Available online at www.sciencedirect.com
SCIENCE DIRECT®
BioSystems 79 (2005) 11–20
BioSystems
www.elsevier.com/locate/biosystems

Dynamics of pruning in simulated large-scale spiking neural networks

Javier Iglesias[a,b,c,*], Jan Eriksson[b], François Grize[a], Marco Tomassini[a], Alessandro E.P. Villa[b,c,d]

[a] *Information Systems Department, University of Lausanne, Lausanne, Switzerland*
[b] *Laboratory of Neuroheuristics, University of Lausanne, Lausanne, Switzerland*
[c] *Laboratory of Neurobiophysics, University Joseph-Fourier, Grenoble, France*
[d] *Neuroheuristic Research Group, I.S.I. Foundation, Torino, Italy*

Abstract

Massive synaptic pruning following over-growth is a general feature of mammalian brain maturation. This article studies the synaptic pruning that occurs in large networks of simulated spiking neurons in the absence of specific input patterns of activity. The evolution of connections between neurons were governed by an original bioinspired spike-timing-dependent synaptic plasticity (STDP) modification rule which included a slow decay term. The network reached a steady state with a bimodal distribution of the synaptic weights that were either incremented to the maximum value or decremented to the lowest value. After 1×10^6 time steps the final number of synapses that remained active was below 10% of the number of initially active synapses independently of network size. The synaptic modification rule did not introduce spurious biases in the geometrical distribution of the remaining active projections. The results show that, under certain conditions, the model is capable of generating spontaneously emergent cell assemblies.

Keywords: Locally connected random network; Spike-timing-dependent synaptic plasticity; Spiking neural network; Large-scale simulation

1. Introduction

Massive synaptic pruning following over-growth is a general feature of mammalian brain maturation (Rakic et al., 1986; Zecevic and Rakic, 1991). Pruning starts near time of birth and is completed by time of sexual maturation. Biological mechanisms that regulate pruning mechanisms involve complex neurochemical pathways of cell signaling and are not intended to be reviewed here. Trigger signals able to induce synaptic pruning could be related to dynamic functions that depend on the timing of action potentials. Spike-timing-dependent synaptic plasticity (STDP) is a change in the synaptic strength based on the ordering of pre- and post-synaptic spikes. This mechanism has been proposed to explain the origin of long-term potentiation

* Corresponding author. Tel.: +41 21 692 35 87; fax: +41 21 692 35 85.
E-mail address: Javier.Iglesias@hec.unil.ch (J. Iglesias).

0303-2647/$ – see front matter © 2004 Elsevier Ireland Ltd. All rights reserved.
doi:10.1016/j.biosystems.2004.09.016

(LTP), i.e. a mechanism for reinforcement of synapses repeatedly activated shortly before the occurrence of a post-synaptic spike (Kelso et al., 1986; Bi and Poo, 1998; Froemke and Dan, 2002; Kepecs et al., 2002; Markram et al., 1997). STDP has also been proposed to explain long-term depression (LTD), which corresponds to the weakening of synapses strength whenever the pre-synaptic cell is repeatedly activated shortly after the occurrence of a post-synaptic spike (Karmarkar and Buonomano, 2002).

The glutamatergic NMDA receptors were initially identified as the receptor site with all biological features compatible with LTP induced by coincident pre- and post-synaptic cell discharges (Wigstrom and Gustafsson, 1986). The involvement of NMDA receptors in timing-dependent long-term depression (tLTD) has been recently described (Sjöström et al., 2003). Recent investigations suggest that glutamatergic receptors with AMPA channels and GABAergic receptors may also undergo modifications of the corresponding post-synaptic potentials as a function of the timing of pre- and post-synaptic activities (Engel et al., 2001; Woodin et al., 2003). These studies suggest that several mechanisms mediated by several neurotransmitters may exist at the synaptic level for changing the post-synaptic potential, either excitatory or inhibitory, as a function of the relative timing of pre- and post-synaptic spikes.

The important consequences that changes in synaptic strength may produce for information transmission, and subsequently for synaptic pruning, have raised an interest to simulate the activity of neural networks with embedded synapses characterized by STDP (Lumer et al., 1997; Fusi et al., 2000; Hopfield and Brody, 2004) The relation between synaptic efficacy and synaptic pruning (Chechik et al., 1999; Mimura et al., 2003), suggest that the weak synapses may be modified and removed through competitive "learning" rules. Competitive synaptic modification rules maintain the average neuronal input to a post-synaptic neuron, but provoke selective synaptic pruning in the sense that converging synapses are competing for control of the timing of post-synaptic action potentials (Song et al., 2000; Song and Abbott, 2001).

This article studies the synaptic pruning that occurs in a large network of simulated spiking neurons in the absence of specific input patterns. The originality of our study stands on the size of the network, up to 10,000 units, the duration of the experiment, 1,000,000 time units (one time unit corresponding to the duration of a spike), and the application of an original bioinspired STDP modification rule compatible with hardware implementation (Eriksson et al., 2003; Tyrrell et al., 2003). The network is composed of a mixture of excitatory and inhibitory connections that maintain a balanced input locally connected in a random way.

STDP is considered an important mechanism that modifies the gain of several types of synapses in the brain. In this study the synaptic modification rule was applied only to the excitatory–excitatory connections. This plasticity rule might produce the strengthening of the connections among neurons that belong to cell assemblies characterized by recurrent patterns of firing. Conversely, those connections that are not recurrently activated might decrease in efficiency and eventually be eliminated. The main goal of our study is to determine whether or not, and under which conditions, such cell assemblies may emerge from a large neural network receiving background noise and content-related input organized in both temporal and spatial dimension. In order to reach this goal the first step consisted in characterizing the dynamics of synaptic pruning in absence of content-related input. This first step is described here.

2. Models and methods

2.1. Network connectivity

The network is a 2D lattice folded as a torus to limit the edge effect where the units near the boundary received less input. The size of the network varies between 10×10 to 100×100 units. Several types of units may be defined. In this study we define two types, $q \in \{1, 2\}$, 80% of Type I ($q = 1$) units and 20% of Type II ($q = 2$) units are uniformly distributed over the network according to a space-filling quasi-random Sobol distribution (Press et al., 1992, Fig. 7.7.1). A unit of either type may project to a unit of either type, but self-connections are not allowed.

Each unit is assumed to be at the center of a relative 2D map, with coordinates $x = 0$, $y = 0$. The probability that another unit located at coordinates (x, y) receives a projection is provided by the following density function

$$G(x, y) = \alpha_{[q]} \exp\left(\frac{-2\pi(x^2 + y^2)}{\sigma^2_{[q]}}\right) + \phi_{[q]}$$

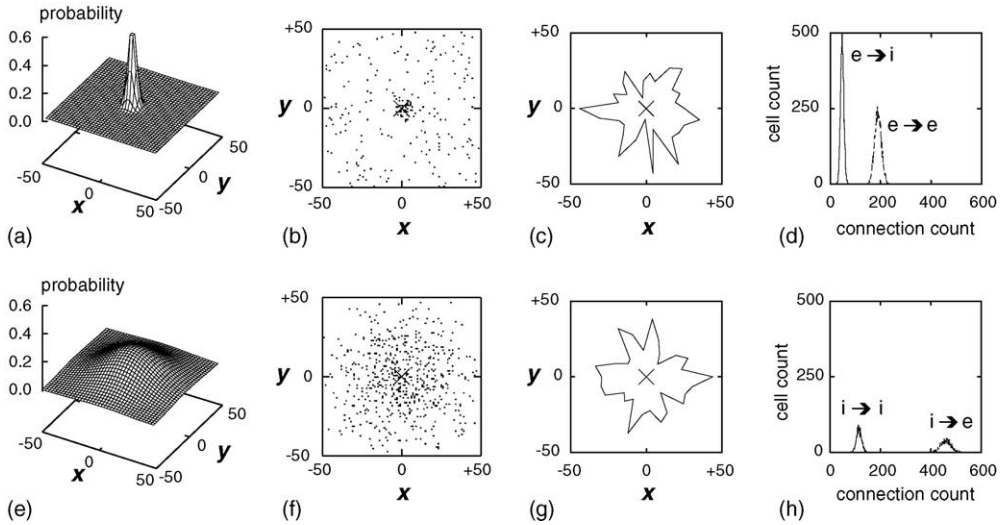

Fig. 1. Main features of the connectivity for Type I unit (upper row) and Type II unit (lower row). (a, e) Density function of the connectivity for a unit located at coordinates 0,0 on a 100 × 100 2D lattice; (b, f) Example of two projecting units, one for each class, located at the center of the 2D map. Each dot represents the location of a target unit connected by the projecting unit. (c, g) Orientation map of the projections of the same example units with polar coordinates smoothed with a bin equal to 12°. A circular line would represent a perfect pattern of isotropic connections. (d, h) Cumulative distributions of the connections. Type I are assumed to represent excitatory units (**e** →) and Type II inhibitory units (**i** →).

where $\alpha_{[q]}$ is the scaling factor for maximal probability of establishing a connection with the closest neighbors, $\sigma_{[q]}$ is a scaling factors for the skewness and width of the Gaussian shaped function, and $\phi_{[q]}$ is an uniform probability (Hill and Villa, 1997). The density function defining the probability of the connections is different for each type of unit and is illustrated in Fig. 1a and e. The values of the parameters used for the density functions are indicated in Table 1.

The random selection of the target units is run independently for each unit of either type. An example of the spatial distribution of the projections of one Type I unit, and of one Type II unit, is illustrated in Fig. 1b and f, respectively. In this example, the Type I unit (Fig. 1b) projects to 233 units and the Type II unit (Fig. 1f) projects to 537 units overall. For each unit it is possible to illustrate the orientation of its connections in the 2D lattice by plotting in polar coordinates the de-

Table 1
Parameter list of the main variables used for both types of units for 100 × 100 networks

Variable	Type I	Type II	Short description
	80	20	Proportion in network (%)
ϕ	2	0	Uniform connection probability (%)
α	60	20	Gaussian maximal probability (%)
σ	10	75	Gaussian distribution width
P	0.84	−1.60	Post-synaptic potential (mV)
V_{rest}	−78	−78	Membrane resting potential (mV)
θ_i	−40	−40	Membrane threshold potential (mV)
t_{refract}	1	1	Absolute refractory period (ms)
τ_{mem}	7	7	Membrane time constant (ms)
τ_{syn}	14	14	Synaptic plasticity time constant (ms)
τ_{act}	11000	11000	Activation time constant (ms)

See text for details.

viation from a perfect isotropic distribution. In case of an isotropic distribution the orientations would be illustrated by a circular line around the center. If such line is not circular it shows that some orientations have been selected preferentially by chance, as it may occur in a random selection procedure. The orientations of the projections of the two example units are illustrated in Fig. 1c and g. It appears that at the single unit level a large degree of anisotropy exists in the connection topology.

Fig. 1d shows the cumulative distribution of all connections established by Type I units projecting to either type. The modes of the histograms show that on average one unit of Type I is projecting to 50 units of Type II and to 190 units of Type I. Fig. 1h illustrates the cumulative distribution of all connections established by Type II units and shows that on average one unit of Type II is projecting to 115 units of Type II and to 460 units of Type I.

2.2. Neuromimetic model

All units of the network are simulated by leaky integrate-and-fire neuromimes. At each time step, the value of the membrane potential of the ith unit, $V_i(t)$, is calculated such that

$$V_i(t+1) = V_{\text{rest}[q]} + B_i(t) + (1 - S_i(t)) \times ((V_i(t) - V_{\text{rest}[q]})k_{\text{mem}[q]}) + \sum_j w_{ji}(t)$$

where $V_{\text{rest}[q]}$ corresponds to the value of the resting potential for the units of class type $[q]$, $B_i(t)$ is the background activity arriving to the ith unit, $S_i(t)$ is the state of the unit as expressed below, $k_{\text{mem}[q]} = \exp(-1/\tau_{\text{mem}[q]})$ is the constant associated to the current of leakage for the units of class type $[q]$, $w_{ji}(t)$ are the post-synaptic potentials of the jth units projecting to the ith unit.

The state of a unit $S_i(t)$ is a function of the membrane potential $V_i(t)$ and a threshold potential $\theta_{[q]_i}$, such that $S_i(t) = \mathcal{H}(V_i(t) - \theta_{[q]_i})$. $\mathcal{H}$ is the Heaviside function, $\mathcal{H}(x) = 0 : x < 0$, $\mathcal{H}(x) = 1 : x \geq 0$. In addition, the state of the unit depends on the refractory period $t_{\text{refract}[q]}$, such that

$$S_i(t + \Delta t) = \frac{(t_{\text{refract}[q]} - \Delta t)}{t_{\text{refract}[q]}} S_i(t)$$

for any $\Delta t < t_{\text{refract}[q]}$. For a refractory period equal to 1 time unit, the state $S_i(t)$ is a binary variable. In this simulation we assume that the refractory period is the same for all units of either type. It is assumed that a unit can generate a spike only for $S_i(t) = 1$. The parameter values used for the simulations are listed in Table 1.

2.3. Synaptic connections

The post-synaptic potential w_{ji} is a function of the state of the pre-synaptic unit S_j, of the "type" of the synapse $P_{[q_j,q_i]}$, and of the activation level of the synapse A_{ji}. This is expressed by the following equation

$$w_{ji}(t+1) = S_j(t)A_{ji}(t)P_{[q_j,q_i]}.$$

Notice that the "type" of the synapse is a parameter that depends on the types of units in the network. In the current study we assume that $P_{[1,1]}$, i.e. (Type I → Type I), and $P_{[1,2]}$ connections, i.e. (Type I → Type II), are of the same kind. The same assumption was made for $P_{[2,1]}$ and $P_{[2,2]}$ connections.

In order to maintain a balanced level of depolarization (excitatory) and hyperpolarization (inhibitory) the Type I unit was considered as excitatory and Type II as inhibitory. We set $P_{[1,1]} = P_{[1,2]} = 0.84\,\text{mV}$ and $P_{[2,1]} = P_{[2,2]} = -1.6\,\text{mV}$.

2.4. Synaptic modification rule

It is assumed a priori that modifiable synapses are characterized by activation levels $[A]$ with N attractor states $[A_1] < [A_2] < \cdots < [A_N]$. Activation levels of type $[1, 1]$ synapses are *integer-valued levels* $A_{ji}(t)$, with $A_{ji}(t) \in \{[A_1] = 0, [A_2] = 1, [A_3] = 2, [A_4] = 4\}$. Index j is referred to as the pre-synaptic unit and index i as the post-synaptic unit. We assume that post-synaptic potentials generated by synapses of type $[1, 1]$ correspond to synaptic currents mediated by NMDA glutamatergic receptors. These discrete levels could be interpreted as a combination of two factors: the number of synaptic *boutons* between the pre- and post-synaptic units and the changes in synaptic conductance as a result of Ca^{2+} influx through the NMDA receptors. In the current study we attributed a fixed activation level (that means no synaptic modification) $A_{ji}(t) = 1$, to $exc \rightarrow inh$, $inh \rightarrow exc$, and $inh \rightarrow inh$ synapses.

A *real-valued variable* $L_{ji}(t)$ is used to implement the spike-timing dependent plasticity rule for $A_{ji}(t)$, with integration of the timing of the pre- and post-synaptic activities. The variables $L_{ji}(t)$ are user-defined boundaries of attraction $L_0 < L_1 < L_2 < \cdots < L_{N-1} < L_N$ satisfying $L_{k-1} < [A_k] < L_k$ for $k = 1, \ldots, N$. This means that whenever $L_{ji} > L_k$ the activation variable A_{ji} jumps from state $[A_k]$ to $[A_{k+1}]$. Similarly, if $L_{ji} < L_k$ the activation variable A_{ji} jumps from state $[A_{k+1}]$ to $[A_k]$. Moreover, after a jump of activation level $[A]$ occurred at time t the real-valued variable L_{ij} is reset to $L_{ij}(t+1) = L_k + L_{k+1}/2$.

Spike-timing dependent plasticity (STDP) defines how the value of L_{ji} at time t is changed by the arrival of pre-synaptic spikes, by the generation of post-synaptic spikes and by the correlation existing between these events. On the generation of a post-synaptic spike (i.e., when $S_i = 1$), the value L_{ji} receives an *increment* which is a decreasing function of the elapsed time from the previous pre-synaptic spike at that synapse. Similarly, when a spike arrives at the synapse, the variable L_{ji} receives a *decrement* which is likewise a decreasing function of the elapsed time from the previous post-synaptic spike (i.e., when $S_j = 1$). This rule is summarized by the following equation: $L_{ji}(t+1) = L_{ji}(t) + (S_i(t)M_j(t)) - (S_j(t)M_i(t))$, where $S_i(t)$, $S_j(t)$ are the state variables of the ith and jth units and $M_i(t)$, $M_j(t)$ are interspike decay functions. $M_i(t)$ may be viewed as a "memory" of the latest interspike interval,

$$M_i(t+1) = S_i(t)M_{\max[q_i]} + (1 - S_i(t))(M_i(t)\ \exp(-t/\tau_{\mathrm{syn}[q_i]}))$$

where $\tau_{\mathrm{syn}[q_i]}$ is the synaptic plasticity time constant characteristic of each type of unit and $M_{\max[q_i]}$ was set $M_{\max[q_i]} = 2$ for all units of either type in this study. In the case that neither the pre- nor the post-synaptic unit is firing a spike, the real-valued variable will decay with a time constant $k_{\mathrm{act}[q_j,q_i]} = \exp(-1/\tau_{\mathrm{act}[q_j,q_i]})$ characteristic for each type of synapse, such that the final equation is the following:

$$L_{ji}(t+1) = L_{ji}(t)k_{\mathrm{act}[q_j,q_i]} + (S_i(t)M_j(t)) - (S_j(t)M_i(t)).$$

In the present study the differences between the user defined boundaries L_k were all equal, such that $\Delta L_k = L_k - L_{k-1} = 20$ for any attractor state $[A_k]$. At the begin of the simulation all modifiable synapses were set to activation level $[A_3] = 2$. Fig. 2a illustrates a case when the synaptic link receives a potentiation determined by the STPD rule described above. The activation variable jumps from $[A_3]$ to $[A_4]$ and stabilizes at highest activation level. Fig. 2b illustrates a case when the synapse is continuously depressed such that the activation variable jumps from $[A_3]$ to $[A_2]$, and then from $[A_2]$ to $[A_1]$ faster than its spontaneous decay, determined by time constant $k_{\mathrm{act}[q_j,q_i]}$. Fig. 2c illustrates a case when the synapse is neither depressed nor potentiated and the activation level spontaneously decay down to the minimal level.

2.5. Synaptic pruning

No generation of new projections is allowed in the present study, although specific rules could be defined

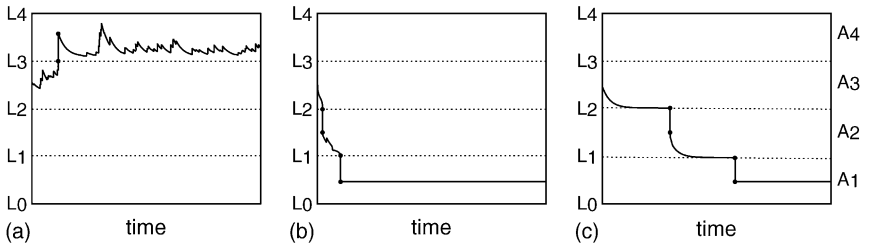

Fig. 2. Pruning dynamics. The *real-valued variable* L_{ji} is increased or decreased according to the STDP rule. If the value L_{ji} reaches one of the L_k user-defined boundaries a jump occurs in the *integer-valued variable* $[A_k]$. At the begin all (**e** → **e**) synapses have been set at activation level $[A_3]$. (a) Example of potentiation with an increase in synaptic strength that is stabilized on the long term. (b) Example of depression with a fast decrease in synaptic activation level down to its minimal level, $[A_1] = 0$ which provokes the elimination of the synapse. (c) Example of a synaptic link which neither affected by potentiation nor by depression, but its efficacy decays down to $[A_1] = 0$ according to the time constant τ_{act}.

to this purpose. Synaptic pruning occurs when the activation level of a synapse reaches a value of zero. This means that synaptic pruning may occur only for synaptic connections of type [1, 1], which are also the most abundant, when the activation level A_{ji} decreases to its minimal value, i.e. $[A_1] = 0$. In this case the synapse $[i, j]$ is eliminated from the network connectivity.

2.6. Background activity

The background activity $B_i(t)$ is used to simulate the input of afferents to the ith unit that are not explicitly simulated within the network. Let us assume that each type of unit receives $n_{\text{ext}[q_i]}$ external afferents. In the present study we simplify by setting that all units receive the same number of external projections and that all of them are excitatory. Namely, we assume $n_i \equiv n \equiv 50$ and that the post-synaptic potential generated by these external afferents is fixed to a value equal to $P_{[1,1]}$. In the current case (see Table 1) each external afferent generates an excitatory post-synaptic potential equal to 0.84 mV.

We assume that the external afferents are correlated among them. This means that each time a unit is receiving a correlated input from 50 external afferents its membrane potential is depolarized to an extent that will generate a spike. Such external input is distributed according to a Poisson process which is independent for each unit and with mean rate λ_i. The rate of external background activity is a critical parameter. We found that with all previous parameters being kept constant a rate of background activity $\lambda_i < 8$ spikes/s is unable to sustain any activity at all. In the present study we have set the Poisson input to a rate $\lambda_i = 10$ spikes/s.

2.7. Network size

We investigated the pruning dynamics with networks of different sizes. The smallest network was defined by 10×10 units and the largest network studied here was 100×100 units, i.e. $(10 \times N)^2$, with $N \in \{1, \ldots, 10\}$. To compensate for the changes in the balance between excitation and inhibition induced by the change of size, we introduced the scaling factor $f = \sqrt[3]{10^4/(10 \times N)^2}$, where N is the size as described before. The uniform probability for an excitatory unit to project to any other unit of the network was scaled according to $\phi^*_{[1]} = f\phi_{[1]}$, leading to a larger number of excitatory connections at the beginning of the simulations for smaller networks. The level of post-synaptic depolarization for excitatory–excitatory synapses was scaled as $P^*_{[1,1]} = (1 + f - 1/2)P_{[1,1]}$, so that the strength of these connections was larger for smaller networks. Table 2 lists the scaled values of $\phi^*_{[1]}$ and $P^*_{[1,1]}$ we used. Note that the values for $N = 10$ correspond to those listed in Table 1.

Table 2
The scaled parameter values for each network size N

N	Size	$\phi^*_{[1]}$ (%)	$P^*_{[1,1]}$ (mV)
1	10×10	9.28	2.36
2	20×20	5.84	1.64
3	30×30	4.46	1.35
4	40×40	3.68	1.19
5	50×50	3.17	1.08
6	60×60	2.81	1.01
7	70×70	2.53	0.95
8	80×80	2.32	0.90
9	90×90	2.14	0.87
10	100×100	2.00	0.84

See text for details.

2.8. Simulation tools

The simulator was a custom developed, Open Source, C program that relies on the GNU Scientific Library (GSL) for random number generation and quasi-random Sobol distribution implementations. A MySQL database back end stored all the configuration details and logs for later reference. This information was extracted and formatted by a set of PHP Web scripts to monitor the status of the simulations and create new ones. With our current implementation at the University of Lausanne, a 10,000 units network simulation for a duration of 1×10^6 time steps lasted approximatively 8 h, depending on the network global activity.

3. Results

All synapses of type [1, 1], i.e. ($exc \rightarrow exc$) were initialized with $A_{ij}(t = 0) = [A_2]$. In presence of background activity only, most synapses were characterized by a decrement of the activation level. After a long time, $t = t_{\text{steady}}$, the network activity is stabilized and STDP does not modify any more the activation level of the synapses. At time $t = t_{\text{steady}}$ most modifi-

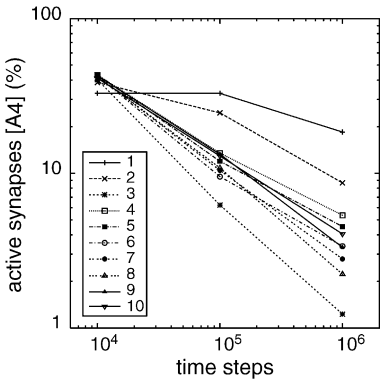

Fig. 3. Pruning dynamics averaged over $n = 10$ simulations for each network size. With the proposed size compensation factor, the pruning dynamics are comparable for network sizes $N \in \{4, \ldots, 10\}$. Simulations for $N = 1$ and 2 saturated, suggesting that the compensation factor was too large for those two specific sizes.

able synapses were eliminated and almost all remaining active synapses were characterized by the highest possible activation level, i.e. $[A_4]$. We observed that $t = t_{\text{steady}}$ could be as long as $t = 1 \times 10^6$ in several simulation runs.

3.1. Network size

It is interesting to notice that the final ratio of active synapses ($R_{\max[A]}$) represented only few percents of the initial number of synapses (Fig. 3). In addition, it is important to notice that those connections that reach the maximal activation level do not necessarily remain active until $t = t_{\text{steady}}$. Several synapses reached level $[A_4]$ after some delay, then their activation level decreased down to $[A_1] = 0$, at variable speed, and the synapse was eventually eliminated.

It is important that a network be attuned to work in a range such that background activity is unable to create spurious attractors by STDP. This means that background activity alone should not create stable connections that would shape the topology of cell assemblies. The size of the network is critical if the goal is to detect the emergence of chains of interconnected units embedded in a large network. Fig. 3 shows that the ratio of active synapses with activation level equal to $[A_4]$ could be as high as 50% of all initial synapses. The final percentage of active units is much less variable and at t_{steady} it is always less than 10% for networks that did not saturate.

3.2. "Seed" effect

A simulation study that relies on large use of randomly generated numbers may fall into local minima or spurious attractors simply by chance. It was necessary to assess the effect of the seed of the random number generator on our simulation. The most critical effect of the randomization may occur at the very beginning, when the initial network topology is created according to the density functions of the connections of the different types of units. The very same simulation, with the parameter set described in Table 1 was repeated 100 times using different random seeds with the largest network size, i.e. 100×100 units.

The choice of the seed had a significant impact on the value of $R_{\max[A]}$ at time t_{steady}, as it could vary in the range [1.30, 6.03]%. However, as shown by the distribution of $R_{\max[A]}$ (Fig. 4), about 90% of these values were comprised between 3.0 and 6.0%. Moreover, we never observed cases with absence of convergence at delays as large as $t = 1 \times 10^6$. This indicates that a "seed effect" exists but this does not cause changes in the overall dynamics of synaptic pruning.

A bias in the orientation of the connections could occur by random choices. In order to test this hypothesis two cases of extreme values of $R_{\max[A]}$ observed in the distribution of Fig. 4 were selected. The first case corresponds to $R_{\max[A]}$ as low as $R_1 = 1.97\%$. The second case corresponds to $R_{\max[A]}$ as high as $R_2 = 6.03\%$.

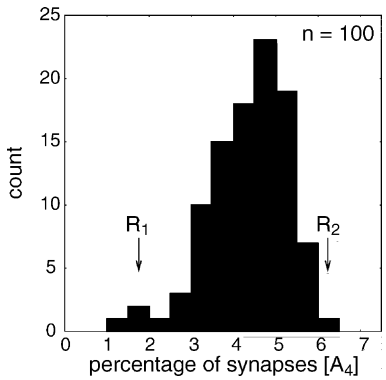

Fig. 4. Random seed effect on the number of synapses that remain after pruning. After $n = 100$ simulations that used different seeds the distribution of $R_{\max[A]}$ at time $t = t_{\text{steady}} = 1 \times 10^6$, with bin = 0.5, shows that in the majority of the runs synaptic pruning left 3.0–6.0% active synapses at the maximum level $[A_4]$.

For both cases we calculated the deviation from an isotropic connection, as defined previously for Fig. 1c and g, for all active synapses, i.e. with an activation level not equal to $[A_1] = 0$. Then we calculated an average deviation plot that corresponds to the mean of the orientations computed over all active synapses at given times t, namely at $t_1 = 1 \times 10^5$, $t_2 = 2 \times 10^5$, and $t = t_{\text{steady}} = 1 \times 10^6$.

Fig. 5a shows the evolution of the orientation map in the case R_1, when the network stabilizes with a low level of active connections. In this example the initial number n of connections at time t_0 was $n = 1,517,240$. At t_1 330,920 synapses remained, at t_2 172,503 synapses remained, and eventually the network stabilized with 30,864 excitatory–excitatory synapses. The orientation map shows that the deviations from an isotropic distribution were equally distributed in all directions. Another factor that could be affected by the random choice is the distance from source-to-target (calculated as an Euclidean distance over the 2D lattice) of the remaining projections. The histogram of the distribution of these distances Fig 5b was normalized with respect to the probability distribution of establishing a connection. In this normalized histogram a ratio of 1 means that the count is perfectly determined by the probability distribution. In the case of R_1 we observe that there was a tendency to some deviation from the original probability function, but

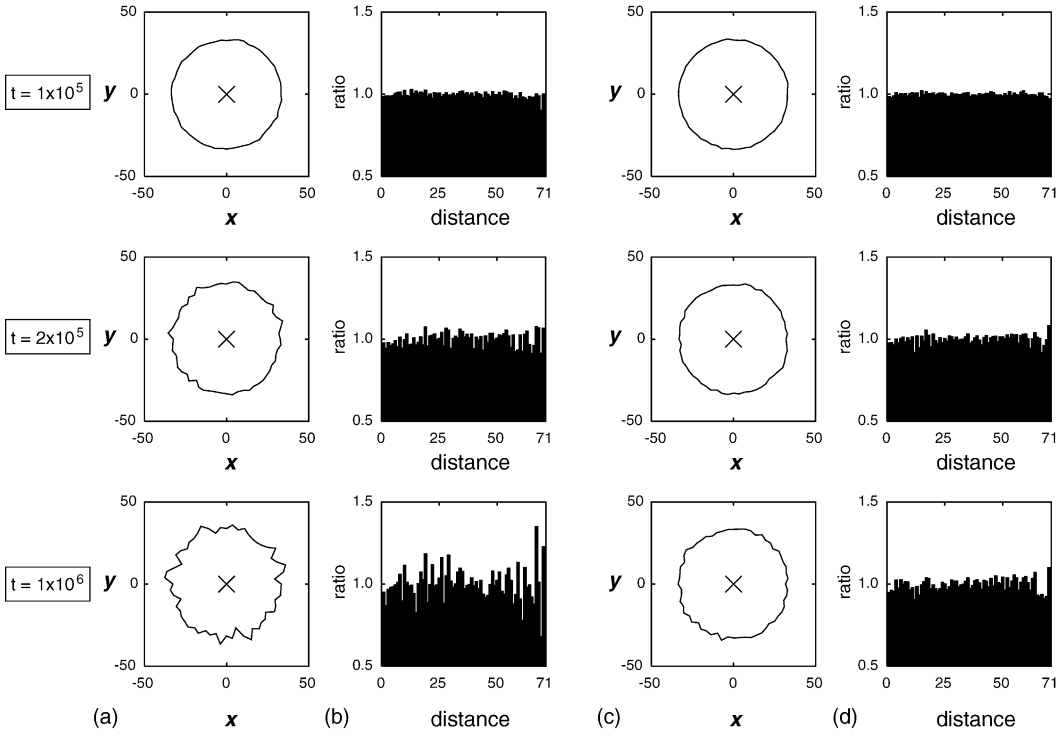

Fig. 5. Random seed effect on the orientation and length of active connections. (a) Average orientation map. A circular line indicates an isotropic orientation of the projections of a unit ideally located at the center marked by a cross. This simulation correspond to run R_1 of Fig. 4. The average deviation from isotropy for all active connections is plotted at various times. The last line correspond to t_{steady}: 30,864 synapses remained active, all with active state $[A_4]$, which corresponded to 1.97% of the initial number of synapses at time t_0. (b) Normalized histogram of the length of the source-to-target projections measured in Euclidean distance in the 2D lattice for simulation run R_1. A flat line at $ratio = 1$ indicates that the distances are totally predicted by the modified 2d Gaussian distribution function described in the text (cf. Section 2.1). (c) Average orientation map corresponding to run R_2 of Fig. 4. At t_{steady} 93,346 synapses remained active, all with active state $[A_4]$, which corresponded to 6.03% of the initial number of synapses at time t_0. (d) Normalized histogram of the length of the source-to-target projections measured in Euclidean distance in the 2D lattice for simulation run R_2.

this variance was the same for any source-to-target distance. In the case R_2, the number of initial synapses was $n = 1,512,634$ and at t_{steady} 93,346 active synapses remained. This analysis shows that the "seed" effect does not introduce significant biases neither in the orientation, nor in the length of the connections that were selected by pruning.

4. Discussion

We assumed a number of simplified hypotheses about the presence of only two types of units, their leaky integrate-and-fire dynamics, their distribution and the dynamics of the transfer functions of the synapses that connect these units. With all these assumptions we observed that the network reached a steady state when the synaptic weights were either incremented to the maximum value or decremented to the lowest value. Our result is in agreement with the bimodal distribution of synaptic strengths observed with a different STDP-based model (Chechik et al., 1999; Song et al., 2000; Song and Abbott, 2001). This effect is interpreted as the effect of the STDP rule that leads pre-synaptic neurons to compete for the capacity to drive the post-synaptic unit to produce an all-or-none output signal, akin to an action potential.

The choice of a 2D lattice topology allowed us to study the effect of incrementing the network size from 10×10 to 100×100 units. It was interesting to observe that the final ratio of synapses that remained active (that we labeled here $R_{\max[A]}$) was below 10% of the number of initially active synapses. This ratio varied only slightly with changes in network size (MacGregor et al., 1995) and the effect of different random seeds at the initialization was also limited. On the other hand, we observed that the ratio of active synapses characterized by the maximum weight could transiently reach a proportion as high as 50% of all initial synapses. This could indicate that a very large network may not be necessary for recurrent networks to emerge. Interconnected sizable "modules" up to 50×50 or 60×60 units embedded in larger networks may offer a more efficient way to recruit active synapses that compete for generating a post-synaptic spike.

A bias in the geometrical orientation of the synapses at the network level might produce important effects on the global dynamics as it could introduce singularities in the network topology. These singularities could sustain attractors with unbalanced excitatory/inhibitory inputs if they were the consequence of content-related inputs (Hill and Villa, 1997). In presence of only background noise these attractors would be spurious and could mask input-related features. We observed that the reinforcement of few synapses occurred without geometric distortions in both direction and source-to-target distance over the 2D lattice. Synaptic pruning proceeded in an homogeneous and isotropic way across the network. This result suggests that the implementation of the current STDP rule is equivalent to random pruning and does not introduce spurious biases.

The present work is currently extended in two directions that will be reported in future articles. The first direction consists in studying the effect of different synaptic transfer functions without changing the other parameters of the simulation. This would introduce temporally asymmetrical STDP (Bi and Wang, 2002), where both time window and efficacy changes are different for LTD than for LTP (Stuart and Hausser, 2001). In addition to the modifiable synaptic rule it would be interesting to introduce more realistic transfer functions in all types of synapse, in particular in the inhibitory synapses, to account for modulation frequencies and pre-synaptic spike interval distribution (Segundo et al., 1995a,b).

The second extension of this work is the introduction of content-related inputs, i.e. spatiotemporal patterns of discharges associated to selected stimuli (Abeles, 1991; Hopfield and Brody, 2000; Villa, 2000). Preliminary observations carried out with a simplified version of this simulation have demonstrated that these types of network may store the traces of time-varying stimuli such that similar stimuli blurred with noise can evoke an activity pattern close to the original one (Eriksson et al., 2003; Torres et al., 2003).

Acknowledgements

The authors thank Dr. Yoshiyuki Asai for discussions and comments on the manuscript. This work is partially funded by the Future and Emerging Technologies program (ISTFET) of the European Community, under grant IST-2000-28027 (POETIC), and under grant OFES 00.0529-2 by the Swiss Government.

References

Abeles, M., 1991. Corticonics: Neural Circuits of the Cerebral Cortex. Cambridge University Press.

Bi, G.Q., Poo, M.M., 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. J. Neurosci. 18 (24), 10464–10472.

Bi, G.Q., Wang, H.X., 2002. Temporal asymmetry in spike timing-dependent synaptic plasticity. Physiol. Behav. 77 (4/5), 551–555.

Chechik, G., Meilijson, I., Ruppin, E., 1999. Neuronal regulation: a mechanism for synaptic pruning during brain maturation. Neural Computation 11, 2061–2080.

Engel, D., Pahner, I., Schulze, K., Frahm, C., Jarry, H., Ahnert-Hilger, G., Draguhn, A., 2001. Plasticity of rat central inhibitory synapses through GABA metabolism. J. Physiol. 535 (2), 473–482.

Eriksson, J., Torres, O., Mitchell, A., Tucker, G., Lindsay, K., Rosenberg, J., Moreno, J.-M., Villa, A.E.P., 2003. Spiking neural networks for reconfigurable POEtic tissue. Lecture Notes Comput. Sci. 2606.

Froemke, R.C., Dan, Y., 2002. Spike-timing-dependent synaptic modification induced by natural spike trains. Nature 416 (6879), 433–438.

Fusi, S., Annunziato, M., Badoni, D., Salamon, A., Amit, D.J., 2000. Spike-driven synaptic plasticity: theory, simulation, VLSI implementation. Neural Comput. 12, 2227–2258.

Hill, S.L., Villa, A.E.P., 1997. Dynamic transitions in global network activity influenced by the balance of excitation and inhibition. Network: Computat. Neural Syst. 8, 165–184.

Hopfield, J.J., Brody, C.D., 2000. What is a moment? "Cortical" sensory integration over a brief interval. Proc. Natl. Acad. Sci. USA 97 (25), 13919–13924.

Hopfield, J.J., Brody, C.D., 2004. Learning rules and network repair in spike-timing-based computation networks. Proc. Natl. Acad. Sci. USA 101 (1), 337–342.

Karmarkar, U.R., Buonomano, D.V., 2002. A model of spike-timing dependent plasticity: one or two coincidence detectors? J. Neurophysiol. 88 (1), 507–513.

Kelso, S.R., Ganong, A.H., Brown, T.H., 1986. Hebbian synapses in hippocampus. Proc. Natl. Acad. Sci. USA 83 (14), 5326–5330.

Kepecs, A., van Rossum, M.C.W., Song, S., Tegner, J., 2002. Spike-timing-dependent plasticity: common themes and divergent vistas. Biol. Cybernet. 87, 446–458.

Lumer, E.D., Edelman, G.M., Tononi, G., 1997. Neural dynamics in a model of the thalamocortical system. ii. The role of neural synchrony tested through perturbations of spike timing. Cerebral Cortex 7 (3), 228–236.

MacGregor, R.J., Ascarrunz, F.G., Kisley, M.A., 1995. Characterization, scaling, and partial representation of neural junctions and coordinated firing patterns by dynamic similarity. Biol. Cybernet. 73 (2), 155–166.

Markram, H., Lubke, J., Frotscher, M., Sakmann, B., 1997. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275 (5297), 213–215.

Mimura, K., Kimoto, T., Okada, M., 2003. Synapse efficiency diverge due to synaptic pruning following over-growth. Phys. Rev. E Stat. Nonlin. Soft. Matter. Phys. 68.

Press, W.H., Flannery, B.P., Teukolsky, S.A., Vetterling, W.T., 1992. Numerical Recipes in C: The Art of Scientific Computing, second ed. Cambridge University Press.

Rakic, P., Bourgeois, J.P., Eckenhoff, M.F., Zecevic, N., Goldman-Rakic, P.S., 1986. Concurrent overproduction of synapses in diverse regions of the primate cerebral cortex. Science 232 (4747), 232–235.

Segundo, J.P., Stiber, M., Vibert, J.F., Hanneton, S., 1995a. Periodically modulated inhibition and its postsynaptic consequences. i. General features, influence of modulation frequency. Neuroscience 68 (3), 657–692.

Segundo, J.P., Stiber, M., Vibert, J.F., Hanneton, S., 1995b. Periodically modulated inhibition and its postsynaptic consequences. ii. Influence of modulation slope, depth, range, noise and of postsynaptic natural discharges. Neuroscience 68 (3), 693–719.

Sjöström, P.J., Turrigiano, G.G., Nelson, S.B., 2003. Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. Neuron 39, 641–654.

Song, S., Abbott, L.F., 2001. Cortical development and remapping through spike timing-dependent plasticity. Neuron 32 (2), 339–350.

Song, S., Miller, K.D., Abbott, L.F., 2000. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nat. Neurosci. 3, 919–926.

Stuart, G.J., Hausser, M., 2001. Dendritic coincidence detection of epsps and action potentials. Nat. Neurosci. 4 (1), 63–71.

Torres, O., Eriksson, J., Moreno, J.M., Villa, A.E.P., 2003. Hardware optimization of a novel spiking neuron model for the POEtic tissue. Lecture Notes Comput. Sci. 2687, 113–120.

Tyrrell, A.M., Sanchez, E., Floreano, D., Tempesti, G., Mange, D., Moreno, J.-M., Rosenberg, J., Villa, A.E.P., 2003. POEtic: An integrated architecture for bio-inspired hardware. Lecture Notes Comput. Sci. 2606.

Villa, A.E.P., 2000. Time and the brain. In: Miller, R. (Ed.), Empirical Evidence About Temporal Structure in Multi-Unit Recordings, vol. 2. Harwood Academic Publishers, pp. 1–51.

Wigstrom, H., Gustafsson, B., 1986. Postsynaptic control of hippocampal long-term potentiation. J. Physiol. 81 (4), 228–236.

Woodin, M.A., Ganguly, K., Poo, M., 2003. Coincident pre- and post-synaptic activity modifies GABAergic synapses by postsynaptic changes in Cl^- transporter activity. Neuron 39, 807–820.

Zecevic, N., Rakic, P., 1991. Synaptogenesis in monkey somatosensory cortex. Cerebral Cortex 1 (6), 510–523.