \sum_{\mathbf{x}, y} R(\mathbf{x}, y) P(y|\mathbf{x}) P(\mathbf{x}) \quad (\text{B.1.3})$$

where \mathbf{x} and y are respectively the pre- and postsynaptic spike trains on the interval $[0, T]$. The gradient ascent learning rule Eq. (B.1.3) yields

$$\Delta w_j = \alpha \sum_{\mathbf{x}, y} R(\mathbf{x}, y) \frac{\partial \log P(y|\mathbf{x})}{\partial w_j} P(y|\mathbf{x}) P(\mathbf{x}) \quad (\text{B.1.4})$$

where α is the learning rate. For convenience, let us recall here that for a Poisson neuron $\log P(y|\mathbf{x})$ can be written as (see Eq. (2.6) of paper II):

$$\log P(y|\mathbf{x}) = \int_0^T \log(\rho(s)) y(s) - \rho(s) ds \quad (\text{B.1.5})$$

where $\rho(s) = g(u(s))$ corresponds to the instantaneous firing rate and is a function of the membrane potential $u(s)$. The mapping between Eq. (B.1.1) and Eq. (B.1.4) can be seen if we consider the following reward scheme

$$R(\mathbf{x}, y) = \begin{cases} 0 & \text{if } \mathbf{x} \in C^\oplus \text{ and } N_y \geq 1 \\ 0 & \text{if } \mathbf{x} \in C^\ominus \text{ and } N_y = 0 \\ -1 & \text{if } \mathbf{x} \in C^\oplus \text{ and } N_y = 0 \\ -1 & \text{if } \mathbf{x} \in C^\ominus \text{ and } N_y \geq 1 \end{cases} \quad (\text{B.1.6})$$

where $N_y = \int_0^T y(s) ds$ is the number of postsynaptic spikes during the interval $[0, T]$. In other words the reward is 0 if the tempotron classifies well and -1 in case of misclassification.

Replacing Eq. (B.1.6) into Eq. (B.1.4), we get

$$\begin{aligned} \Delta w_j = & - \alpha \sum_{\mathbf{x} \in C^\oplus} \frac{\partial \log P(0|\mathbf{x})}{\partial w_j} P(0|\mathbf{x}) P(\mathbf{x}) \\ & - \alpha \sum_{\mathbf{x} \in C^\ominus} \frac{\partial \log(1 - P(0|\mathbf{x}))}{\partial w_j} (1 - P(0|\mathbf{x})) P(\mathbf{x}) \end{aligned} \quad (\text{B.1.7})$$

By replacing in the above equation the exact averaging by the empirical averaging, we get

$$\Delta w_j = \alpha^\pm \sum_{t_j} \int_0^T \rho'(s) \epsilon(s - t_j) ds \quad (\text{B.1.8})$$

where $\rho'(s) = g'(u(s)) = \frac{dg}{du}|_{u=u(s)}$. It is now clear that this last equation is identical to the “biological” version of the tempotron (c.f. Eq. (B.1.2)) if $f(u) = g'(u)$.

In summary, the biological version of the tempotron learning rule is a specific case of the reinforcement learning rule proposed in paper II. Indeed, among the enormous number of possible spike trains that could lead to different rewards, the tempotron considers two classes of output spike trains y : the first class which only contains the case with no spike and the second class containing all spike trains with at least one spike. From this point of view, the tempotron considers a spike-code in the input and a kind of rate code in the output.

We should note that the tempotron has been formulated in a deterministic framework whereas spike generation in papers I and II is stochastic. It is, however, possible to consider the deterministic limit of the function $g(u)$ and therefore get a purely deterministic framework, identical to the one of the tempotron; e.g. if $g(u) = \exp((u - \theta)/\Delta u)$, the deterministic limit is obtained if $\Delta u \rightarrow 0$.

It is also interesting to note that the kernel $\epsilon(s)$ assumed in the learning rule of Gütig and Sompolinsky (2006) is identical to the one that describes the EPSP time course. This is consistent with papers I-III and V which state that the potentiation part of the optimal learning window is given by $\epsilon(s)$.

B.2 Triplets and the SMT Learning Rule

In the introduction of this thesis (section 1.2.2), we have described the Senn-Markram Tsodyks (SMT) model of synaptic plasticity. The aim of this appendix is to show how we can simplify it and get a learning rule which is similar (but not identical) to the triplet

learning rules developed in paper VI and VII. Let us recall here for convenience the dynamics of the SMT model:

$$\dot{N}^u(t) = -\frac{N^u(t)}{\tau_{N^u}} + r^u N^r(t) X(t) \quad (\text{B.2.1})$$

$$\dot{N}^d(t) = -\frac{N^d(t)}{\tau_{N^d(t)}} + r^d N^r(t) Y(t) \quad (\text{B.2.2})$$

$$\dot{S}^u(t) = -\frac{S^u(t)}{\tau_{S^u}} + r^S N^u(t) (1 - S^u(t - \epsilon)) Y(t) \quad (\text{B.2.3})$$

$$\dot{S}^d(t) = -\frac{S^d(t)}{\tau_{S^d}} + r^S N^d(t) (1 - S^d(t - \epsilon)) X(t) \quad (\text{B.2.4})$$

$$\begin{aligned} \dot{w}^\infty(t) &= r^P (1 - w^\infty(t - \epsilon)) [S^u(t + \epsilon) - \theta^u]^+ Y(t) \\ &\quad - r^D w^\infty(t - \epsilon) [S^d(t + \epsilon) - \theta^d]^+ X(t) \end{aligned} \quad (\text{B.2.5})$$

$$\dot{w}(t) = \frac{w^\infty(t) - w(t)}{\tau_M} \quad (\text{B.2.6})$$

In order to see the relation between the triplet model and the SMT model, we have to do some simplifying assumptions:

1. (H1). Almost all NMDA receptors are in the recovery state, i.e. $N^r \simeq 1$
2. (H2). The number of secondary messengers in the *up* state and *down* state is small, i.e. $S^u \simeq 0$, $S^d \simeq 0$.
3. (H3). The threshold of the secondary messengers are zero, i.e. $\theta^u = \theta^d = 0$.
4. (H4). The weight changes are instantaneous, i.e. $\tau_M \rightarrow 0$.

From those hypothesis, we get the simplified SMT model named as sSMT:

$$\dot{N}^u(t) = -\frac{N^u(t)}{\tau_{N^u}} + r^u X(t) \quad (\text{B.2.7})$$

$$\dot{N}^d(t) = -\frac{N^d(t)}{\tau_{N^d(t)}} + r^d Y(t) \quad (\text{B.2.8})$$

$$\dot{S}^u(t) = -\frac{S^u(t)}{\tau_{S^u(t)}} + r^S N^u(t) Y(t) \quad (\text{B.2.9})$$

$$\dot{S}^d(t) = -\frac{S^d(t)}{\tau_{S^d}} + r^S N^d(t) X(t) \quad (\text{B.2.10})$$

$$\dot{w}^\infty(t) = r^P (1 - w(t - \epsilon)) S^u(t + \epsilon) Y(t) - r^D w(t - \epsilon) S^d(t + \epsilon) X(t) \quad (\text{B.2.11})$$

Note that the update of the weight occurs *after* the update of S^u and S^d because of the dependence upon $t + \epsilon$ where $\epsilon > 0$. Because of this precise sequence of updates, the first term on the r.h.s. of Eq. (B.2.11) captures pre-post pairs of spikes as well as pre-post-post triplets of spikes (c.f. Fig. B.2). If the order is reversed, only pre-post-post triplet of spikes are considered. Similarly, the second term on the r.h.s. of Eq. (B.2.11) corresponds to a post-pre sequence as well as a post-pre-pre sequence (see Fig. B.1). The main difference with our triplet model is that in the SMT model, the pre-post-post sequence leading to potentiation is imposed whereas in papers VI and VII, the triplet potentiation term corresponds to a pre-post-post sequence *as well as* a post-pre-post sequence. Indeed, in the SMT model, the secondary messengers depend on the number of NMDA receptors in the up and down state whereas in papers VI and VII, the pre- and postsynaptic detectors are independent; c.f. Eqs. (1) and (2) in paper VII.

Apart from the order distinction we made in the previous paragraph, Eqs. (B.2.7)-(B.2.11) correspond to the All-to-All triplet terms of paper VI and VII. Note that if the assumption H2 is not valid, the level of secondary messengers saturates and hence their dynamics correspond to a Nearest-Past-Spike interaction scheme. More precisely, the pre-post part of the pre-post-post term would be of the type All-to-All and the post-post part of the pre-post-post term would be of the type Nearest-Past-Spike (see Fig. B.1). Similar arguments apply to the triplet depression term. This mixture of All-to-All and Nearest-Past-Spike schemes is similar to the type of interactions assumed in the model of Froemke and Dan (2002). In this paper, an efficacy value is attributed to each a presynaptic (postsynaptic) spike as a function of the distance to the last presynaptic (postsynaptic) spike; this correspond to Nearest-Past-Spike interactions. Then, the pre- and postsynaptic spikes interact in an All-to-All interaction scheme.

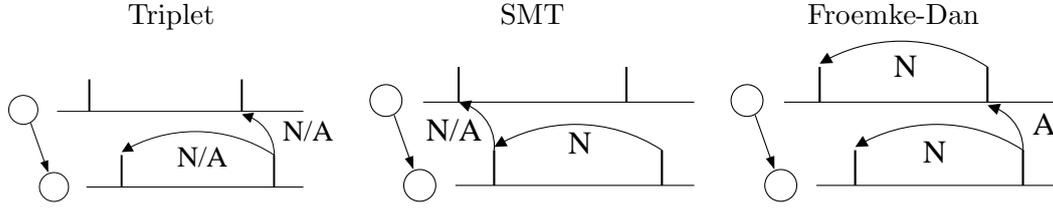


Figure B.1: Comparison between the (*left*) triplet model of paper VI and VII, (*middle*) the simplified SMT (sSMT) model, and (*right*) the Froemke-Dan model (2002). N and A denote respectively the Nearest-Past-Spike and the All-to-All interaction scheme. C.f. Fig 1.3. N/A denotes that both types of interactions are possible. In the sSMT model, the pre-post interaction is All-to-All if the hypothesis H2 (low level of secondary messengers) is valid and corresponds to Nearest-Past-Spike if H2 is not valid. Only the triplet potentiation term is represented.

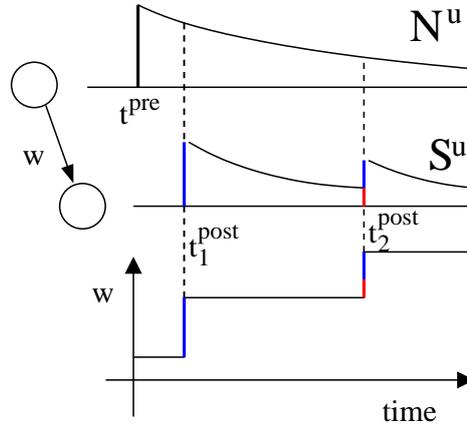


Figure B.2: Illustration of the pair and triplet effects in the SMT model. The presynaptic spike (at time t_1^{pre}) causes an increase of N^u . The two postsynaptic spikes (at time t_1^{post} and t_2^{post}) cause two subsequent increases of S^u . The vertical blue lines correspond to pair effects while the vertical red lines to triplet effects. The weight w increases at time t_1^{post} because the value of S^u is read just after the postsynaptic spike, i.e. at time $t_1^{post} + \epsilon$: this is a pure pair effect. At time t_2^{post} , there is a pair effect as well as a triplet effect.

Bibliography

- A. Artola and S. Bröcher and W. Singer** (1990). Different voltage dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347, 69–72.
- F. Attneave** (1954). Some informational aspects of visual perception. *Psychological Review* 61, 183–193.
- H. B. Barlow** (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith (Ed.), *Sensory Communication*, pp. 217–234. MIT Press.
- G. Barrionuevo and F. Schottler and G. Lynch** (1980). The effects of repetitive low frequency stimulation on control and “potentiated” synaptic responses in the hippocampus. *Life Sciences* 27(24), 2385–2391.
- A.J. Bell and T.J. Sejnowski** (1995). An information maximization approach to blind separation and blind deconvolution. *Neural Computation* 7, 1129–1159.
- C.C. Bell and V. Han and Y. Sugawara and K. Grant** (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387, 278–281.
- G.Q. Bi and M.M. Poo** (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- G-Q. Bi** (2002). Spatiotemporal specificity of synaptic plasticity: cellular rules and mechanisms. *Biological Cybernetics* 87, 319–332.
- E.L. Bienenstock and L.N. Cooper and P.W. Munro** (1982). Theory of the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience* 2, 32–48. reprinted in Anderson and Rosenfeld, 1990.
- T.V.P. Bliss and A.R. Gardner-Medwin** (1973). Long-lasting potentiation of synaptic

transmission in the dentate area of unanaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 357–374.

T.V.P. Bliss and T. Lomo (1973). Long-lasting potentiation of synaptic transmission in the dentate area of anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232, 551–356.

T.V.P. Bliss and M.A. Lynch (1988). Long-term potentiation of synaptic transmission in the hippocampus: properties and mechanisms. In P. Landfield and S. Deadwyler (Eds.), *Long-term potentiation in the hippocampus: from biophysics to behavior*, New York, pp. 3–72. Alan R. Liss.

T. V. P. Bliss and G. L. Collingridge and R.G.M. Morris (2003). Long-term potentiation: enhancing neuroscience for 30 years. *Phil. Trans. R. Soc. Lond. B* 358, 607–611.

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

A. N. Burkitt and M. H. Meffin and D.B. Grayden (2004). Spike-timing-dependent plasticity: The relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Computation* 16, 885–940.

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

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

Y. Frégnac and D. E. Shulz and S. Thorpe and E. Bienenstock (1992). Cellular analogs of visual cortical epigenesis. I plasticity of orientation selectivity. *Journal of Neuroscience* 12(4), 1280–1300.

R.C. Froemke and Y. Dan (2002). Spike-timing dependent plasticity induced by natural spike trains. *Nature* 416, 433–438.

Froemke, RC and Poo, MM and Dan, Y. (2005). Spike-timing-dependent synaptic

- plasticity depends on dendritic location. *Nature* 434(7030), 221–5.
- R.C. Froemke and I.A. Tsay and M. Raad and J.D. Long and Y. Dan** (2006). Contribution of individual spikes in burst-induced long-term synaptic modification. *Journal of Neurophysiology* 95, 1620–1629.
- Y.-X. Fu and K. Djupsund and H. Gao and B. Hayden and K. Shen and Y. Dan** (2002). Temporal specificity in the cortical plasticity of visual space representation. *Science* 296, 1999–2003.
- W. Gerstner and R. Kempter and J. Leo van Hemmen and H. Wagner** (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383, 76–78.
- W. Gerstner and W. K. Kistler** (2002). *Spiking Neuron Models*. Cambridge University Press.
- S.J. Gould and R.C. Lewontin** (1979). Spandrels of san-marco and the panglossian paradigm - a critique of the adaptationist program. *Proceedings of the Royal Society of London Series B-Biological Sciences* 295, 581–598.
- R. Gütiğ and R. Aharonov and S. Rotter and H. Sompolinsky** (2003). Learning input correlations through non-linear temporally asymmetry hebbian plasticity. *J. Neuroscience* 23, 3697–3714.
- Gütiğ, R. and Sompolinsky, H.** (2006). The tempotron: a neuron that learns spike timing-based decisions. *Nature Neuroscience* 9, 420–428.
- D. O. Hebb** (1949). *The Organization of Behavior*. New York: Wiley.
- A. L. Hodgkin and A. F. Huxley** (1952). A quantitative description of ion currents and its applications to conduction and excitation in nerve membranes. *J. Physiol. (London)* 117, 500–544.
- D. H. Hubel and T. N. Wiesel** (1959). Receptive fields of single neurons in the cat's striate cortex. *J. Physiol. (London)* 148, 574–591.
- D. H. Hubel and T. N. Wiesel** (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (London)* 160, 106–154.
- N. Intrator and L. N. Cooper** (1992). Objective function formulation of the BCM theory of visual cortical plasticity. *Neural Networks* 5, 3–17.

- E.M. Izhikevich** (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks* 14, 1569–1572.
- William James** (1890). *Psychology (Briefer Course)*, ch. 16. New York: Holt.
- E.R. Kandel** (2001). The molecular biology of memory storage: a dialogue between genes and synapses. *Science* 294, 1030–1038.
- U.R. Karmarkar and D.V. Buonomano** (2002). A model of spike-timing dependent plasticity: one or two coincidence detectors. *J. Neurophysiology* 88, 507–513.
- Kasai, H. and Matsuzaki, M. and Noguchi, J. and Yasumatsu, N. and Nakahara, H.** (2003). Structure-stability-function relationships of dendritic spines. *Trends Neurosci* 26(7), 360–8.
- R. Kempter and W. Gerstner and J. L. van Hemmen** (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514.
- W. M. Kistler and J. Leo van Hemmen** (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic potentials. *Neural Comput.* 12, 385–405.
- K.S. Lashley** (1924). Studies of cerebral function in learning. vi. the theory that synaptic resistance is reduced by the passage of the nerve impulse. *Psychol. Rev.* 31, 369–375.
- S. Laughlin** (1981). A simple coding procedure enhances a neuron’s information capacity. *Z. Naturforsch. C* 36, 910–912.
- W. B. Levy and D. Stewart** (1983). Temporal contiguity requirements for long-term associative potentiation/depression in hippocampus. *Neurosci*, 8, 791–797.
- Ralph Linsker** (1988). Self-organization in a perceptual network. *Computer* 21, 105–117.
- J. Lisman** (1989). A mechanism for hebb and anti-hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA* 86, 9574–9578.
- T. Lomo** (1966). Frequency potentiation of excitatory synaptic activity in the dentate area of the hippocampal formation. *Acta. Physiol. Scand.* 68, 128.
- E. Mach** (1886). *The Analysis of Sensations*. Chicago and London: Open Court. Trans. by C. M. Willimans in 1914 from 1st German ed.

- R. C. Malenka and J.A. Kauer and R.S. Zucker and R. A. Nicoll** (1988). Post-synaptic calcium is sufficient for potentiation of hippocampal synaptic transmission. *Science* 242, 81–84.
- H. Markram and J. Lübke and M. Frotscher and B. Sakmann** (1997). Regulation of synaptic efficacy by coincidence of postsynaptic AP and EPSP. *Science* 275, 213–215.
- Mayer, ML and MacDermott, AB and Westbrook, GL and Smith, SJ and Barker, JL** (1987). Agonist-and voltage-gated calcium entry in cultured mouse spinal cord neurons under voltage clamp measured using arsenazo III. *J. Neurosci* 7, 3230–3244.
- Meffin, H. and Besson, J. and Burkitt, AN and Grayden, DB** (2006). Learning the structure of correlated synaptic subgroups using stable and competitive spike-timing-dependent plasticity. *Phys Rev E* 73, 041911.
- C.D. Meliza and Y. Dan** (2006). Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron* 49, 183–189.
- M. L. Minsky and S. A. Papert** (1969). *Perceptrons*. Cambridge Mass.: MIT Press.
- J.P. Nadal and N. Brunel and N. Parga** (1998). Nonlinear feedforward networks with stochastic outputs: infomax implies redundancy reduction. *Network: Computation in Neural Systems* 9, 207–217.
- Nadal, J.P.A. and Parga, N.A.** (1994). Nonlinear neurons in the low-noise limit: a factorial code maximizes information transfer. *Network: Computation in Neural Systems* 5(4), 565–581.
- J. P. Nadal and N. Parga** (1997). Redundancy reduction and independent component analysis: Conditions on cumulants and adaptive approaches. *Neural Computation* 9, 1421–1456.
- E. Oja** (1982). A simplified neuron model as a principal component analyzer. *J. Mathematical Biology* 15, 267–273.
- Erkki Oja and Juha Karhunen and Aapo Hyvärinen** (1997). From neural principal components to neural independent components. In *ICANN '97: Proceedings of the 7th International Conference on Artificial Neural Networks*, London, UK, pp. 519–528.

Springer-Verlag.

- K. Pearson** (1892). *The Grammar of Science*. London: Scott.
- P.D. Roberts** (1999). Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. *J. Computational Neuroscience* 7, 235–246.
- Rubin, J.E. and Gerkin, R.C. and Bi, G.Q. and Chow, C.C.** (2005). Calcium Time Course as a Signal for Spike-Timing-Dependent Plasticity. *Journal of Neurophysiology* 93, 2600–2613.
- J. Rubin and D. D. Lee and H. Sompolinsky** (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Physical Review Letters* 86, 364–367.
- W. Senn and H. Markram and M. Tsodyks** (2001). An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Computation* 13, 35–67.
- C.E Shannon** (1948). A mathematical theory of communication. *The Bell System Technical Journal* 27, 379–423, 623–656.
- H. Z. Shouval and M. F. Bear and L. N. Cooper** (2002). A unified model of NMDA receptor dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA* 99, 10831–10836.
- P.J. Sjöström and G.G. Turrigiano and S.B. Nelson** (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32, 1149–1164.
- S. Song and K.D. Miller and L.F. Abbott** (2000). Competitive Hebbian learning through spike-time-dependent synaptic plasticity. *Nature Neuroscience* 3, 919–926.
- G.S. Stent** (1973). A physiological mechanism for Hebb’s postulate of learning. *Proc. Natl. Acad. Sci. USA* 70, 997–1001.
- Suri, R.E.A.** (2004). A computational framework for cortical learning. *Biological Cybernetics* 90(6), 400–409.
- Tzounopoulos, T. and Kim, Y. and Oertel, D. and Trussell, L.O.** (2004). Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nature Neuroscience* 7(7), 719–725.

- M. C. W. van Rossum and G. Q. Bi and G. G. Turrigiano** (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience* 20, 8812–8821.
- Volterra, A. and Meldolesi, J.** (2005). Astrocytes, from Brain Glue to Communication Elements: the Revolution Continues. *Nature Reviews Neuroscience* 6, 626–640.
- V. Volterra** (1930). *Theory of Functionals and of Integral and Integro-Differential Equations*. New York: Dover.
- H. X. Wang and R. C. Gerkin and D. W. Nauen and G. Q. Bi** (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nature Neuroscience* 8, 187–193.
- X. Xie and S. Seung** (2004). Learning in neural networks by reinforcement of irregular spiking. *Phys. Rev. E* 69, 041909.
- H. Yao and Y. Dan** (2001). Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron* 32, 315–323.
- H. Yao and Y. Dan** (2004). Intracortical mechanism of stimulus-timing-dependent plasticity in visual cortical orientation tuning. *Proc. Natl. Acad. Sci. USA* 101, 5081–5086.
- L.I. Zhang and H.W. Tao and C.E. Holt and W.A. Harris and M.-M. Poo** (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395, 37–44.

EPFL LCN

AAB 1 21 (Bâtiment AAB)

Station 15

CH-1015 Lausanne

Tel.: +41 21 693 5593

Fax: +41 21 693 9600

e-mail: jean-pascal.pfister@a3.epfl.ch

webpage: <http://diwww.epfl.ch/~jpfister>

28, married, 1 child, Swiss

Professional experience

- August 2002 - June 2006: **Postgraduate research** at the Laboratory of Computational Neuroscience (Prof. Wulfram Gerstner), Brain Mind Institute and School of Computer and Communication Science, EPFL, Switzerland. Research on Spike-Timing Dependent Plasticity.
- March 2002 - June 2006: **Teaching** at the Collège secondaire intercommunal de la Planta, Switzerland.
- February 2003 - Sept. 2003: **Supply Teaching** in mathematics at the Ecole d'ingénieur du canton de Vaud, Switzerland.
- Oct. 1999 - July 2000: **Research in the industry** for Moving Magnet Technologies SA, Besançon, France.
- Oct. 2000 - June 2000: **Undergraduate research** at the Laboratory of Artificial Intelligence, EPFL, Switzerland.

Education

- July 2006: **PhD in Computational Neuroscience**, EPFL, Switzerland.
- Aug. 2003: **Advanced course: Methods and Models in Neurophysics**, Les Houches, France.
- March 2002: **M.Sc in Physics**, EPFL, Switzerland.
- 1999-2000: Third year of physics at **The University of Nottingham**, United Kingdom.
- June 1996: **Swiss Federal Maturity in Science** (Baccalauréat) with Greek option, Bienne, Switzerland.

Languages

- French: mother tongue
- English: fluent
- German: good knowledge

Computer science skills

- OS: Windows 9x/XP, UNIX (Solaris), Linux.
- Languages: \LaTeX , Pascal, Perl, C, HTML.
- Others: Matlab, Maple, Mathematica.

Personal interests

- Sports: Alpinism, Rock climbing, Triathlon (Polyathlon 2005 and 2006, Ironleman 2005), Scuba Diving.
- Others: Leading youth groups, Camp Leader, President of the Biblical Group (GBU) of the EPFL in 1998-1999.

Publications

1. **Pfister JP and Barber D and Gerstner W.** Optimal Hebbian Learning: A Probabilistic Point of View. *Artificial Neural Networks and Neural Information Processing - ICANN/ICONIP 2003*, edited by O. Kaynak, E. Alpaydin, E. Oja and L. Xu. Berlin: Springer-Verlag, 92-98, 2003.
2. **Toyoizumi T and Pfister JP and Aihara K and Gerstner W.** Generalized Bienenstock-Cooper-Munro rule for spiking neurons that maximizes information transmission *Proceedings of the National Academy of Science USA*. 102, 5239-5244, 2005.
3. **Toyoizumi T and Pfister JP and Aihara K and Gerstner W.** Spike-Timing Dependent Plasticity and Mutual Information Maximization for a Spiking Neuron Model. *Advances in Neural Information Processing Systems 17*, edited by L.K. Saul and Y. Weiss and L. Bottou, MIT Press, Cambridge MA, 1409-1416, 2005.
4. **Pfister JP and Gerstner W.** Beyond Pair-Based STDP: a Phenomenological Rule for Spike Triplet and Frequency Effects. *Advances in Neural Information Processing Systems 18*, edited by Y. Weiss and B. Schölkopf and J. Platt, MIT Press, Cambridge MA, 1083-1090, 2006.
5. **Pfister JP and Toyoizumi T and Aihara K and Gerstner W.** Optimal Spike-Timing Dependent Plasticity for Precise Action Potential Firing in Supervised Learning. *Neural Computation*, 18, 1309-1339, 2006.
6. **Toyoizumi T and Pfister JP and Aihara K and Gerstner W.** Optimality Model of Unsupervised Spike-Timing Dependent Plasticity: synaptic memory and weight distribution. *Accepted to Neural Computation*.
7. **Pfister JP and Gerstner W.** Why Triplets of Spike are Necessary in Models of Spike-Timing-Dependent Plasticity. *Submitted to The Journal of Neuroscience*.