

Instabilities in Attractor Networks with Fast Synaptic Fluctuations and Partial Updating of the Neurons Activity

J.J. Torres⁺, J. Marro⁺, J.M. Cortes^{+§*}, and B. Wemmenhove[§]

⁺Institute “Carlos I” for Theoretical and Computational Physics, and
Departamento de Electromagnetismo y Física de la Materia,
University of Granada, 18071 Granada, Spain

[§]Department of Biophysics and SNN, Radboud University,
6525 EZ Nijmegen, The Netherlands

^{*}Institute for Adaptive and Neural Computation, School of Informatics,
University of Edinburgh, EH1 2QL, UK.

July 15, 2008

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 the 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 et al., 2006; Torres et al., 2007; Marro et al., 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 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 et al., 2002; Pantic et al., 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 $\sigma \equiv \{\sigma_i\}$ and synaptic weights $\mathbf{w} \equiv \{w_{ij} \in \mathbb{R}\}$, where $i, j = 1, \dots, N$, and assume a presynaptic current $h_i(\sigma, \mathbf{w})$ on each neuron due to the weighted action of the others. At each time unit, one updates the activity of n neurons, $1 \leq n \leq N$. This induces evolution in discrete time, t , of the state probability distribution according to

$$P_{t+1}(\sigma) = \sum_{\sigma'} R(\sigma' \rightarrow \sigma) P_t(\sigma'), \quad (1)$$

where the transition rate is a superposition:

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

Here, $\tilde{\varphi}_n(\sigma_i \rightarrow \sigma'_i) \equiv \varphi(\sigma_i \rightarrow \sigma'_i) \left[1 + \left(\delta_{\sigma'_i, -\sigma_i} - 1 \right) \delta_{n,1} \right]$ 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 and Droz, 1998), for $n = N$ or, macroscopically, $\rho \equiv n/N \rightarrow 1$, while updating proceeds sequentially, as in kinetic Ising-like models (Marro and 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 (Pantic et al., 2002; Marro and Dickman, 1999; Abbott and Kepler, 1990). 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 and 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 $p_n(\mathbf{x}) = \binom{N}{n}^{-1} \delta(\sum_i x_i - n)$ in (2). For the sake of simplicity, we also need to assume that the currents are such that $h_i(\sigma, \mathbf{w}) = h[\pi(\sigma), \xi_i]$, where $\xi_i \equiv \{\xi_i^\mu = \pm 1; \mu = 1, \dots, M\}$ are some given, *stored patterns* (realizations of the set of activities) and $\pi \equiv \{\pi^\mu(\sigma)\}$. Here, $\pi^\mu(\sigma) = N^{-1} \sum_i \xi_i^\mu \sigma_i$ 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 $\pi_{t+1}^\mu(\sigma) = \rho N^{-1} \sum_i \xi_i^\mu \tanh(h_i^t) + (1 - \rho) \pi_t^\mu(\sigma)$, where $h_i^t \equiv \beta h_i[\pi_t(\sigma), \xi_i]$, 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 $h_i(\sigma, \mathbf{w}) = \sum_{j \neq i} w_{ij} \sigma_j$ 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 $P_{t \rightarrow \infty}(\sigma) \propto \exp(\beta \sum_i h_i \sigma_i)$ and, for high enough β , the *stored* patterns ξ are attractors of dynamics (Amit, 1989). We checked that, in agreement with some indications (Herz and 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 (Ferster, 1996; Dobrunz and Stevens, 1997; Abbott and Regehr, 2004). 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 (Tsodyks et al., 1998; Thomson and Deuchars, 1994; Abbott et al., 1997; Thomson et al., 2002; Cook et al., 2003). The consequences of this have already been analyzed in various contexts (Pantici et al., 2002; Cook et al., 2003; Bibitchkov et al., 2002; Cortes et al., 2006; Marro et al., 2007; 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:

$$w_{ij} = [1 - (1 - \Phi) q(\pi)] N^{-1} \sum_{\mu} \xi_i^{\mu} \xi_j^{\mu}, \quad (3)$$

where $q(\pi) \equiv \frac{1}{1+\alpha} \sum_{\mu} \pi^{\mu}(\sigma)^2$. 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 π 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 $q(\pi)$. Therefore, any positive $\Phi < 1$ simulates synaptic depression if $q(\pi)$ is large. This is in agreement with the fact that, the greater π is, more activity will in the average arrive to a particular postsynaptic neuron i in the network and, therefore, this neuron will be more depressed. Although the magnitude $q(\pi)$ 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; 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; Torres et al., 2002; Pantic et al., 2003) corresponds to certain choices of Φ and $q(\pi)$ in (3) concerning steady states. For instance, a possible choice for $M = 1$ and $\rho = 1$ is $\Phi = 1 - \gamma/\gamma_0$ and

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

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 $|\partial F/\partial \pi| < 1$; $F(\pi; \rho, \Phi) \equiv \rho \tanh\{\beta\pi[1 - (1 - \Phi)\pi^2]\} + (1 - \rho)\pi$. The fixed point is therefore independent of ρ , but stability demands that $\rho < \rho_c$ with

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

(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, one of the *antipatterns*, namely, the negative of one of the given patterns—for small ρ ; (b) fully irregular behavior with positive Lyapunov exponent for $\rho > \rho_c$; (c) regular oscillation between one attractor and its negative for $\rho > \rho_c$; (d) onset of chaos as ρ is further increased; and (e) rapid and ordered periodic oscillations between one pattern and its antipattern when, finally, all the neurons are active. The cases (b) and (d) are examples of instability-induced switching phenomena, in which the system activity chaotically *visits*

different attractors by describing heteroclinic paths and remaining different time intervals in the neighborhood of each attractor. This kind of behavior was previously observed for $\rho \rightarrow 1$ at certain values of Φ (Cortes et al., 2006). The interesting new facts are that this requires a minimum of synchronized neurons, that this minimum—as well as many other details—depend on Φ , and that, as we show in the caption of figure 1, varying ρ above $\rho_c(\Phi)$ seems to induce further intriguing qualitative changes.

It is also to be remarked that chaotic switching or itinerancy requires that the system is in a specially sensitive state first described in (Cortes et al., 2006; Torres et al., 2007). This is accomplished in the present case by means of the activity-dependent fast noise modelled in (3). One should expect that variations of this assumption on the weights may result in an equivalent sensitive state. As a matter of fact, we found that changing the sign of Φ does not affect our main observations. However, the case $\Phi = 1$, in which the weights are fixed, does not exhibit interesting behavior, and ρ turns then into an irrelevant parameter. On the other hand, the model in (Pantic et al., 2002; Torres et al., 2002; Pantic et al., 2003) does not seem to involve sufficient sensitivity for the purpose (see figure 3), in spite of the fact that it includes an activity-dependent depression mechanism. The explanation is the following. Assuming that the dynamics can be written as $\pi_{t+1} = G(\pi_t)$, the gain function $G(\pi)$ in the model in (Torres et al., 2002) is a nonlinear one which behaves monotonically for all values of the depression parameter. In our case, however, a non-monotonic type of gain function occurs for some values of Φ and ρ (see comparison in figure 4). This has been reported to be important to originate a chaotic dynamics among the attractors (Dominguez and Theumann, 1997; Caroppo et al., 1999).

Monitoring activity trajectories as one varies ρ in the case of several stored patterns provides the following qualitative picture for arbitrary patterns. As far as $\rho < \rho_c$, the activity remains wandering around one of the patterns. The pattern selected depends on the initial condition, and the trajectory visits a neighborhood of it whose volume increases slightly with ρ . The trajectory seems to tend to densely fill this volume with time. Increasing ρ , however, the system may escape from the initially chosen pattern and, eventually, will tend to visit all the patterns. In addition, one observes that the trajectory is rather structured. That is, there are many jumps between the more correlated patterns but only very few to the less correlated ones if the system is close to the edge of chaos, and the system attention to all the patterns tends to be balanced as ρ is increased within a chaotic window. Increasing ρ further, the network surpasses equiprobability of patterns and, eventually, abandons the chaotic regime to fall into a limit cycle, where pe-

riodically oscillates between a pattern and its antipattern. This confirms and details the behavior shown in figure 1.

This behavior, which is clearly observed in Monte Carlo simulations, can also be obtained under a mean field theory. Assume, for instance, random patterns with $p(\xi_i^\nu) = \frac{1+a}{2}\delta(\xi_i^\nu - 1) + \frac{1-a}{2}\delta(\xi_i^\nu + 1)$ where $\langle \xi^\nu \rangle = a$, even for $0 < |a| \ll 1$. In the simplest case of two patterns this mean field dynamics is determined by

$$\begin{aligned}\pi_{t+1}^1 &= \rho \frac{1+a^2}{2} \tanh[B(\pi_t)(\pi_t^1 + \pi_t^2)] + \rho \frac{1-a^2}{2} \tanh[B(\pi_t)(\pi_t^1 - \pi_t^2)] \\ &\quad + (1-\rho) \pi_t^1 \\ \pi_{t+1}^2 &= \rho \frac{1+a^2}{2} \tanh[B(\pi_t)(\pi_t^1 + \pi_t^2)] - \rho \frac{1-a^2}{2} \tanh[B(\pi_t)(\pi_t^1 - \pi_t^2)] \\ &\quad + (1-\rho) \pi_t^2,\end{aligned}\tag{6}$$

where $B(\pi) \equiv \beta[1 - (1 - \Phi)q(\pi)]$. It may be noticed that only in the non-interesting case of orthogonal patterns, namely $a = 0$, the mean field dynamics (6) gives chaotic switching between a particular pattern and its antipattern but not between different patterns. Otherwise, the situation is of chaotic switching among the stored patterns.

4 Discussion

This paper deals with ANN in which the density ρ of neurons that are updated at each time step is a parameter, so that the limit $\rho \rightarrow 0$ (1) corresponds to sequential (parallel) updating. Our main motivation is that previous studies of ANN in these limits revealed qualitatively different behavior, and that analysis in which the number of updated neurons is systematically varied are rare in the literature, e.g., (Herz and Marcus, 1993). It is worth to remark also that there are several arguments which suggest studying changes with ρ . One is simply the suspicion, born outside biology, that a network could perhaps like to maintain inert some of the nodes during operation, and not necessarily for economy but in order to gain efficiency. As a matter of fact, as one may get convinced by oneself by looking at our expressions for the currents h_i , hushing some of the nodes may be equivalent to modifying the wiring topology, and this is recognized as a method to enhance a network efficiency (Torres et al., 2004; Eguiluz et al., 2005).

More specifically within biology, one may notice that assuming cells that are stimulated only in the presence of a neuromodulator such as dopamine, ρ could stand for the fraction of neurons modulated each cycle. There is no input on the other $1 - \rho$, so that information from the previous state is maintained, which was argued to be a basis for working memories (Egorov et al., 2002; LeBeau et al., 2005). On the other hand, varying ρ may also be relevant to simulate various situations of persistent activity (Wagenaar et al., 2006), the observed variability of the neurons threshold (Azouz and Gray, 2000), and the possible existence of *silent neurons* (Olshausen and Field, 2004; Shoham et al., 2006), for instance.

The fact is that varying ρ in our model turns out to be very intriguing. However, ρ is relevant only if the network is enough *sensitive*. Such a condition occurs in our case as a consequence of activity-dependent *fast synaptic* noise as modelled in (3). The parameter ρ is irrelevant in other cases as, in particular, for the model in (Pantic et al., 2002; Pantic et al., 2003) which is based on the depression mechanism introduced in (Tsodyks et al., 1998), and also when the synaptic weights are fixed, even heterogeneously as in a Hopfield–Hebb network. On the contrary, the model here exhibits kind of *dynamic* association, namely, the activity either goes to one attractor or else, for large enough ρ , visits possible attractors. The visits may abruptly become chaotic. Besides synchronization of a minimum of neurons, this requires careful tuning of ρ . As a matter of fact, as shown by equation (5) and figure 1, a complex situation makes it difficult to predict the result for slight changes of ρ .

Another interesting feature of our model is illustrated in figure 5. This shows time series of the mean firing rate, $m = \frac{1}{2N} \sum_i (1 + \sigma_i)$, in a case study with six patterns exposed to two different stimuli of the same intensity and duration (between 3000 and 4000 n Monte Carlo trials). Each pattern is a string of N bits. Three patterns are randomly generated with 40, 50 and 60% of the bits set to 1, and the other three with the 1s at the first 70, 50 and 25% positions, respectively; the rest of the bits are set to -1 . The bottom graph shows the baseline activity without stimulus (BS) and the activity level under stimulus $\mu = 1$ (SA1) and $\mu = 2$ (SA2), i.e., two of the patterns. The behavior which exhibits the system in this case (which we found for other parameter values as well) is amazingly alike to observations in a comparable (but true, not computer) experimental setting concerning the odor response of the projection neurons in the locust antennal lobe (Rabinovich et al., 2001; Mazor and Laurent, 2005).

Interesting enough, the switching which shows our model due to stimulus destabilization in the simulation of figure 5 occurs for $\rho < \rho_c$. In fact, a

similar phenomenon was observed also for $\rho \rightarrow 0$ (Cortes et al., 2006). This shows that, at least in this case, an efficient adaptation to a changing environment does not require chaos. However, the chaotic itinerancy we described above allows for a more efficient search of the attractors space in a way that was believed to hold in relevant systems under a critical condition (Chialvo, 2006). Our model thus illustrates a mechanism that makes chaos extremely beneficial. This confirms expectations (Korn and Faure, 2003; Glass, 2002; Ashwin and Timme, 2005) that the instability inherent to chaos facilitates moving to any pattern at any time. The present model system illustrates a specific mechanism which allows for this. As ρ increases in a chaotic region, it is more likely that the activity will visit all the attractors, not only the most correlated ones. The number and diversity of attractors it visits then increases with ρ , and we observed that the time spent in the attractor also varies with ρ . The system in this way may perform family discrimination and classification by tuning ρ . We finally remark that our model allows for describing a coupling of ρ to the activity, which may be quite a realistic setting in some cases. No doubt it would be interesting to study other related model situations.

We thank I. Erchova, P.L. Garrido and H.J. Kappen for very useful comments, and financial support from FEDER–MEC project FIS2005-00791, JA projects P06–FQM–01505 and P07–FQM–02725, and EPSRC–COLAMN project EP/CO 10841/1.

References

- Abbott, L. F. and Kepler, T. B. (1990). Model neurons: From hodgkin-huxley to hopfield. *Lecture Notes in Physics*, 368:5–18.
- Abbott, L. F. and Regehr, W. G. (2004). Synaptic computation. *Nature*, 431:796–803.
- Abbott, L. F., Varela, J. A., Sen, K., and Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275:220–224.
- Amit, D. J. (1989). *Modeling brain function: The world of attractor neural networks*. Cambridge University Press.
- Ashwin, P. and Timme, M. (2005). Nonlinear dynamics: when instability makes sense. *Nature*, 436:36–37.
- Azouz, R. and Gray, C. M. (2000). Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proc. Natl. Acad. Sci. USA*, 97:8110–8115.
- Bibitchkov, D., Herrmann, J. M., and Geisel, T. (2002). Pattern storage and processing in attractor networks with short-time synaptic dynamics. *Network: Comput. Neural Syst.*, 13:115–129.
- Caroppo, D., Mannarelli, M., Nardulli, G., and Stramaglia, S. (1999). Chaos in neural networks with a nonmonotonic transfer function. *Phys. Rev. E*, 60:2186–2192.
- Chialvo, D. R. (2006). Are our senses critical? *Nature Phys.*, 2:301–302.
- Chopard, B. and Droz, M. (1998). *Cellular Automata Modeling of Physical Systems*. Cambridge University Press.
- Cook, D. L., Schwindt, P. C., Grande, L. A., and Spain, W. J. (2003). Synaptic depression in the localization of sound. *Nature*, 421:66–70.
- Cortes, J. M., Torres, J. J., Marro, J., Garrido, P. L., and Kappen, H. J. (2006). Effects of Fast Presynaptic Noise in Attractor Neural Networks. *Neural Comp.*, 18:614–633.
- Dobrunz, I. E. and Stevens, C. F. (1997). Heterogeneity of release probability, facilitation and depletion at central synapses. *Neuron*, 18:995–1008.

- Dominguez, R. R. C. and Theumann, W. (1997). Generalization and chaos in a layered neural network. *J. Phys. A: Math. Gen.*, 30:1403-1414.
- Eckmann, J. P. and Ruelle, D. (1985). Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.*, 57:617-656.
- Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., and Alonso, A. A. (2002). Graded persistent activity in entorhinal cortex neurons. *Nature*, 420:173-178.
- Eguiluz, V., Cecchi, G., Chialvo, D. R., Baliki, M., Apkarian A. V. (2005). Scale-free brain functional networks. *Phys. Rev. Lett.*, 94:018102.
- Ferster, D. (1996). Is neural noise just a nuisance? *Science*, 273:1812-1815.
- Glass, L. (2002). Chaos in biological systems. In Arbib, M. A., editor, *Handbook of Brain Theory and Neural Networks*, page 205. MIT Press.
- Herz, A. V. M. and Marcus, C. M. (1993). Distributed dynamics in neural networks. *Phys. Rev. E*, 47:2155-2161.
- Hopfield, J. J. (1982). Neural Networks and Physical Systems with Emergent Collective Computational Abilities. *Proc. Natl. Acad. Sci. USA*, 79:2554-2558.
- Korn, H. and Faure, P. (2003). Is there chaos in the brain? ii. experimental evidence and related models. *C. R. Biologies*, 326:787-840.
- LeBeau, F. E., Manira, A. E., and Griller, S. (2005). Tuning the network: Modulation of neuronal microcircuits in the spinal cord and hippocampus. *Trends Neurosci.*, 28:552-561.
- Marro, J. and Dickman, R. (1999). *Nonequilibrium phase transitions in lattice models*. Cambridge University Press.
- Marro, J., Torres, J. J., and Cortes, J. M. (2007). Chaotic hopping between attractors in neural automata. *Neural Networks*, 20:230-235.
- Mazor, O. and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron*, 48:661-673.
- Olshausen, B. A. and Field, D. J. (2004). Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.*, 14:481-487.

- Pantic, L., Torres, J., and Kappen, H. (2003). Coincidence detection with dynamic synapses. *Network: Comput. Neural Syst.*, 14:17–33.
- Pantic, L., Torres, J. J., Kappen, H. J., and Gielen, S. C. A. M. (2002). Associative Memory with Dynamic Synapses. *Neural Comp.*, 14:2903–2923.
- Rabinovich, M., Volkovskii, A., Lecanda, P., Huerta, R., Abarbanel, H., and Laurent, G. (2001). Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Phys. Rev. Lett.*, 87:068102.
- Shoham, S., O’Connor, D. H., and Segev, R. (2006). How silent is the brain: Is there a *Dark Matter* problem in neuroscience? *J. Compar. Physiol. A*, 192:777–784.
- Thomson, A. M., Bannister, A., Mercer, A., and Morris, O. (2002). Target and temporal pattern selection at neocortical synapses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 357:1781–1791.
- Thomson, A. M. and Deuchars, J. (1994). Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci.*, 17:119–126.
- Torres, J. J., Cortes, J. M., Marro, J., and Kappen, H. J. (2007). Competition between synaptic depression and facilitation in attractor neural networks. *Neural Comp.*, 19: 2739–2755.
- Torres, J. J., Muñoz, M. A., Marro, J., and Garrido, P. L. (2004). Influence of topology on the performance of a neural network. *Neurocomputing*, 58:229–234.
- Torres, J. J., Pantic, L., and Kappen, H. J. (2002). Storage capacity of attractor neural networks with depressing synapses. *Phys. Rev. E*, 66: 061910.
- Tsodyks, M., Pawelzik, K., and Markram, H. (1998). Neural networks with dynamic synapses. *Neural Comp.*, 10:821–835.
- Wagenaar, D. A., Nadasdy, Z., and Potter, S. M. (2006). Persistent dynamic attractors in activity patterns of cultured neuronal networks. *Phys. Rev. E*, 73:051907.

Figure Captions

Figure 1: The Lyapunov exponent (solid curve), showing transitions from regular ($\lambda < 0$) to chaotic ($\lambda > 0$) as the *synchronization parameter* $\rho = n/N$ is varied, as obtained analytically from the saddle-point solution for $\Phi = 0.005$, $M = 1$ patterns, and $\beta = 50$. The chaotic windows here were precisely confirmed using related Monte Carlo simulations with $N = 3600$ neurons. The minimum fraction of active neurons needed to start the period-doubling route to chaotic behavior, ρ_c , is shown. This picture is strongly dependent on Φ ; there is a rather broad range of Φ values, including negative ones, for which the behavior is qualitatively similar. The dashed curve is the Hopfield–Hebb case $\Phi = 1$. The inset details the interesting region showing chaotic behavior.

Figure 2: The overlap as a function of time (in units of n Monte Carlo trials) after $t = 1920$, for $N = 1600$, $\beta = 20$, $\Phi = -0.4$, $M = 3$ uncorrelated patterns and, from bottom to top, $\rho = 0.08, 0.50, 0.65, 0.92$ and 1.00 , respectively. In this case, $\rho_c = 0.085$.

Figure 3: Time variation of the mean firing rate $m \equiv (1 + \pi)/2$ in an attractor neural network which stores a single pattern with depressing synapses, as modeled in (Tsodyks et al., 1998; Pantic et al., 2002), under partial updating in the oscillatory regime. Panels show, from top to bottom, the cases $\rho = 1, 0.7, 0.3, 0.1$. This (which corresponds to certain model parameters) reveals that, except for scaling of the typical temporal scale for the oscillations, partial updating does not introduce new phenomenology in this model, contrary to the case presented in this paper.

Figure 4: This compares the gain function in the model in this paper, for $\rho = 1$ and varying Φ (left panel) and the gain function in the model in (Torres et al., 2002) for varying γ (right panel). In both cases β was set to 3. Different curves in the left case are for $\Phi = 1$ (non-depressed case), 0.6, 0.2 and 0 (high depression); the curves in the right case occur when the corresponding parameter $\gamma = 0$ (non-depressed case), 0.5, 3, 10 (high depression case) This shows how the gain function can be non-monotonic for some values of the depression parameter Φ in the model in this paper. This allows for non-zero fixed point solutions, namely, the points that intersect the diagonal, with negative slopes (whose absolute value is larger than one) which leads to a period-doubling route to chaos.

Figure 5: Itinerancy induced by external stimuli. Mean firing rates as a function of time (bottom) and *phase-space* trajectories (top) trying to recreate an experimental observation concerning odor responses (Mazor and Laurent, 2005). The graphs show two Monte Carlo simulations of our system

with $N = 1600$, $\beta = 4$, $\Phi = -0.45$, $\rho = 3/64 < \rho_c$, and six stored patterns, for different stimuli, corresponding to green and red colors, respectively. The top graph involves a standard false-neighbor method (Eckmann and Ruelle, 1985) with *embedding dimension* $d_e = 5$, and the time delay is $\tau = 20$.

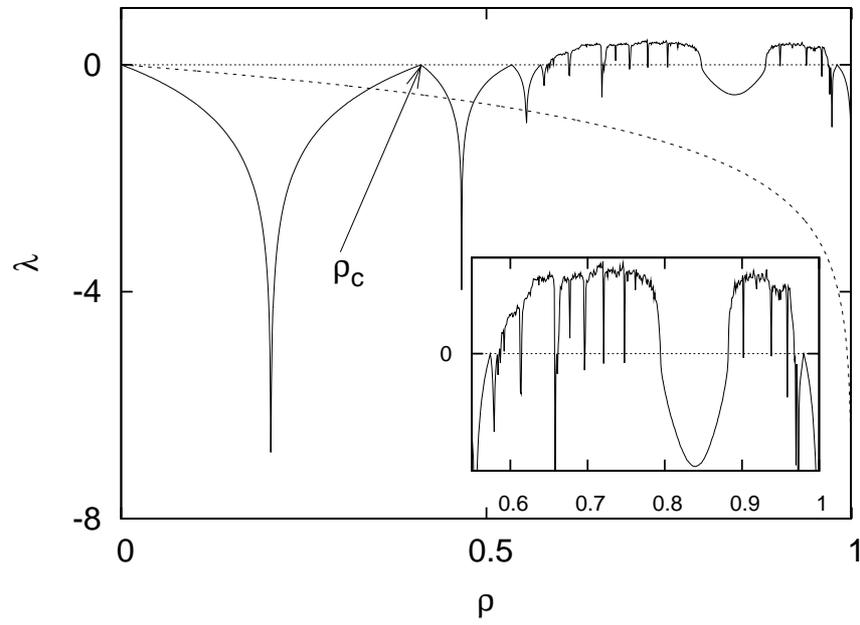


Figure 1: Torres et al.

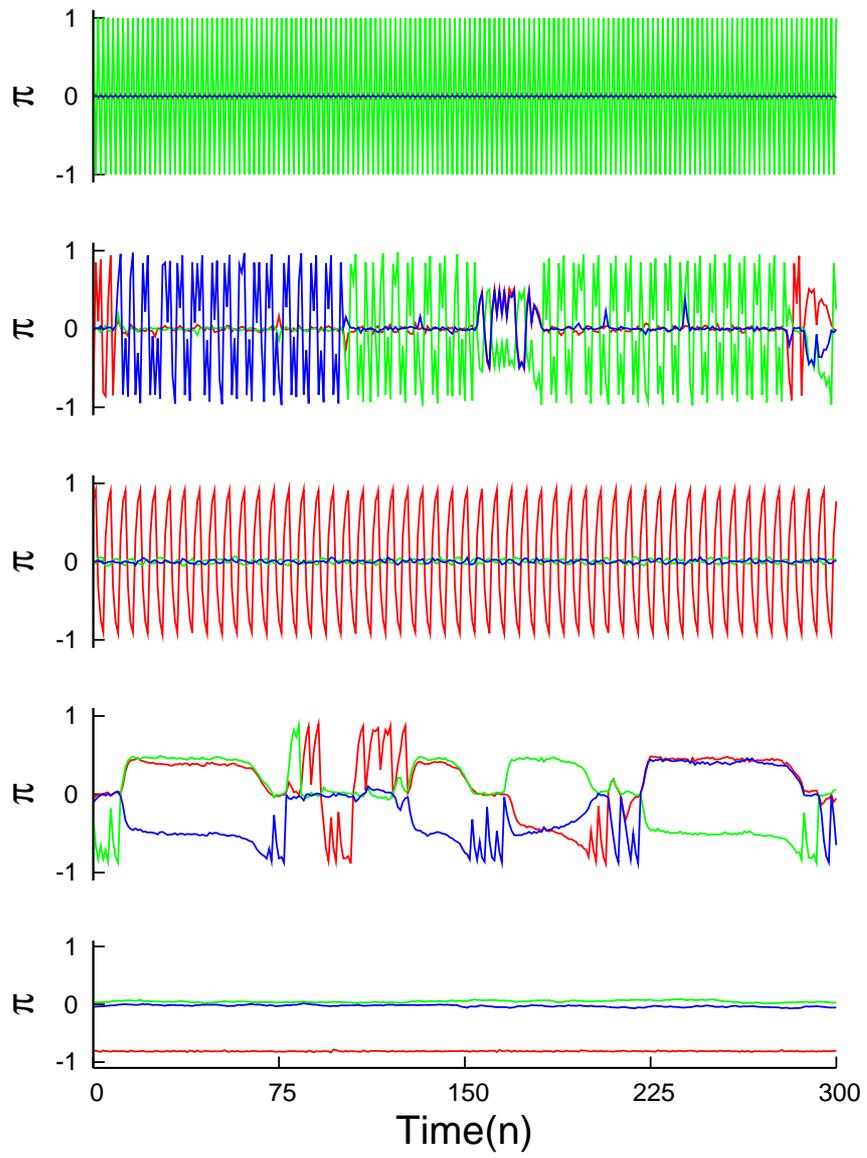


Figure 2: Torres et al.

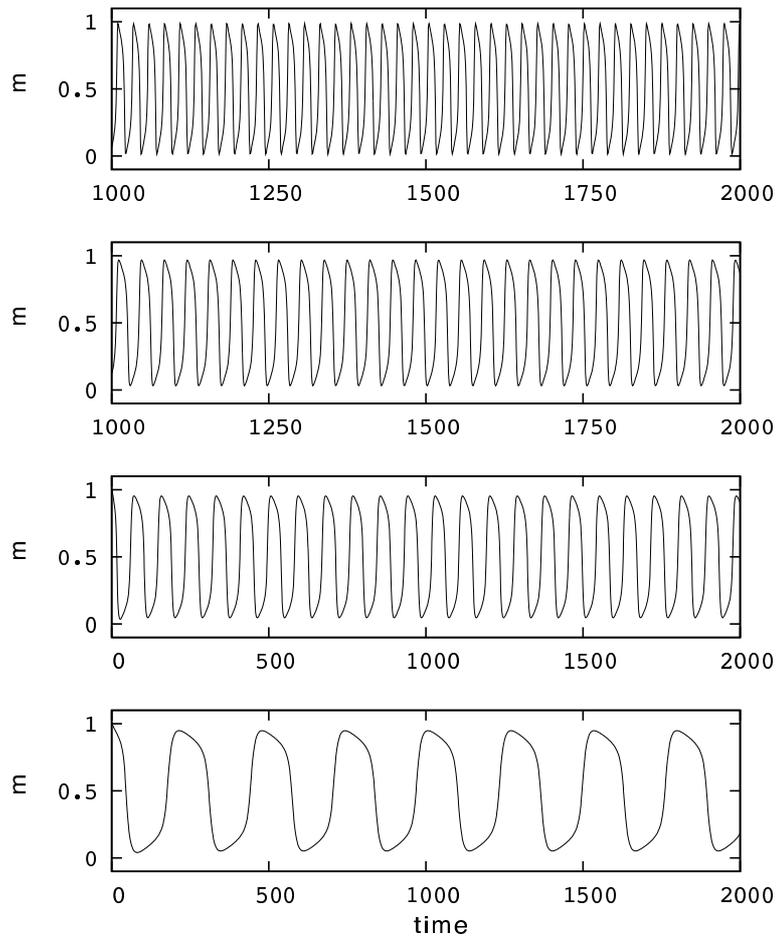


Figure 3: Torres et al.

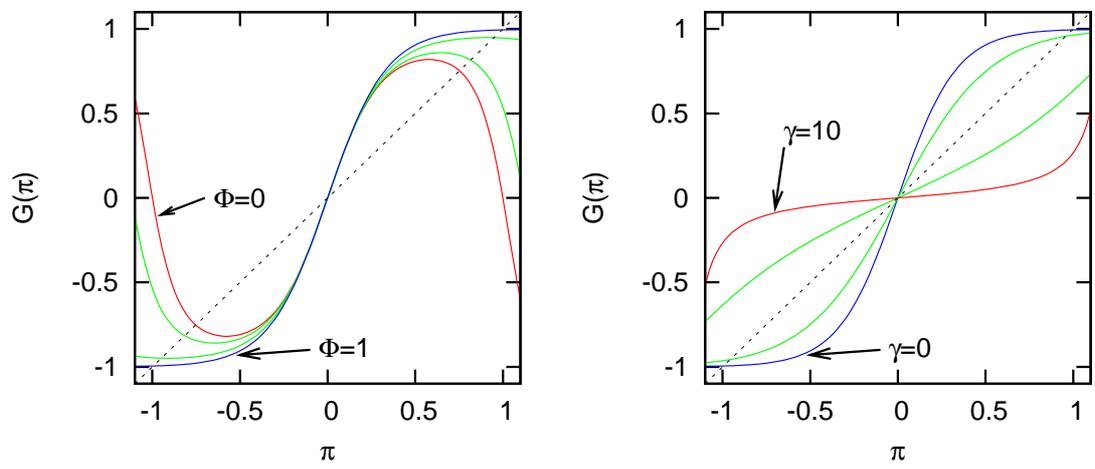


Figure 4: Torres et al.

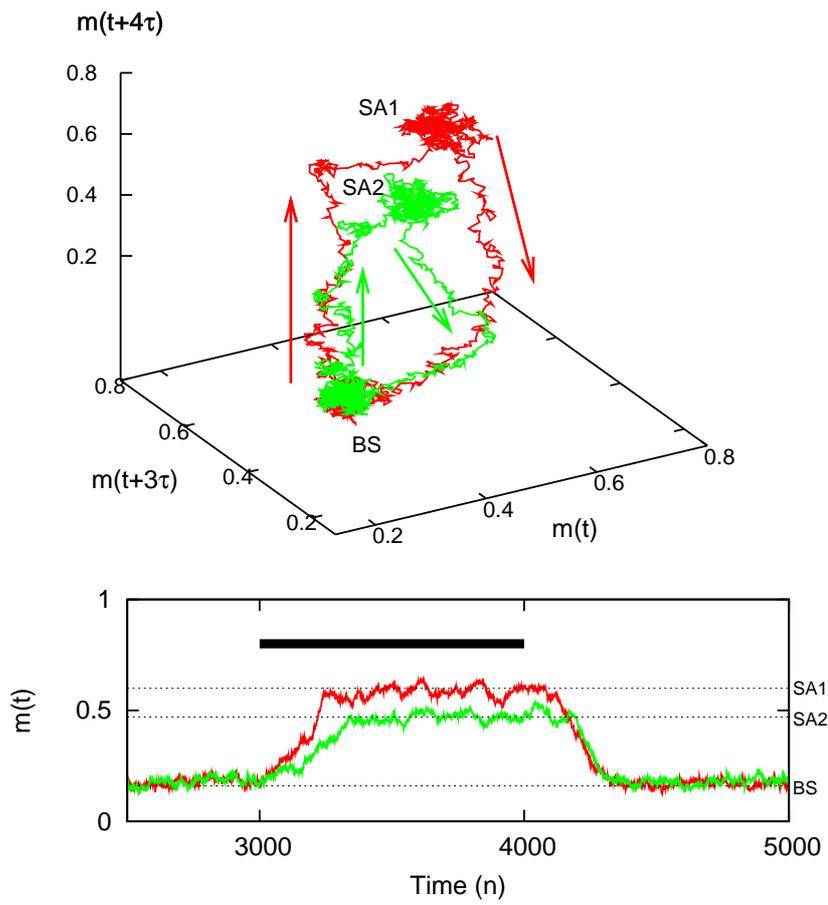


Figure 5: Torres et al.