

SPIKE-TIMING DEPENDENT PLASTICITY AND REGIME TRANSITIONS IN RANDOM RECURRENT NEURAL NETWORKS

Frédéric Henry

Movement Sciences Institute (UMR 6233)
University of the Mediterranean
163, avenue de Luminy, CP910
13288 Marseille Cedex 9, France
email: frederic.henry@etumel.univmed.fr

Emmanuel Daucé

Centrale Marseille/UMR 6233
Technopôle de Château-Gombert
13383 Marseille Cedex 13, France
email: edauce@ec-marseille.fr

ABSTRACT

In this paper, we investigate how Spike-Timing Dependent Plasticity, when applied to a random recurrent network of leaky integrate-and-fire neurons, can affect its dynamical regime. We show that in an autonomous network with self-sustained activity, STDP has a regularization effect and simplifies the dynamics.

We then look at two different ways to present stimuli to the network: potential-based input and current-based input. We show that in the first case STDP can lead to either synchronous or asynchronous periodical activity, depending on the network's internal parameters. However, in the latter case, synchronization can only appear when the input is presented to a fraction of the neurons instead of the whole.

KEY WORDS

Spike-Timing Dependent Plasticity, Synchronization, Random Recurrent Neural Networks.

1 Introduction

The Spike-Timing Dependent Plasticity (STDP) rule comes from biological observations showing that the conductance of a synapse is modified according to the precise timing between the presynaptic and postsynaptic spikes [1, 2, 3]. In most observed cases, on excitatory synapses, the synapse is potentiated if the EPSP (Excitatory Post-Synaptic Potential) is shortly followed by the emission of an action potential at the soma, and depressed in the opposite case (when the AP is followed by the EPSP).

The role of this plasticity rule is still conjectural. Depending on the network structure and activity, the effects on the global response varies. In feedforward networks, STDP is found to reduce the latency of a neuron's response to a given input [4, 5]. In the brain, however, most of the neuron inputs come from recurrent connections. This is why we mostly insist here on the properties of recurrent neural networks, since they can display various dynamical regimes, and as such various qualitative responses. According to [6, 7], the dynamical regimes of random recurrent networks of spiking neurons can be classified in four cat-

egories (synchronous/asynchronous, regular/irregular), depending on the initial network parameters (balance between excitation and inhibition, gain of the cells responses...). The degree of synchrony a neural network can display is supposed to play a prominent role both in the local transmission of sensory information [8] and in the formation of large-scale patterns of activity [9]. The question remains how recurrent neural networks may regulate their degree of synchrony in order to adapt their response to various sensory situations.

We suggest here that STDP may participate to such a regulation and we address the question of STDP-induced regime transitions. It has been shown for instance that STDP could have a decoupling effect in a synchronized neural network, whereas anti-STDP made the activity switch to a synchronized regime [10]. It has also been shown, both in simulation [11] and in biology [8] that STDP could enhance and stabilize the synchronous transmission of an already synchronized input in feed-forward neural networks. Those contradictory results suggest to investigate further the putative effect of STDP on the dynamical regimes, despite few rigorous results exist on that matter. In order to decipher this question, we propose here a simple simulation setup in order to test various conditions in which synchronous and/or periodical regimes can emerge in a network where the initial activity is aperiodical and asynchronous. We then address the question of STDP-related perception and sensory coding in recurrent neural networks.

2 Random recurrent neural networks framework

We simulate random and recurrent neural networks, implemented with discrete leaky integrate-and-fire neurons.

Consider a set of neurons labelled by indexes $i \in \{1, \dots, N\}$. The neuron activity $\{S_i(s)\}_{s < t}$ is defined as a sum of Diracs corresponding to the series of spikes emitted up to the time t (see [12]). Taking into account the absolute refractory period τ_r , the firing threshold θ , the current




Figure 1. Activity of a recurrent neural network with self-sustained activity. $\sigma_J = 3.0$; $\tau_r = 2$ ms. Raster plot with time on x-axis and neuron index on y-axis. A vertical black bar appears when a neuron fires.

activity $S_i(t)$ is defined the following way:

$$S_i(t) = \begin{cases} \delta(0) & \text{if } \max_{s \in [t - \tau_r, t]} (S_i(s)) = 0 \text{ and } V_i(t) \geq \theta \\ 0 & \text{elsewhere} \end{cases} \quad (1)$$

where $\delta(0)$ is a Dirac impulse and $V_i(t)$ is the neuron's membrane potential, defined according to the Leaky Integrate-and-Fire (LIF) differential scheme:

$$\begin{cases} V_i(t) = V_i^{in}(t) + V_i^{ext}(t) \\ \frac{dV_i^{in}}{dt} = -\frac{V_i^{in}(t)}{\tau_m} - S_i(t)V_i^{in}(t) + \sum_{j=1}^N J_{ij}S_j(t - \tau_{ij}) + I_i^{ext}(t) \end{cases} \quad (2)$$

where V_i^{ext} is a superimposed potential. The transmission delay is τ_{ij} , J_{ij} is the synaptic weight from neuron j to neuron i , τ_m is the membrane time constant and I_i^{ext} is an external current.

Since S_i is a sum of Dirac distributions, its presence in the derivative leads to sudden steps in the membrane potential. Thus, when the neuron fires, the potential is reset to zero because of the $-S_i(t)V_i^{in}(t)$ term, whereas presynaptic spikes cause sudden increases or decreases corresponding to the synaptic weight.

We set $\tau_r \in \{1, 2, 3\}$ ms, $\tau_m = 10$ ms and $\theta = 1$. In the simulations, we use a simple first order integration with resolution $\delta t = 1$ ms.

The individual weights and delays are independent and strongly heterogeneous.

The synaptic weights are set according to a Gaussian draw $\mathcal{N}\left(\frac{\mu_J}{N}, \frac{\sigma_J^2}{N}\right)$. We take $\mu_J = 0$ so that the excitatory influences compensate the inhibitory ones (balanced networks configuration). The weights sum standard deviation σ_J represents the internal coupling. The axonal transmission delays are set according to a Poisson draw of expected value $\lambda_\tau = 10$ ms. The simulations take place on rather small neural networks composed of $N = 100$ neurons, but can be extended to larger sizes with the same global parameters.

With the range of parameters we choose, the activity of such networks is irregular (aperiodical and asynchronous), as can be seen on figure 1. First, the *asynchrony*

of the activity directly results from the balance between the excitatory and inhibitory influences [6]. Second, the irregularity of the activity is a well-known feature of recurrent heterogeneous networks [13, 14].

3 Effects of STDP on the activity

Model The STDP rule is classically implemented as an anti-symmetrical function of the pre-synaptic/post-synaptic spike times. It can be applied either to all pairs of spikes emitted by pre-synaptic and post-synaptic neurons (“all-to-all” implementation) [4], or only to the “nearest neighbours” (first pre-synaptic spike after the firing of post-synaptic neuron and first post-synaptic spike after the firing of pre-synaptic neuron) [15, 16].

Here, in order to have an “all-to-all” implementation with a low memory cost (see also [17]), we use a local trace of the most recent spikes ε_i , whose decay corresponds to the membrane time constant : $\tau_m \frac{d\varepsilon_i}{dt} = -\varepsilon_i + S_i(t)$. This trace is consistent with the modelling of the PSP $J_{ij}\varepsilon_j(t - \tau_{ij})$ taking place at the $j \rightarrow i$ synapse at time t , and allows to define a simple additive “all to all” STDP rule:

$$\frac{dJ_{ij}(t)}{dt} = \alpha [S_i(t)\varepsilon_j(t - \tau_{ij}) - \varepsilon_i(t)S_j(t - \tau_{ij})] \quad (3)$$

Here again the presence of sums of Dirac distributions in the derivative leads the weight to change suddenly when there is a pre-synaptic or post-synaptic spike. The left term corresponds to the potentiation effect of the PSP preceding the post-synaptic spike, and the right term corresponds to the depression effect of the post-synaptic spike memory $\varepsilon_i(t)$ when a pre-synaptic spike hits the synapse. The synapse is thus potentiated when the PSP arrives shortly before the neuron spike (ε_j is high when the neuron fires at t) and depressed in the opposite case. The rule is globally balanced as the expectation of the two terms is the same if the mean levels of activities do not vary in time. The rule only amplifies the transitions (or fluctuations) taking place in the activity, when a burst of activity takes place either at the pre-synaptic or at the post-synaptic levels. In a recurrent neural network, the global effect of the rule is not easy to anticipate.

Some trends can however be drawn. For instance, starting from an irregular regime, the STDP rule tends to produce more regularity and homogeneity, as measured thereafter.

Regularization In order to characterize the regularity of the self-sustained activity, we use an estimation of the effective number of Degrees of Freedom (#DOF) based on a Principal Components Analysis [18, 19], see figure 2.

Our data set is composed of the membrane potentials of all the neurons over sliding windows of 100 ms. A Principal Components Analysis is first applied to the data set, followed by a calculation of the entropy of the




Figure 2. -A- Estimated number of degrees of freedom during the first 5 seconds of simulation ($\sigma_J = 3.0$, $\tau_r = 2$ ms). -B- Evolution of #DOF when STDP is applied after 5 s of simulation with a learning coefficient $\alpha = 0.1$.




Figure 3. Evolution of the average firing rate when STDP is applied for different values of internal coupling. $\tau_r = 2$ ms.

normalised principal values p_i of the transformation matrix: $S = -\sum_{i=1}^N p_i \ln(p_i)$. This value is considered as an approximate log count of significant principal components weighted by their respective size, so that $\#DOF = e^S$ is an approximation of the effective number of degrees of freedom.

When Spike-Timing Dependent Plasticity is applied to a balanced random recurrent neural network with spontaneous, self-sustained activity, the number of degrees of freedom rapidly decreases (see figure 2B). This corresponds to a strong simplification of the activity, meaning that every neuron displays a closely similar response.

Homeostatic effect We now test the effect of STDP on several networks with spontaneous activity and various internal connection strengths (σ_J).

The effect of STDP on the average firing rate of the network depends on the internal weights (figure 3). Indeed, it seems that STDP has an homeostatic effect: networks with strong connections start with a high activity which is reduced by STDP, whereas networks with weak connections and low activity undergo an increase in their mean firing rate.

4 Networks under external stimulation

The most prominent effect of STDP on recurrent neural networks is the raise of periodicity in the activity. This periodicity can take several forms, from a purely synchronized regime to asynchronous periodic regime. The final outcome of the STDP application depends on the initial parameters, on the level activity, but also on the way external stimulations are sent to the network.

The inputs we send to the network are distributed among every neuron. We define a static stimulus P as a random vector of size N and whose values are randomly set according to a gaussian draw $\mathcal{N}(0, \sigma_I^2)$.

We actually define two different ways to send a stimulus to a network: the inputs are either directly added to the neuron potential (potential-based presentation):

$$V_i^{ext}(t) = P_i \text{ and } I_i^{ext}(t) = 0 \quad (4)$$

or to the current (current-based presentation):

$$V_i^{ext}(t) = 0 \text{ and } I_i^{ext}(t) = P_i \quad (5)$$

We use $\sigma_I = 1$ for potential-based input and $\sigma_I = 0.2$ for current-based input (the current input needs to be lower since it is integrated at each step).

Potential-based inputs First, we examine the case of potential-based inputs. We find that, in this case, while the initial activity is still irregular (upper part of the figure 4), the application of STDP with a strong enough learning coefficient always leads to a regularization of the network's activity toward a periodical regime. Two final outcomes can however be observed: either a synchronization of the neurons, where they all fire in a short window of time and remain silent afterward until the next burst of activity (see figure 4A), or a quasi-periodical "synfire" asynchronous firing (see figure 4B).

The refractory period and the internal coupling strength seem to play an important role into the triggering of a synchronous regime or not, as can be observed in the figure 5.

High internal coupling leads to asynchrony, while increasing the refractory period enhances synchronization (synchronization could not be obtained with a refractory period of $\tau_r = 1$ ms).

Current-based inputs With a gaussian current-based input presentation, we never observe the apparition of synchrony; the activity is "only" periodical (figure 6A).

However, if we use binary inputs, exciting only 20% of the neurons so they fire at their maximal frequency, some form of synchrony can be observed (figure 6B) between the excited neurons. The neurons which are not excited also fire in synchrony but in antiphase with the excited neurons, even if there are a few exceptions of neurons firing in phase with the "wrong" group.




Figure 4. Examples of activity, before and after STDP. A black bar is plotted when a neuron fires at a given time. Neurons are ordered by their input levels: the first neurons have the highest input and the last ones have the lowest ones. Upper figures represent activities before STDP, and the lower figures after STDP. -A- Asynchronous aperiodical activity before applying STDP and synchronous periodical answer after STDP. $\sigma_J = 2.0$; $\tau_r = 3$ ms. -B- Irregular activity before STDP and asynchronous periodical activity after. $\sigma_J = 5.0$; $\tau_r = 3$ ms.




Figure 5. Synchronized outcome depends on the parameters. Here, the resulting regime is a function of the internal connection strength and of the refractory period. This figure is based on observation for 20 different networks.




Figure 6. -A- Activity before and after applying STDP to a network with current-based input. $\sigma_J = 2.0$; $\tau_r = 3$ ms. -B- Activity before and after applying STDP to a network where the current-based input only stimulates 20% of the neurons (the bottom ones). $\sigma_J = 2.0$; $\tau_r = 3$ ms.

Whereas there is synchronisation, we must emphasize that the emergence of two distinct groups in antiphase is very different from the synchrony resulting from STDP with potential-based inputs. Moreover, the initial period is very different: around 20 ms for current-based input, versus around 60 ms for potential-based.

5 Discussion

We showed that our implementation of Spike-Timing Dependent Plasticity makes possible the transition from a dynamical regime to another one, which is more ordered and periodical. Moreover, whether the resulting answer is synchronized or not seems to depend both on the network's parameters (refractory period and weights standard deviation) and on the input presentation type.

The results obtained on the current-based inputs indicate that a form of synchronization appears more easily if the neurons are explicitly divided in several categories (excited neurons / non-excited neurons). The potential-based input presentation naturally creates such categories, even with gaussian distribution of the input strength, since neurons receiving enough stimulation on their potential (i.e., more than the firing threshold) tonically fire at their maximum rate as long as they do not receive inhibition; however this distinction between tonic and phasic neurons could be less effective when the internal coupling is strong (see 4B versus 4A).

This aim of this study is to give an insight in the com-

plex interactions between the self-sustained activity and the synaptic adaptation. Despite the strong heterogeneity in synaptic weights and delays, a regular response is rapidly obtained in various conditions, but not every response is interesting in terms of sensory encoding. To our mind, the most interesting outcome is a *selective* synchronous regime, i. e. a fine tuning of the synchrony as a function of the “relevance” of the input (see also [20]). The conditions under which a selective synchronous response can be obtained in a balanced neural network (not too strong internal coupling, contrasted inputs, strong cell refractoriness) need more investigation. We need to estimate more precisely which of those conditions really matter, and then propose more realistic schemes of sensory encoding in recurrent neural networks, in order to compare with biological sensory systems.

References

- [1] Guo-Qiang Bi and Mu-Ming Poo. Synaptic modifications in cultured hippocampal neurons : Dependence on spike timing, synaptic strength, and post-synaptic cell type. *The Journal of Neuroscience*, 18(24):10464–10472, 1998.
- [2] Henry Markram, Joachim Lübke, Micheal Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275:213–215, 1997.
- [3] Li I. Zhang, Huizhong W. Tao, Christine E. Holt, William A. Harris, and Mu ming Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395:37–44, 1998.
- [4] L.F. Abbott Sen Song, Kenneth D. Miller. Competitive hebbian learning through spike-timing dependent synaptic plasticity. *Nature*, 2000.
- [5] Rudy Guyonneau, Rudy VanRullen, and Simon Thorpe. Neurons tune to the earliest spikes through stdp. *Neural Computation*, 17:559–879, 2005.
- [6] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8:183–208, 2000.
- [7] E. Daucé, O. Moynot, O. Pinaud, and M. Samuelides. Mean-field theory and synchronization in random recurrent neural networks. *Neural Processing Letters*, 14:115–126, 2001.
- [8] Stijn Cassenaer and Gilles Laurent. Hebbian stdp in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature*, 448:709–714, 2007.
- [9] E. Rodriguez, N. George, J.-P. Lachaux, J. Martinerie, B. Renault, and F.J. Varela. Perception’s shadow: long-distance synchronization of human brain activity. *Nature*, 397(6718):430–433, 1999.
- [10] Evgeniy V. Lubenov and Athanassios G. Siapas. Decoupling through synchrony in neuronal circuits with propagation delays. *Neuron*, 58:118–131, 2008.
- [11] Roland E. Suri and Terrence J. Sejnowski. Spike propagation synchronized by temporally asymmetric hebbian learning. *Biological Cybernetics*, 87:440–445, 2002.
- [12] W. Gerstner and W. Kistler. *Spiking Neuron Models. Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002.
- [13] B. Cessac. Increase in complexity in random neural networks. *Journal de Physique I*, 5:409–432, 1995.
- [14] Hédi Soula, Guillaume Beslon, and Olivier Mazet. Spontaneous dynamics of asymmetric random recurrent spiking neural networks. *Neural Computation*, 18:60–79, 2006.
- [15] M. C. W. van Rossum, G. Q. Bi, and G. G. Turrigiano. Stable hebbian learning from spike timing-dependent plasticity. *The Journal of Neuroscience*, 20(23):8812–8821, December 2000.
- [16] Eugene M. Izhikevich, Joseph A. Gally, and Gerald M. Edelman. Spike-timing dynamics of neuronal groups. *Cerebral Cortex*, 14:933–944, 2004.
- [17] Jean-Pascal Pfister and Wulfram Gerstner. Triplets of spikes in a model of spike timing-dependent plasticity. *The Journal of Neuroscience*, 26(38):9673–9682, 2006.
- [18] W. Andy Wright, Robert E. Smith, Martin Danek, and Phillip Greenway. A generalisable measure of self-organisation and emergence. In G. Dorffner, H. Bischof, and K. Hornik, editors, *Artificial Neural Networks - ICANN 2001*, 2001.
- [19] Alexandra Penn. Steps towards a quantitative analysis of individuality and its maintenance : A case study with multi-agent systems. In *Fifth German Workshop on Artificial Life: Abstracting and Synthesizing the Principles of Living Systems*, 2002.
- [20] C.A. Skarda and W.J. Freeman. How brains make chaos in order to make sense of the world. *Behav. Brain Sci.*, 10:161–195, 1987.