Specialisation in recurrent neural networks with spike-timing-dependent plasticity

Matthieu Gilson, David B. Grayden, Doreen A. Thomas, Anthony N. Burkitt, J. Leo van Hemmen

SPECIALISATION IN RECURRENT NEURAL NETWORKS WITH SPIKE-TIMING-DEPENDENT PLASTICITY

Matthieu Gilson
Department of Electrical and Electronic Engineering
The University of Melbourne
Melbourne, 3010 VIC, Australia
email: mgilson@bionicear.org

David B. Grayden, Doreen A. Thomas and Anthony N. Burkitt
Department of Electrical and Electronic Engineering
The University of Melbourne
Melbourne, 3010 VIC, Australia

J Leo van Hemmen
Physik Department T35
TU München
85747 Garching bei München, Germany

Abstract

We examine how Spike-Timing-Dependent Plasticity (STDP) can strengthen recurrent excitatory connections in a network of Poisson neurons stimulated by two pools of external inputs, where only one pool has spike-time correlation. We derive conditions on the STDP and network parameters such that the network exhibits a stable activity in terms of spiking-rates and show how the evolution of the recurrent weights is determined by the interplay between STDP, the synaptic mechanisms and the input correlation structure. The results indicate that STDP induces competition that causes the emergence of a bimodal distribution of the weights, while the mean of the weights rapidly reaches a stable equilibrium. This behaviour can be related to the emergence of functional areas in recurrently connected networks, such as self-organising maps inspired by the evolution of the visual cortex in the first weeks of new-born mammalian animals.

KEY WORDS
Learning, STDP, Pulse Recurrent Network, Self-organisation.

1 Introduction

The hypothesis that neural networks become functionally organised due to changes in the synaptic efficacies (or weights) according to the neuronal activity has received considerable experimental support. Spike-Timing-Dependent Plasticity (STDP) [1, 2, 3] is a fruitful candidate for this mechanism, which relies on the correlations between pre- and post-synaptic action potentials (or spikes). It has been observed in various brain areas that exhibit local recurrent connectivity, such as the cortex or the hippocampus. Thus, understanding the weight dynamics induced by STDP in recurrent architectures is a crucial step to investigate learning in the brain. We build our work on previous studies of the impact of STDP in feed-forward networks [4, 5] and more recently in recurrently connected architectures [6]. In our model, the "fast" neural dynamics (evolution of the spiking activity of the neurons) depends on the synaptic weights and on the network topology; then the synaptic weights are in turn changed by STDP according to the neural activity on a slower time scale.

2 Models

2.1 Additive STDP

Our version of STDP describes the change of the synaptic weight due to the precise timing of single spikes and of pairs of spikes [1, 4]. For a pre-synaptic spike and a post-synaptic spike at times $t^{\rm in}$ and $t^{\rm out}$ resp., the change in the synaptic strength (or weight) J is the sum of three contributions

$$\delta J \propto \begin{cases} w^{\rm in} & \text{at time } t^{\rm in} \\ w^{\rm out} & \text{at time } t^{\rm out} \\ W(t^{\rm in}-t^{\rm out}) & \text{at time } \max(t^{\rm in},t^{\rm out}) \end{cases} \quad (1)$$

The rate-based contributions $w^{\rm in}$ and $w^{\rm out}$ account for the effects of each pre-synaptic and each post-synaptic spike resp. The STDP window function W determines the correlation contribution in terms of the spike-time difference $t^{\rm in}-t^{\rm out}$ [1, 4]. These contributions are scaled by a learning rate η, typically very small ($\eta \ll 1$) to ensure that the learning processes are very slow compared to the other neuronal mechanisms.

A typical choice for the learning window function W is illustrated in Fig. 1, although the present analysis can account for any arbitrary shape of W. The left side, corresponding to $t^{\rm in}-t^{\rm out}<0$, has positive values of W, which induces potentiation when the pre-synaptic spike occurs before the post-synaptic spike, in accordance to Hebb's postulate. The right side, corresponding to the converse situation $t^{\rm in}-t^{\rm out}>0$, is negative and thus induces depression. We also assume $W(0)=0$ here, but this does not play any role in the dynamics in continuous time. We use an additive version of STDP [1, 4], in contrast to the multiplicative version of STDP [7] and other variants [5]. Note that in nu-

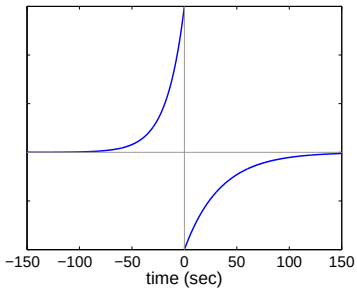

Figure 1. Example of STDP window function W. It consists of one decaying exponential for potentiation (*left*), with characteristic time constants equal to 17 ms. Likewise for depression (*right*) with a time constant of 34 ms.

merical simulation, this model requires to use bounds on the weights due to their divergence [4].

2.2 Poisson neuron

In this model, the firing mechanism of a given neuron i is approximated by an inhomogeneous Poisson process driven by an intensity parameter $\rho_i(t)$, in order to generate the output spike-time series $S_i(t)$ (or Dirac comb) [4]. Roughly speaking, a Poisson neuron generates a noisy output spike train with a given spiking-rate (or intensity). The Poisson parameter $\rho_i(t)$ can be related to the mean potential of the soma of neuron i and it evolves over time according to the spike-time series $\hat{S}_k(t)$ of pre-synaptic sources (indexed by k) as illustrated in Fig. 2

$$\rho_i(t) \;=\; \nu_0 + \sum_k \Big[J_k(t) \sum_n \epsilon\,(t - t_{k,n} - d_k) \Big]. \qquad (2)$$

The constant ν_0 is the spontaneous spiking-rate (identical for all the neurons). As in [4], each pre-synaptic spike at synapse k induces a variation of $\rho_i(t)$, called a post-synaptic potential (PSP), whose time course is determined by the normalised post-synaptic response kernel ϵ, and scaled by the synaptic efficacies (or weights) J_k. The delay d_k accounts for the axonal propagation duration on the pre-synaptic neuron. In order to preserve causality, $\epsilon(u) = 0$ for $u < 0$. The overall synaptic influx is the summation of the PSPs over all spike times $t_{k,n}$ (arriving at the k^{th} synapse, indexed by n). Note that we only consider positive weights here, i.e. excitatory synapses.

2.3 Link with physiology

In our model, the neurons have spontaneous activity and are also excited by external inputs (other neurons) that convey "neural information" encoded in their spiking-rate and correlation structures. We thus have two sources of stochasticity: that of the external inputs (spike trains described by their spiking-rate and correlation structures), which can

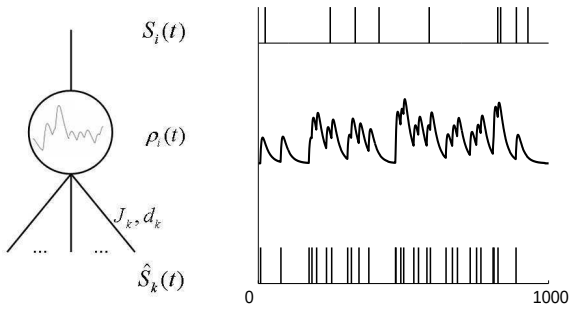

(a) Schematic view (b) Activity (time in ms)

Figure 2. Poisson neuron model. (a) Schematic view of the neuron, with the cell body or soma (circle), the dendritic inputs (below) and the output axon (above). (b) Illustration of the variation of $\rho_i(t)$ (*middle plot*) for a given succession of pre-synaptic spikes $\hat{S}_k(t)$ (*bottom spike train*) and the resulting output $S_i(t)$ (*top spike train*).

be related to the variability observed in physiological data; and that which can be understood as the impact of the background activity upon the firing of each neuron.

The Poisson neuron is a coarse approximation to biological activation mechanisms. However, we focus on the weight dynamics in the neural network, which is very slow compared to the activation dynamics of these neurons (cf. Sec. 2.1). Therefore, the activation dynamics need not be precisely accurate in order to qualitatively capture the main drift of the weight evolution [4, 6].

3 Theoretical analysis

3.1 Presentation of the network

Let us consider a network of N Poisson neurons stimulated with M Poisson spike trains with constant spiking-rates (see Fig. 3). In this paper, we consider the situation where only the recurrent weights J evolve according to STDP (cf. Eq. (1)); the input weights K are fixed. Note that we do not allow self-connections of neurons. In order to describe the neural activity, the variables of importance are the spiking-rates and the correlations, here averaged over a given time period T, as well as the synaptic weights [4, 6].

We incorporate the delays in the analysis via the covariance coefficients C_{ij}^W between neurons i and j,

$$\begin{aligned} C_{ij}^W(t) &:= \int_{-\infty}^{+\infty} W(u) C_{ij}(t, u - d_{ij})\,\mathrm{d}u \qquad (3)\\ C_{ij}(t,u) &:= Q_{ij}(t,u) \;-\; \nu_i(t)\,\nu_j(t) \\ Q_{ij}(t,u) &:= \frac{1}{T}\int_{t-T}^{t} \big\langle S_i(t') S_j(t'+u) \big\rangle\,\mathrm{d}t' \\ \nu_i(t) &:= \frac{1}{T}\int_{t-T}^{t} \big\langle S_i(t') \big\rangle\,\mathrm{d}t', \end{aligned}$$

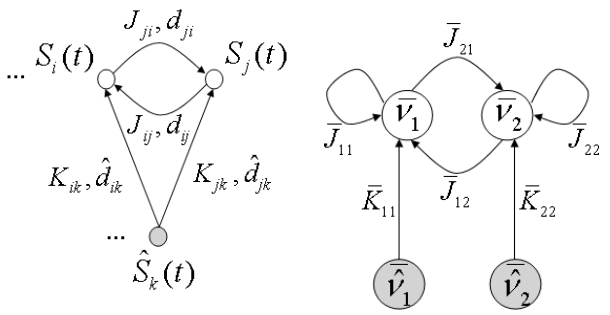

(a) Schematic view of three neurons in the network

(b) Network topology

	external	cross	internal
spiking-rates	$\hat{\nu}_k(t)$		$\nu_i(t)$
covar. coef.	$\hat{C}_{kl}(t,u)$	$F_{ik}(t,u)$	$C_{ij}(t,u)$
weights		$K_{ik}(t)$	$J_{ij}(t)$
delays		$\hat{d}_{ik}$	d_{ij}

(c) Table of the network variables

Figure 3. Presentation of the network and the notation. (a) The N internal neurons (open circles, indexed by $1 \leq i,j \leq N$ and with output spike trains denoted by $S_i(t)$) receive inputs from M external input spike trains $\hat{S}_k(t)$ ($1 \leq k \leq M$). (b) The internal neurons are split into two groups (open circles), which each are stimulated by only one of the two external input pools (filled circles). Here the overlined variables with subscripts 1 and 2 correspond to the mean values (related to $\hat{\nu}$, ν, K, J, etc.) over each pool of external inputs, each group of internal neurons or pair of these. (c) The table shows the variables that describe the neural activity (spiking-rates $\hat{\nu}$ and ν, covariance coefficients $\hat{C}$, F and C), and those related to the synaptic connections (weights K and J; delays $\hat{d}$ and d).

where Q_{ij} is the correlation between neurons i and j, and ν_i the time-averaged spiking-rate of neuron i. The delay d_{ij} in Eq. (3) accounts for the axonal propagation of the spikes and the neurotransmitter diffusion, i.e. a spike fired at time t_0 by neuron j arrives at the post-synaptic site of the connection $j \to i$ at time $t^{\rm in} = t_0 + d_{ij}$ (cf. Eq. (1)); we neglect the dendritic propagation here.

3.2 Dynamical system that describes the neural dynamics

A dynamical system of matrix equations can be derived to model the evolution of the neural activity in the network (or neural dynamics) in a manner similar to [6], using the following assumptions

- the expected values of external input spiking-rates and pairwise covariances are constant in time, which is here equivalent to constant time-averaged spiking-rates and covariances for any realization of the spiking history;
- the separation of the time scales of the activation mechanisms and the learning dynamics, the latter happening on a slower time scale (T is chosen in between these two time scales);
- the expected values of the neuron spiking-rates and pairwise covariances are quasi-constant in time, i.e. they only vary due to the slow learning on the weights;
- the input and recurrent delays are sharply distributed: $\hat{d}_{ik} \simeq \hat{d}$ for all suitable indices i and k, $d_{ij} \simeq d$ for all i and j;
- the PSP kernel ϵ and the recurrent delays (d) have a short duration compared to that of W.

After some calculations inspired by [8], we obtain the following equations that approximate the evolution of the network activity

$$\begin{aligned} \nu &= \left(\mathbb{1}_N - J\right)^{-1} \left(\nu_0\, E + K\, \hat{\nu}\right) \qquad (4)\\ C^W &= \left(\mathbb{1}_N - J\right)^{-1} K\, \hat{C}^{W*\zeta}\, K^{\mathbf{T}} \left(\mathbb{1}_N - J\right)^{-1\,\mathbf{T}} \\ \dot{J} &= \Phi_J\big(w^{\rm in}\, E\, \nu^{\mathbf{T}} + w^{\rm out}\, \nu\, E^{\mathbf{T}} + \widetilde{W}\, \nu\, \nu^{\mathbf{T}} + C^W\big). \end{aligned}$$

Time has been rescaled to remove η and the dependence on time t of the variables is omitted here and in the remainder of the paper. The $N \times N$ matrix $\mathbb{1}_N$ is the identity, the N-column vector E is filled with ones and the superscript $\mathbf{T}$ denotes transposition. Φ_J is a projector in the space of the matrices J that nullifies non-existing connections, in particular self-connections [6]. The scalar $\widetilde{W}$ is the integral value of the STDP window function W

$$\widetilde{W} = \int W(u)\,\mathrm{d}u. \qquad (5)$$

The function ζ describes the filtering of the short-term neuronal mechanisms on the input covariance $\hat{C}$ to obtain the neuron covariance C

$$\zeta(r) \simeq \int \epsilon(r + r' + d)\epsilon(r')\,\mathrm{d}r'. \qquad (6)$$

The coefficient $\hat{C}^{W*\zeta}$ is defined similarly to C^W in Eq. (3) by replacing W by $W*\zeta$ (where $*$ denotes the convolution) and C by $\hat{C}$. With correlated inputs generated as described in [5] for a spiking-rate $\hat{\nu}_0$ and a correlation level $\hat{c}$, we have

$$\hat{C}^{W*\zeta}_{kl} = \begin{cases} \hat{c}\,\hat{\nu}_0\,[W*\zeta](0) & \text{for } k \text{ and } l \text{ in same pool,} \\ 0 & \text{otherwise.} \end{cases} \qquad (7)$$

The system Eq. (4) differs from that in [6] in the use of "covariance" coefficients instead of "correlation" coefficients, and in the presence of ζ. Note that this dynamical system only captures the main drift (i.e. "first-order") of the stochastic evolution of the weights. However the present framework can still be used to evaluate higher orders of the network dynamics, such as the variance of the weights [4, 6].

4 Analysis of the evolution of the recurrent weights

We focus on a particular network topology: the external inputs are divided into two independent pools of the same size, with homogeneous spiking-rates and correlations within each; each input pool excites half of the recurrently connected neurons as illustrated in Fig. 3(b). The recurrent connectivity is full except for self-connections. This situation can be obtained, for example, by symmetry breaking on the input weights K due to a previous learning with STDP on the K [5]. We can work in a reduced-dimension space, averaging the spiking-rates and the weights over each pool and group (cf. Fig. 3(b)). We define

$$\begin{aligned} \widetilde{\nu} &:= \nu_0\, E + K\,\hat{\nu} \\ \widetilde{C} &:= K \hat{C}^{W*\zeta} K^{\mathbf{T}}, \end{aligned} \tag{8}$$

which are determined by the fixed external inputs.

4.1 Homeostatic equilibrium

Using a similar analysis to that in [6] for the case of no external inputs, we examine the conditions of equilibrium on the mean values of the spiking-rates $\nu_{\rm av} = (\bar{\nu}_1 + \bar{\nu}_2)/2$ and of the weights $J_{\rm av} = (\bar{J}_{11} + \bar{J}_{21} + \bar{J}_{12} + \bar{J}_{22})/4$ averaged over all the neurons and recurrent connections in the whole network, cf. Fig. 3(b). The homeostatic equilibrium, i.e. the situation when $\nu_{\rm av}$ and $J_{\rm av}$ correspond to a zero drift $\dot{J}_{\rm av} = 0$, corresponds to

$$\begin{aligned} \nu^*_{\rm av} &= -\frac{w^{\rm in} + w^{\rm out}}{\widetilde{W} + \widetilde{C}_{\rm av}/\widetilde{\nu}^2_{\rm av}} \\ J^*_{\rm av} &= \frac{\nu^*_{\rm av} - \widetilde{\nu}_{\rm av}}{n^J_{\rm av}\nu^*_{\rm av}}, \end{aligned} \tag{9}$$

where the subscript 'av' denotes the means of the variables over the whole network, and $n^J_{\rm av}$ is the average number of pre-synaptic recurrent connections for the neurons. The integral value $\widetilde{W}$ in Eq. (5) of the STDP window function W reflects the balance between potentiation and depression. The homeostatic equilibrium is stable provided that

$$\widetilde{W} + \frac{\widetilde{C}_{\rm av}}{\widetilde{\nu}^2_{\rm av}} \;<\; 0. \tag{10}$$

For weak correlation, this expression reduces to $\widetilde{W} < 0$, which implies more depression than potentiation induced by STDP for uncorrelated inputs. Note that it is realisable only if the equilibrium mean spiking-rate $\nu^*_{\rm av} > 0$, which requires that $w^{\rm in} + w^{\rm out} > 0$. These last two conditions are the same as those obtained in the case of no external inputs [6].

4.2 Learning the input correlation structure

When input pool #1 has correlation and pool #2 has none, it can be shown that the weight means $\bar{J}_{11}$ and $\bar{J}_{21}$ (i.e. weights coming from the neuron group #1, see Fig. 3(b)) in the dynamical system of Eq. (4) will evolve in the same direction, opposite to that for $\bar{J}_{12}$ and $\bar{J}_{22}$ (i.e. from group #2). This results from two concomitant effects of STDP: the mean pre-synaptic weight stabilises for each neuron (i.e. the spiking-rate of each neuron reaches a stable equilibrium), which implies in particular that $\bar{J}_{11} + \bar{J}_{12} = \text{const.}$, while the input correlation causes the difference $\bar{J}_{11} - \bar{J}_{12}$ to grow either positively or negatively, depending on the sign of $[W * \zeta](0)$. Likewise for $\bar{J}_{21}$ and $\bar{J}_{22}$. The weights $\bar{J}_{11}$ and $\bar{J}_{21}$ are potentiated if $[W * \zeta](0) > 0$ (cf. Eq. (7)), which means that neuron group #1 dominates group #2 in the recurrent network, as illustrated in Fig. 4. The converse situation happens when $[W * \zeta](0) < 0$.

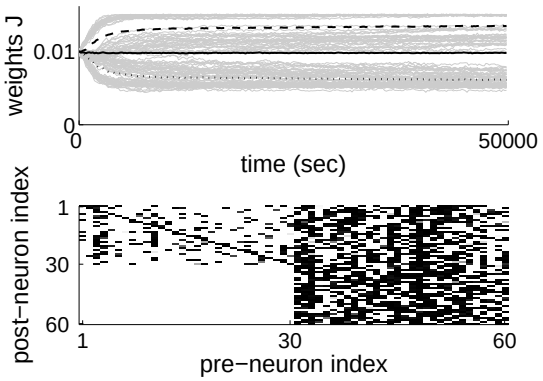

Figure 4. Evolution of the recurrent weights J due to correlated input stimulation. The network corresponds to Fig. 3(b) with two groups of $N = 30$ neurons each stimulated by only one pool of 30 inputs. Only input pool #1 has non-zero spike correlation. The STDP parameters corresponded to $\widetilde{W} < 0$ and $[W * \zeta](0) > 0$. *Top*: For each neuron, two grey lines represent the means of its pre-synaptic recurrent weights coming from group #1, and from group #2; the thick lines are the corresponding means over all the neurons, i.e. $(\bar{J}_{11} + \bar{J}_{21})/2$ for the dashed line and $(\bar{J}_{12} + \bar{J}_{22})/2$ for the dotted line (the solid black line is $\bar{J}_{\rm av}$). The weights coming from group #1 became potentiated (dashed line), while the weights coming from group #2 (dotted line) became depressed. *Bottom*: Weight matrix $\bar{J}$ after the emergence of the structure. Lighter pixels indicate potentiated weights: the left side (with means $\bar{J}_{11}$ and $\bar{J}_{21}$) has become more potentiated than the right side ($\bar{J}_{12}$ and $\bar{J}_{22}$). The black diagonal corresponds to the non-existent self-connections.

Starting from initial homogeneous recurrent weights, STDP causes them to split into a bimodal distribution, which corresponds to the emergence of a feed-forward pathway in the recurrent connections. The asymptotic values of the weight means for each subset ($\bar{J}_{11}$, ...) depend on the bounds in a complex manner, since the weights means $\bar{J}_{11}, \ldots$, are separated while $J_{\rm av}$ is constrained to a stable value. It was observed in general that the potentiation of $\bar{J}_{21}$ was higher than that of $\bar{J}_{11}$, but further examination is necessary to understand the reason, which does not clearly follow from the analysis of the system of equations.

5 Conclusion

We have shown that a simple model of STDP can cause the recurrently connected weights to self-organise in a network of Poisson neurons stimulated by two pools of external inputs, where only one pool has spike-time correlation. The results indicate that while the mean of the weights rapidly reaches a stable equilibrium, the input correlation structure is captured by STDP so that a bimodal distribution of the recurrent weights emerges. This behaviour occurs for a broad range of STDP parameters, in particular $w^{\text{in}} + w^{\text{out}} > 0$ and $\widetilde{W} < 0$, i.e. STDP is asymmetric and induces more depression than potentiation for uncorrelated inputs. Moreover, the STDP window function W, the PSP kernel ϵ and the recurrent delay d (assumed identical) interplay through the sign of $[W * \zeta](0)$ to determine the splitting of the recurrent weights in the configuration we consider.

Our results can be related to self-organising maps [9, 10, 11]: after symmetry breaking on the input connections [4, 5], the areas sensitive to a given input pathway that conveys information (here correlation) develop stronger within-group feedback (i.e. more synchronisation) and takes over other areas which receive uncorrelated inputs ("no information"), under the condition $[W * \zeta](0) > 0$. In our model where the PSP kernel ϵ and the recurrent delays ($d_{ij} \simeq d$) have a short duration compared to that of W, the condition $[W * \zeta](0) > 0$ is satisfied for example when the potentiation induced by W for small time lags (near the origin) is stronger than the depression for similar small time lags, i.e. $W(0-) > |W(0+)|$, as illustrated in Fig. 1. A feed-forward pathway then emerges among the recurrent connections.

Further studies are needed to gain a better understanding of the weights dynamics due to STDP in recurrent networks, in particular in the interesting case of two balanced correlated input pools. Minimal assumptions were deliberately made in order to obtain such a self-organisation scheme; incorporating a cortically-realistic spatial distribution of connections (including inhibitory ones) could lead to richer dynamics and consists of a natural extension of our theoretical framework, towards a more faithful representation of the cortical physiology. Studying the impact of a change of neuron model, such as the Integrate-and-Fire neuron, would also be of great interest since we showed here that the short-time neuronal mechanisms take part in determining the qualitative evolution of the recurrent weights.

Acknowledgments

The authors are grateful to C. Trengove, S. Byrnes, H. Meffin, M. Eager, I. Mareels, K. Borovkov, D. Nesic and B. Hughes for very helpful discussions that significantly improved the manuscript. MG is funded by scholarships from the University of Melbourne and from NICTA. ANB and DBG acknowledge funding from the Australian Research Council (ARC Discovery Projects #DP0453205 and #DP0664271), the Brockhoff Foundation, and The Bionic Ear Institute.

References

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

[2] H. Markram, J. Lubke, M. Frotscher, A. Roth, and B. Sakmann. Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *Journal of Physiology-London*, 500(2):409–440, 1997.

[3] G. Q. Bi and M. M. Poo. Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review of Neuroscience*, 24:139–166, 2001.

[4] R. Kempter, W. Gerstner, and J. L. van Hemmen. Hebbian learning and spiking neurons. *Physical Review E*, 59(4):4498–4514, 1999.

[5] R. Gutig, R. Aharonov, S. Rotter, and H. Sompolinsky. Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *Journal of Neuroscience*, 23(9):3697–3714, 2003.

[6] A. N. Burkitt, M. Gilson, and J. L. van Hemmen. Spike-timing-dependent plasticity for neurons with recurrent connections. *Biological Cybernetics*, 96(5):533–546, 2007.

[7] M. C. W. van Rossum and G. G. Turrigiano. Correlation based learning from spike timing dependent plasticity. *Neurocomputing*, 38:409–415, 2001.

[8] A. G. Hawkes. Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society Series B-Statistical Methodology*, 33(3):438–443, 1971.

[9] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69, 1982.

[10] Y. Choe and R. Miikkulainen. Self-organization and segmentation in a laterally connected orientation map of spiking neurons. *Neurocomputing*, 21(1-3):139–157, Oct 1998.

[11] O. G. Wenisch, J. Noll, and J. L. van Hemmen. Spontaneously emerging direction selectivity maps in visual cortex through stdp. *Biological Cybernetics*, 93(4):239–247, Oct 2005.