

Optimal Hebbian Learning: a Probabilistic Point of View

Jean-Pascal Pfister¹, David Barber², and 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

Abstract. Up to now, 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. The idea is to create a model in order 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 results are consistent with the actual knowledge of back-propagating action potentials

1 Introduction

Since synaptical changes are assumed to underly memory and learning processes, a crucial question is to determine what causes those changes and what is the law 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.

Until now, the existing models are either phenomenological [14] or pseudo-mechanistic [1]. Our model is derived from a probabilistic point of view in the sense that the learning rule should optimise the likelihood of observing a post-synaptic 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 [12]:

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 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).




Fig. 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 from [4].

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 as escape rate [10]. The membrane potential model we take is the Spike Response Model (SRM). The simplest SRM is called SRM_0 [8] 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 $\eta(s)$ is a kernel describing the back-propagating action potential at the location of the synapse, $\epsilon(s)$ 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 $\tilde{\mathbf{t}}$ is

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

Now we can come back to the continuous case. By taking the limit $\Delta t \rightarrow 0$, we have $\bar{\rho}(t) \rightarrow \rho(t)$, $\bar{P}(\mathbf{t} \in \Omega(\tilde{\mathbf{t}}) | \bar{\rho}(s)) \rightarrow P(\mathbf{t} = \mathbf{t}' | \rho(s))$, $\sum_n \log(1 - \bar{\rho}(t_n) \Delta t) \rightarrow -\int_0^T \rho(t) dt$ 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 weights 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

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

and κ is the learning rate. Since $g(u(t)) = \exp(\beta(u(t) - \theta))$ is a reasonable choice [10], 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 $\beta = 1$:

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

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

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

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

where Θ is the usual Heaviside step function with $\Theta(s) = 1$ for $s > 0$ and $\Theta(s) = 0$ else.

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 if back-propagating action potential amplitude η_0 is positive (say $\eta_0 = 1$) and not if it is negative. This is consistent with the experiments of Stuart and Sakmann [22].




Fig. 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 back-propagating action potential is given by $\eta = -1$ for **(a)** and $\eta = 1$ for **(b)**. Note the different vertical scales.

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)).

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

Let us note that the 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 fix point for w if $t_i < t_j$.

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 if the back-propagating action potential leads to a long-lasting depolarisation of the dendrite after a somatic spike. The duration of the back-propagating spike 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 exact location of the synapse on the dendrite, in the same way as back-propagating action potential may change along the way.

It is of course possible to make the model more complex by not using the SRM_0 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. 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.
7. R. Froenke and Y. Dan: Spike-timing dependent plasticity induced by natural spike trains. *Nature*, 2002, Vol. 416, 433-438.
8. W. Gerstner: Time Structure of the Activity in Neural Network Models. *Phys. Rev. E*, 1995, 51, Vol. 1, 738-758.
9. 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.
10. W. Gerstner and W.M. Kistler: *Spiking Neuron Models*. Cambridge University Press, 2002.
11. R. Gütiğ and R. Aharonov and S. Rotter and H. Sompolinsky: Learning input correlations through non-linear temporally asymmetry Hebbian plasticity. To appear.
12. D.O. Hebb: *The Organization of Behavior*. Wiley, 1949, New York.
13. R. Kempter and W. Gerstner and J. L. van Hemmen: Hebbian learning and spiking neurons. *Phys. Rev. E*, 1999, Vol. 59, 4, 4498-4514
14. 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.
15. 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.
16. 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.
17. P.D. Roberts and C.C. Bell: Spike timing dependent synaptic plasticity in biological systems. *Biol. Cybernetics*, 2002, Vol. 87, 392-403.

18. 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.
19. 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.
20. 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.
21. 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
22. G.J. Stuart and B. Sakmann: Active propagation of somatic action-potentials into neocortical pyramidal cell dendrites. *Nature*, 1994, Vol. 367, 69-72.
23. 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.
24. M.C.W. van Rossum and G.Q. Bi and G.G. Turrigiano: table Hebbian learning from spike timing-dependent plasticity. *J. Neuroscience*, 2000, Vol. 20, 8812-8821.