Neural Networks 21 (2008) 1272–1277

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

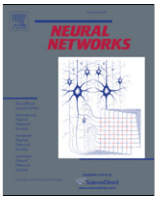

Instabilities in attractor networks with fast synaptic fluctuations and partial updating of the neurons activity[☆]

J.J. Torres [a,*], J. Marro [a], J.M. Cortes [b], B. Wemmenhove [c]

[a] Institute "Carlos I" for Theoretical and Computational Physics, and Departamento de Electromagnetismo y Física de la Materia, Universidad de Granada, E-18071, Granada, Spain
[b] Institute for Adaptive and Neural Computation, School of Informatics, University of Edinburgh, EH1 2QL, UK
[c] Department of Biophysics and SNN, Radboud University, 6525 EZ Nijmegen, The Netherlands

ARTICLE INFO

Article history:
Received 18 April 2007
Received in revised form
15 July 2008
Accepted 19 July 2008

Keywords:
Atractor neural networks
Partial updating of neurons
Chaotic itinerancy among attractors

ABSTRACT

We present and study a probabilistic neural automaton in which the fraction of simultaneously–updated neurons is a parameter, $\rho \in (0, 1)$. For small ρ, there is relaxation towards one of the attractors and a great sensibility to external stimuli and, for $\rho \geq \rho_c$, itinerancy among attractors. Tuning ρ in this regime, oscillations may abruptly change from regular to chaotic and vice versa, which allows one to control the efficiency of the searching process. We argue on the similarity of the model behavior with recent observations, and on the possible role of chaos in neurobiology.

1. Introduction

Attractor neural networks (ANN) are a paradigm for the property of associative memory (Hopfield, 1982; Amit, 1989). Nevertheless, concerning practical applications, and also when trying to mold the essence of actual systems, the utility of ANN is severely limited, mainly by the fact that they can only retrieve one memory at a time. In this note, we show that such a limitation may be systematically overcome by simply generalizing familiar model situations. More specifically, we here extend some of our recent work on ANN with fast pre–synaptic noise (Cortes, Torres, Marro, Garrido, & Kappen, 2006; Marro, Torres, & Cortes, 2007; Torres, Cortes, Marro, & Kappen, 2007). The result is a novel mathematically–tractable ANN whose activity eventually describes heteroclinic paths among the attractors. This illustrates, in particular, the possibility of a constructive role of chaos during searching processes.

Our previous related studies essentially considered the same model system, but focussed on two different ways of updating it, namely, (i) sequential and (ii) parallel updating. Interesting enough, the ensuing behavior was qualitatively, even dramatically different. That is, the main observation was, respectively, (i) a great enhancement of the system sensibility to external stimuli as a consequence of rapid synaptic fluctuations which simulate facilitation and/or depression (Cortes et al., 2006; Torres et al., 2007), and (ii) chaotic behavior while the system spontaneously visited all the available attractors (Marro et al., 2007). Each of these two regimes of behavior is to be associated with a different functionality of an essential dynamic instability. Such an important dependence on the updating process is rather unexpected. For instance, we checked that it does not occur in a recent model (Pantic, Torres, Kappen, & Gielen, 2002; Pantic, Torres, & Kappen, 2003) which is based on a different depression mechanism. This situation motivated us to study in detail the changeover between (i) and (ii) as a modification of our previously proposed ANN (Cortes et al., 2006; Marro et al., 2007). That is, we here present neural automata in which the number or density ρ of neurons that are updated at each time step is a parameter. The resulting behavior as one modifies ρ is varied and intriguing. It leads us to argue on the possible relevance of our observations to interpret neurobiological experiments.

2. Definition of model

Let the sets of neuron activities $\boldsymbol{\sigma} \equiv \{\sigma_i\}$ and synaptic weights $\mathbf{w} \equiv \left\{w_{ij} \in \mathbb{R}\right\}$, where $i, j = 1, \dots, N$, and assume a presynaptic current $h_i(\boldsymbol{\sigma}, \mathbf{w})$ on each neuron due to the weighted action of the others. At each time unit, one updates the activity of n neurons, $1 \leqslant n \leqslant N$. This induces evolution in discrete time, t, of the state probability distribution according to

$$P_{t+1}(\boldsymbol{\sigma}) = \sum_{\boldsymbol{\sigma}'} R\left(\boldsymbol{\sigma}' \to \boldsymbol{\sigma}\right) P_t(\boldsymbol{\sigma}'), \qquad (1)$$

[☆] Contributed article.
[*] Corresponding author. Tel.: +34 958244014.
E-mail address: jtorres@onsager.ugr.es (J.J. Torres).

0893-6080/$ – see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.neunet.2008.07.002

where the transition rate is a superposition:

$$R\left(\boldsymbol{\sigma}\rightarrow\boldsymbol{\sigma}'\right)=\sum_{\mathbf{x}}p_n(\mathbf{x})\prod_{\{i|x_i=1\}}\tilde{\varphi}_n\left(\sigma_i\rightarrow\sigma_i'\right)\prod_{\{i|x_i=0\}}\delta_{\sigma_i,\sigma_i'}. \tag{2}$$

Here, $\tilde{\varphi}_n\left(\sigma_i\rightarrow\sigma_i'\right)\equiv\varphi\left(\sigma_i\rightarrow\sigma_i'\right)\left[1+\left(\delta_{\sigma_i',-\sigma_i}-1\right)\delta_{n,1}\right]$ and we denote $\mathbf{x}\equiv\{x_i=0,1\}$ an extra set of indexes which helps one in selecting the desired subset of neurons. The above thus describes *parallel* updating, as in familiar cellular automata (Chopard & Droz, 1998), for $n=N$ or, macroscopically, $\rho\equiv n/N\rightarrow 1$, while updating proceeds sequentially, as in kinetic Ising-like models (Marro & Dickman, 1999), for $n=1$ or $\rho\rightarrow 0$.

We shall consider explicitly the simplest version of this model which happens to be both interesting and mathematically tractable. First, we assume binary neurons, so that $\sigma_i=\pm 1$, which is known to be sufficient in order to capture the essentials of cooperative processes (Abbott & Kepler, 1990; Marro & Dickman, 1999; Pantic et al., 2002). The elementary rate φ is an arbitrary function of $\beta\sigma_i h_i$ (with β an inverse "temperature" or stochasticity parameter) which we assume to satisfy detailed balance. This property is not fulfilled by the superposition (2) for $n>1$, however. Consequently, the resulting steady states are generally out of equilibrium, which is more realistic in practice than thermodynamic equilibrium (Marro & Dickman, 1999). On the other hand, we shall only illustrate the case in which the n neurons are chosen at random out from the set of N, so that one has $p_n(\mathbf{x})=\binom{N}{n}^{-1}\delta\left(\sum_i x_i-n\right)$ in (2). For the sake of simplicity, we also need to assume that the currents are such that $h_i(\boldsymbol{\sigma},\mathbf{w})=h\left[\boldsymbol{\pi}(\boldsymbol{\sigma}),\boldsymbol{\xi}_i\right]$, where $\boldsymbol{\xi}_i\equiv\left\{\xi_i^\mu=\pm 1;\mu=1,\ldots,M\right\}$ are some given, *stored patterns* (realizations of the set of activities) and $\boldsymbol{\pi}\equiv\{\pi^\mu(\boldsymbol{\sigma})\}$. Here, $\pi^\mu(\boldsymbol{\sigma})=N^{-1}\sum_i\xi_i^\mu\sigma_i$ measures the *overlap* between the current state and pattern μ. For $N\rightarrow\infty$ and finite M, i.e., in the limit $\alpha\equiv M/N\rightarrow 0$ (which is not the interesting case, but may serve first for illustrative purposes) the resulting time equation under these conditions is $\pi_{t+1}^\mu(\boldsymbol{\sigma})=\rho N^{-1}\sum_i\xi_i^\mu\tanh\left(h_i^t\right)+(1-\rho)\pi_t^\mu(\boldsymbol{\sigma})$, where $h_i^t\equiv\beta h_i\left[\boldsymbol{\pi}_t(\boldsymbol{\sigma}),\boldsymbol{\xi}_i\right]$, for any μ. The above result is general and valid for any type of patterns. It is to be noticed that the sum over i in this map can be replaced by an average over the distribution of patterns $p(\xi_i^\mu)$. This permits a simple derivation of mean-field dynamical equations for the overlaps, at least for finite M. Note also that Monte Carlo simulations do not require restriction concerning the nature of the stored patterns.

The above allows for different relations between the currents h_i and the weights w_{ij}, and between these and other system properties. The simplest realization corresponds to the Hopfield case (Hopfield, 1982) which follows from the map above for $\rho\rightarrow 0$ and currents given by $h_i(\boldsymbol{\sigma},\mathbf{w})=\sum_{j\neq i}w_{ij}\sigma_j$ with the weights fixed according to the Hebb prescription, namely, $w_{ij}=N^{-1}\sum_\mu\xi_i^\mu\xi_j^\mu$. The symmetry $w_{ij}=w_{ji}$ then assures $P_{t\rightarrow\infty}(\boldsymbol{\sigma})\propto\exp\left(\beta\sum_i h_i\sigma_i\right)$ and, for high enough β, the *stored* patterns $\boldsymbol{\xi}$ are attractors of dynamics (Amit, 1989). We checked that, in agreement with some indications (Herz & Marcus, 1993), the Hopfield–Hebb network exhibits associative memory for any $\rho>0$. However, the situation is more complex, e.g., it depends on ρ, as one goes beyond Hopfield–Hebb, as we show in the next section.

It is well documented that transmission of information and computations in the brain are correlated with activity-induced fast fluctuations of synapses, i.e., our w_{ij}'s (Abbott & Regehr, 2004; Dobrunz & Stevens, 1997; Ferster, 1996). More specifically, it has been observed that there is some efficacy lost after heavy work, so that synapses suffer from *depression*; it is claimed that repeated activation decreases the neurotransmitter release which depresses the synaptic response (Abbott, Varela, Sen, & Nelson, 1997; Cook, Schwindt, Grande, & Spain, 2003; Tsodyks, Pawelzik, & Markram, 1998; Thomson & Deuchars, 1994; Thomson, Bannister, Mercer, & Morris, 2002). The consequences of this have already been analyzed in various contexts (Bibitchkov, Herrmann, & Geisel, 2002; Cook et al., 2003; Cortes et al., 2006; Marro et al., 2007; Pantic et al., 2002; Torres et al., 2007), and a main general conclusion from these studies is that depression importantly affects a network performance reducing, in particular, the stability of the attractors. Motivated by these facts, we shall adopt here the Hopfield currents and the following prescription for the synaptic weights:

$$w_{ij}=\left[1-(1-\Phi)\,q(\boldsymbol{\pi})\right]N^{-1}\sum\nolimits_\mu\xi_i^\mu\xi_j^\mu, \tag{3}$$

where $q(\boldsymbol{\pi})\equiv\frac{1}{1+\alpha}\sum_\mu\pi^\mu(\boldsymbol{\sigma})^2$. Note here that, in addition of static quenched disorder as in the standard Hopfield model, the weights (3) include a time dependence through the overlap vector $\boldsymbol{\pi}$ which is a measure of the network firing activity. These weights, which reduce to the Hebb prescription for $\Phi=1$, amount to assume short-term fluctuations which change synapses by a factor Φ on the average with a probability $q(\boldsymbol{\pi})$. Therefore, any positive $\Phi<1$ simulates synaptic depression if $q(\boldsymbol{\pi})$ is large. This is in agreement with the fact that, the greater π is, more activity will on average arrive to a particular postsynaptic neuron i in the network, and therefore, this neuron will be more depressed. Although the magnitude $q(\boldsymbol{\pi})$ involves a sum over all stored patterns, this will only affect neurons that are active in a particular pattern for not too high correlated patterns. More details concerning these assertions are in Cortes et al. (2006) and Marro et al. (2007).

Our setting here is rather close to the one in previous treatments of depressing synapses in a cooperative environment. As a matter of fact, one may show after some simple algebra that the model in Pantic et al. (2002, 2003) and Torres, Pantic, and Kappen (2002) corresponds to certain choices of Φ and $q(\boldsymbol{\pi})$ in (3) concerning steady states. For instance, a possible choice for $M=1$ and $\rho=1$ is $\Phi=1-\gamma/\gamma_0$ and

$$q(\pi)=\frac{\gamma_0[\gamma(1-\pi^2)+4]}{\gamma^2(1-\pi^2)+4\gamma+4} \tag{4}$$

where γ is the depression parameter defined in Torres et al. (2002) and γ_0 is the value for that parameter at which $\Phi=0$. This type of nonlinearity in $q(\pi)$, however, induces less sensitivity than the choice we are using here (see next section).

For the sake of completeness, we shall be concerned in this paper with both positive and negative values of Φ. A result is that the behavior we are looking for ensues in any of these cases (but only for certain values of Φ).

3. Some main results

In the limite $N\rightarrow\infty$ the (nonequilibrium) stationary state follows from the map for $M=1$ as $\pi_\infty=F(\pi_\infty;\rho,\Phi)$, and local stability requires that $|\partial F/\partial\pi|<1$; $F(\pi;\rho,\Phi)\equiv\rho\tanh\left\{\beta\pi\left[1-(1-\Phi)\pi^2\right]\right\}+(1-\rho)\pi$. The fixed point is therefore independent of ρ, but stability demands that $\rho<\rho_c$ with

$$\rho_c=2\left\{3\beta\pi_\infty^2\left[\left(\frac{4}{3}-\Phi\right)-(1-\Phi)\pi_\infty^2\right]-\beta+1\right\}^{-1} \tag{5}$$

(a condition that cannot be fulfilled in the Hopfield, $\Phi=1$ case). As Fig. 1 shows, $\rho=\rho_c$ marks the period-doubling route to chaos in the saddle–point map. This behavior is confirmed numerically for $M\gg 1$ stored arbitrary patterns, as shown numerically below.

Fig. 2 shows some typical *stationary* Monte Carlo runs, i.e., from bottom to top: (a) convergence towards one attractor – in fact,