Order-disorder transition in the Chialvo-Bak 'minibrain' controlled by network geometry

Joseph Wakeling

Institut de Physique Théorique, Université de Fribourg Pérolles
CH-1700 Fribourg, Switzerland

Abstract

We examine a simple biologically-motivated neural network, the three-layer version of the Chialvo-Bak 'minibrain' [Neurosci. 90 (1999) 1137], and present numerical results which indicate that a non-equilibrium phase transition between ordered and disordered phases occurs subject to the tuning of a control parameter. Scale-free behaviour is observed at the critical point. Notably, the transition here is due solely to network geometry and not any noise factor. The phase of the network is thus a design parameter which can be tuned. The phases are determined by differing levels of interference between active paths in the network and the consequent accidental destruction of good paths.

Key words: Phase transitions, Neural networks, Neuroscience
PACS: 05.70.Fh, 84.35.+i, 87.19.La

1 Introduction

Biological neuronal systems are certainly complex—but the source of this complexity still remains unclear. Bearing in mind the many examples of complex behaviour arising from simple rules [1–6], a number of authors [6–14] have suggested that some such basic mechanics might be responsible, and a number of neural models have been proposed [8–14] which aim to show that even a very simple interpretation of biological neural dynamics can nevertheless produce complex patterns of learning behaviour. Perhaps the most appealing is that created by Chialvo and Bak [13, 14]: rather cheekily dubbed the

Email address: `joseph.wakeling@unifr.ch` (Joseph Wakeling).
URL: `http://neuro.webdrake.net/` (Joseph Wakeling).

Preprint submitted to Physica A *6 February 2003*

'minibrain' [15], it is notable for its combined use of extremal ('winner-takes-all') dynamics [5, 16] and solely negative feedback to create a fast and highly adaptive learning system.

In this paper we examine the learning patterns of the simplest version of this model and present numerical results which indicate that, subject to tuning of a control parameter, a non-equilibrium transition is observed in the learning dynamics between ordered and disordered phases, similar to that observed in traditional Hopfield neural networks [17–23]; scale-free behaviour is observed at the critical point. Notably, the control parameter in the minibrain network is the network geometry: the network phase is therefore *a design parameter of the system*. Although a complete analytic understanding of this phenomenon is still to be developed, the underlying mechanism behind this behaviour is identified as being interference between active paths in the network. The behaviour at the critical point can also be compared to Kohonen's self-organizing map (SOM) [16], where a small stimulus can also lead to scale-free modification of the network structure [24].

2 The Chialvo-Bak 'minibrain'

Although a variety of alternative network topologies are possible [13, 14], the simplest version of the minibrain consists of three separate layers of neurons (see Figure 1): a layer of η_{IP} inputs, a 'hidden' intermediary layer of size η_{IM}, and a layer of η_{OP} output neurons. Each input neuron has a (one-way) synaptic connection to every intermediary neuron, and similarly each intermediary neuron is connected to every output neuron; we will use the terminology

$$\Gamma \equiv \Gamma(\eta_{\mathrm{IP}}, \eta_{\mathrm{IM}}, \eta_{\mathrm{OP}}) \tag{1}$$

to denote such a network geometry.

Each synaptic connection is assigned a strength value, initially determined randomly: when an input neuron is stimulated, a signal is sent along the single strongest connection leading from it, and the intermediary neuron at the end of that connection then fires. In the same manner, the output neuron that fires is the one at the end of the strongest connection leading from the firing intermediary neuron. These *extremal* ('winner-takes-all') dynamics, inspired both by earlier SOC models [5] and the SOM [16], mean that we can identify an *active level* of synaptic connections, the strongest connections leading from each neuron, that are collectively responsible for all decisions taken by the system. As we will see later, the active level is an extremely important aspect of the minibrain network.

Chialvo and Bak assume that synaptic activity leaves behind a tag of some

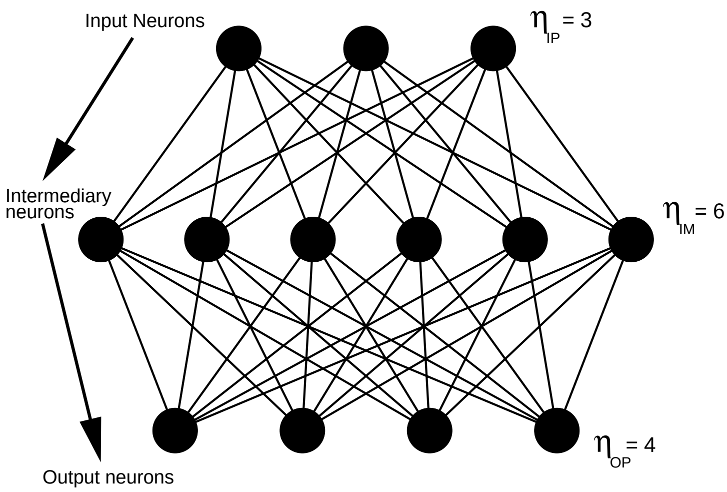

Fig. 1. A three-layer minibrain network with $\eta_{\text{IP}} = 3$ input neurons, $\eta_{\text{IM}} = 6$ intermediary neurons and $\eta_{\text{OP}} = 4$ output neurons. Each input neuron has a (one-way) synaptic connection to every intermediary neuron, and each intermediary neuron has a similar connection to every output neuron. Using the shorthand of Equation (1), this network would be denoted by $\Gamma(3, 6, 4)$.

sort (biologically, perhaps a chemical trace) which identifies that synapse as 'used'; feedback then takes the form of a global signal sent to the system (perhaps the release of some hormone) which reacts only with the tagged synapses. Following the lead of Barto, Jordan et al. [8–10], this global signal is extremely simple, indicating no more than whether the action taken has failed or succeeded. What next? If success is indicated, no further action is taken. *There is no strengthening of 'good' synapses.* If on the other hand the feedback signal is one of failure, the tagged synapses are 'punished' by having their strengths reduced by a random amount.

This selectionist interpretation of biological learning dynamics may be controversial because, in particular, it clashes with the traditional Hebbian learning [7], where good connections are strengthened—and this has traditionally had prime place among theories of biological learning. Chialvo and Bak cite two main reasons for their alternative viewpoint [13, 14]: firstly, *reinforcement of connections that produce good responses leads to addiction* unless some mechanism exists to stop the reinforcement once a correct response has been learned [25], whereas punishment of connections that produce bad responses by definition stops as soon as a good response is achieved; secondly, *learning by mistakes is quicker*, as a system in an unknown situation is more likely to make mistakes than get things right. That selectionism and negative feedback may in fact be the primary biological learning mechanism has also been emphasized by a number of other authors in works both experimental [26, 27] and theoretical [28, 29]. A detailed account of the biological motivation for the model can be found in [13, 14].

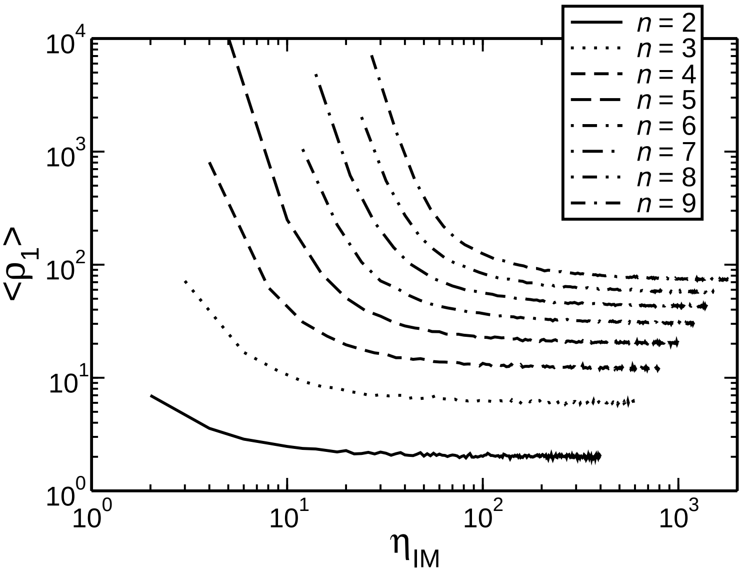

Fig. 2. Mean number, $\langle \rho_1 \rangle$, of applications of negative feedback required to learn an input-output map of size n, plotted as a function of the number $\eta_{\rm IM}$ of intermediary neurons. Data is averaged over 2048 realizations.

The minibrain has already been shown to have interesting learning behaviour [13–15], in particular being able to learn nonlinear functions such as the XOR problem and its generalization, the parity problem. Also, and importantly, it is a highly *adaptive* model, able to quickly unlearn patterns of behaviour that are no longer successful—a direct result of the negative-only feedback [13, 14, 30].

3 Control parameter for learning and order-disorder phase transition

Arguably the simplest task one can set such a system is the learning of an input-output map (e.g., $1 \to 1$, $2 \to 2$, ..., $n \to n$): sequentially stimulating each input neuron and, if the output is unsatisfactory, applying negative feedback as described above. Chialvo and Bak showed, although without explicit quantification, that learning time for such a map (defined as the number of times the network has to run through a complete cycle of inputs before complete success is achieved) decreases as the number $\eta_{\rm IM}$ of intermediary neurons increases [13]. However, 'learning time' cannot really be thought of as a physical quantity within the minibrain network itself, and hence in this paper we will use an alternative (though related) measure of learning efficiency: the total number of times, ρ_1, that negative feedback must be applied to the system—or equivalently, the number of errors made—before we can run through the complete sequence of inputs without any mistakes. This measure can be thought of as analogous to the measure of neural network learning efficiency used in early work on the subject [17–23] using the traditional Hopfield neural networks: the

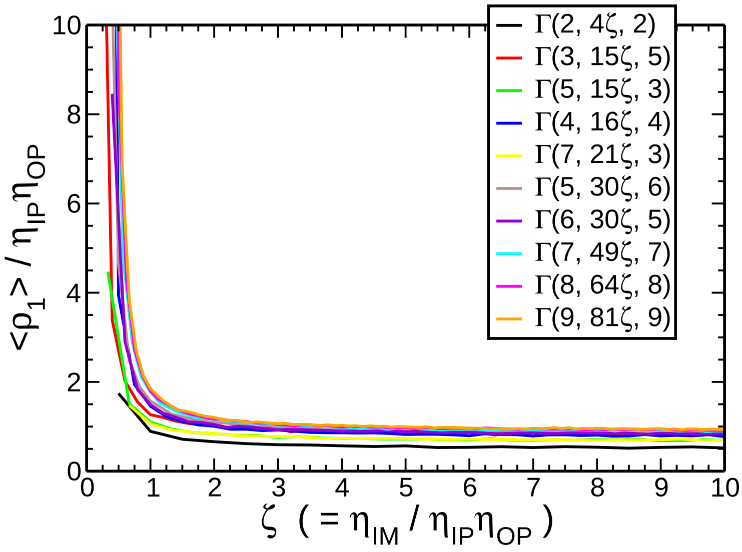

Fig. 3. Mean number, $\langle \rho_1 \rangle$, of applications of negative feedback (in $\eta_{\rm IP}\eta_{\rm OP}$ unit) required to learn input-output maps of varying sizes, plotted as a function of control parameter $\zeta = \eta_{\rm IM}/\eta_{\rm IP}\eta_{\rm OP}$. Note the data collapse as a result of this normalization, and the sharp transition between large and small (asymptotic minimum) values that occurs at $\zeta \approx 1$. The data legend uses the notation of Equation (1). Data is averaged over 2048 realizations.

number of examples the network must see in order to learn a general function.

Fig. 2 shows that in learning an input-output map of the above type (with $\eta_{\rm IP} = \eta_{\rm OP} = n$), the mean number, $\langle \rho_1 \rangle$, of applications of negative feedback required scales qualitatively in a similar manner to Chialvo and Bak's learning time: as the number of intermediary neurons $\eta_{\rm IM}$ increases, $\langle \rho_1 \rangle$ decreases towards an eventual asymptotic minimum whose value appears to depend on n. This behaviour is quantified precisely in Fig. 3, where we investigate the effects of unequal mappings where $\eta_{\rm IP} \neq \eta_{\rm OP}$, e.g., $i \rightarrow i \bmod \eta_{\rm OP}$: both the minimum value of $\langle \rho_1 \rangle$ and the decrease in negative feedback that occurs with increasing $\eta_{\rm IM}$ collapse when $\langle \rho_1 \rangle$ and $\eta_{\rm IM}$ are normalized with respect to $\eta_{\rm IP}\eta_{\rm OP}$, and hence we can define a control parameter for learning time,

$$\zeta = \frac{\eta_{\rm IM}}{\eta_{\rm IP}\eta_{\rm OP}} \tag{2}$$

It can be seen that at $\zeta \approx 1$ a sharp transition occurs in $\langle \rho_1 \rangle$ between the asymptotic minimum and truly immense values. Numerical results (Fig. 4) suggest that this is a genuine phase transition: if once the minibrain has learned a given input-output mapping, we randomly reassign a new output to one of the inputs, and repeat this process, we find that the probability function of negative feedback ρ_1 required to adapt to the new map varies between an ordered phase ($\zeta > 1$) where there is a maximum cutoff to learning times, with ρ_1 distributed exponentially, and a disordered phase ($\zeta < 1$) where

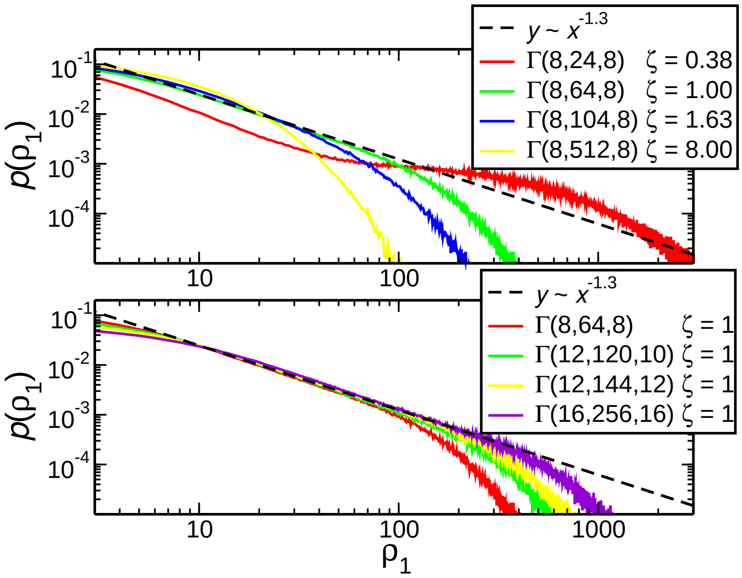

Fig. 4. Probability function $p(\rho_1)$ of the number of applications of negative feedback required to adapt to a new input-output map, based on the control parameter ζ: (a) different distributions of learning time for a network with fixed $\eta_{\rm IP} = \eta_{\rm OP} = 8$ and varying $\eta_{\rm IM}$, with power-law behaviour at the critical point $\zeta = \eta_{\rm IM}/\eta_{\rm IP}\eta_{\rm OP} = 1$; (b) power-law (scale-free) behaviour at the critical point for different system sizes, with $p(\rho_1) \sim \rho_1^{-1.3}$. Data taken over 10^6 adaptions.

learning is almost impossible, with ρ_1 strongly biased towards huge values. The probability function of ρ_1 becomes scale-free at the critical point $\zeta = 1$, scaling as a power law with exponent -1.3.

To understand what this means physically for the network, we have to recall what the meaning of negative feedback is. Decreasing the strength of a synaptic connection almost certainly removes it from the active level of synapses; another connection replaces it. A large number of negative feedback signals therefore corresponds to a large change in the geometry of the active level of the network—an 'avalanche' in the network landscape, so to speak. But what is it that causes such large avalanches?

The answer lies in *path interference*, the *overlapping of the active paths* leading from two or more distinct inputs. In the three-layer minibrain, this is as simple as the strongest connections from two different inputs pointing to the same intermediary neuron. Broadly speaking, if the paths from two or more neurons overlap in this fashion, it is likely that (at least) one of the inputs is connected to the wrong output neuron; we can make this explicitly true if we require that the input-output map to be learned be one-to-one. The negative feedback, however, will affect *both* paths by removing the intermediary-output connection from the active level. Thus, a correct wiring of input to output can be destroyed by such interference.

A brief statistical analysis reveals the behaviour of this factor. Suppose we have an arbitrary minibrain network, $\Gamma(\eta_{\rm IP}, \eta_{\rm IM}, \eta_{\rm OP})$. Labeling the input neurons by $0, 1, 2, \ldots, \eta_{\rm IP} - 1$, let us define, for a given input i, $\mu(i)$ to be the intermediary neuron led to by the strongest synaptic connection from i, and let p_i^Γ be the probability that at least two input neurons in the set $\{0, \ldots, i\}$ share the same value of μ. Thus, the probability of having path interference in the network is given by $P_\Gamma = p^\Gamma_{\eta_{\rm IP}-1}$. Assuming that the values of μ are randomly distributed, it follows that $p_0^\Gamma = 0$ and $p_1^\Gamma = 1/\eta_{\rm IM}$, and more generally (assuming $\eta_{\rm IM} \geq \eta_{\rm IP}$, since obviously $P_\Gamma = 1$ if this is not the case),

$$p^\Gamma_{i+1} = p^\Gamma_i + \frac{i+1}{\eta_{\rm IM}}(1 - p^\Gamma_i) \tag{3}$$

Now, substituting for $\eta_{\rm IM}$ in terms of the control parameter ζ, we have:

$$p^\Gamma_{i+1} = p^\Gamma_i + \frac{i+1}{\eta_{\rm IP}\eta_{\rm OP}\zeta}(1 - p^\Gamma_i) \tag{4}$$

From which we can calculate,

$$p^\Gamma_n = \sum_{i=1}^{n} \left[\binom{n}{i} (-1)^{i-1} \left(\eta_{\rm IP}\eta_{\rm OP}\right)^{-i} \zeta^{-i} \right] \tag{5}$$

Assuming large, fixed $\eta_{\rm IP}$ and $\eta_{\rm OP}$, the higher-order terms in Equation (5) will vanish for $\zeta > 1$, and so P_Γ will scale approximately according to a power law of ζ, with exponent -1. Conversely, if $\zeta < 1$ then the higher-order terms are brought back into play and P_Γ scales approximately exponentially; Fig. 5 shows nicely how the critical point $\zeta = 1$ acts as a transition point between these two scalings. The phase transition can thus be seen as a transition between a high-interference phase where learned patterns are frequently destroyed by accident as a result of mistakes elsewhere, and a low-interference phase where active connections remain largely stable.

It is interesting to compare this behaviour with the learning patterns observed in more traditional neural models, for example the large body of work on Hopfield networks [17–23]. Typically the performance of an N-input network in learning a function is studied with respect to the number P of examples presented to it (analogous to the parameter ρ_1 in this work, which is, recall, the number of times the minibrain must be shown where it is wrong before the error rate decreases to zero). A transition between ordered and disordered phases, similar to that observed here, was observed by Amit et al. [18], occurring as a result of the temperature of the system; however, the behaviour we see in the minibrain is more reminiscent of the behaviour observed by Gardner [19], who showed that, for fixed temperature, the phase transition

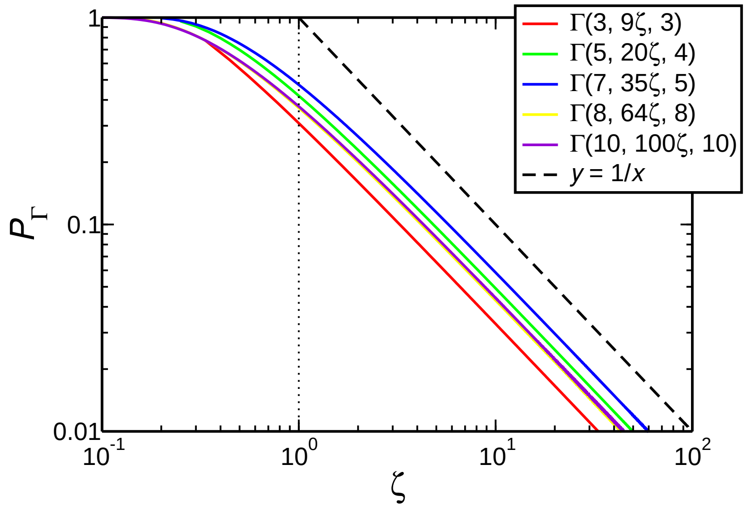

Fig. 5. Probability P_Γ of finding path interference in various networks $\Gamma(\eta_{\mathrm{IP}}, \eta_{\mathrm{IM}}, \eta_{\mathrm{OP}})$, plotted as a function of control parameter ζ. Note the transition at $\zeta \approx 1$ between exponential and power-law behaviour. Values numerically calculated from Equation (4).

in the Hopfield model occurred subject to the connectivity of neurons. In the minibrain, the path interference driving the phase transition is similarly a result solely of the network geometry—but in addition the minibrain is noise- (temperature-) resistant [14]. The network phase can therefore be seen as robust with respect to the environment.

4 Conclusions

While only a very simple interpretation of real neural dynamics, we have seen that the minibrain displays an interesting variety of behaviour, with a non-equilibrium transition between ordered and disordered phases occurring solely on the basis of network geometry. Thus *the network phase is not due to external noise but is instead a design component.* The system can be tuned—for instance, by evolution—to the most advantageous phase.

The question of *which* phase is most advantageous, of course, depends on the minibrain's surrounding environment. This immediately motivates further research into the question of which environments might favour which of the different phases. In addition, considering the biological motivation for the model, one might want to investigate whether behaviour resembling any (or all) of these phases is to be found in real neural systems of different types.

The results here also emphasize the importance of negative feedback as providing behavioural control in learning systems. The accidental destruction of good connections that drives the different phases in the minibrain model relies on the fact that, because of the negative-only feedback, the active and inactive connections in the network are of similar strength [13, 14, 30]. Not only are such systems faster and more adaptive [13, 14], they also place in the hands of the designer aspects which would otherwise be in the hands of the environment.

In summary, the minibrain's combination of negative-only feedback and extremal dynamics can be seen to cause complex behavioural patterns that depend much more on internal characteristics of the network than any external pressures. This simple neural model has already been shown to have interesting implications for learning systems and the importance of selectionist procedures; the results presented here extend this to the wider context of phase transitions on networks and their origins.

Acknowledgments

Special thanks and many fond memories to Per Bak, who took the time despite serious illness to offer encouragement and many helpful suggestions. Grateful thanks also to Maya Paczuski, Ginestra Bianconi, Andrea Capocci, Damien Challet, Dante Chialvo, Paulo Laureti, and Yi-Cheng Zhang, for advice and stimulating conversations, and to anonymous referees for their comments. This work was supported in part by Swiss National Science Foundation grant no. 20-61470.00.

References

[1] S. Kauffman, The Origins of Order: Self-Organization and Selection in Evolution, Oxford University Press, Oxford, 1993.
[2] E. R. Berlekamp, J. H. Conway, R. K. Guy, Winning Ways For Your Mathematical Plays, A. K. Peters, Ltd., 2001.
[3] P. Bak, C. Tang, K. Wiesenfeld, Phys. Rev. Lett. 59 (1987) 381.
[4] P. Bak, C. Tang, K. Wiesenfeld, Phys. Rev. A 38.
[5] P. Bak, K. Sneppen, Phys. Rev. Lett. 71 (1993) 4083.
[6] P. Bak, How Nature Works: The Science of Self-Organized Criticality, Oxford University Press, Oxford, 1997.
[7] D. O. Hebb, The Organization of Behaviour, Wiley, New York, 1949.
[8] A. G. Barto, Hum. Neurobiol. 4 (1985) 229.
[9] A. G. Barto, M. I. Jordan, Proc. IEEE Int. Conf. on Neural Networks 2 (1987) 629.
[10] P. Mazzoni, R. A. Andersen, M. I. Jordan, Proc. Natl. Acad. Sci. USA 88 (1991) 4433.

[11] P. Alstrøm, D. Stassinopoulos, Phys. Rev. E 51 (1995) 5027.
[12] D. Stassinopoulos, P. Bak, Phys. Rev. E 51 (1995) 5033.
[13] D. R. Chialvo, P. Bak, Neurosci. 90 (1999) 1137.
[14] P. Bak, D. R. Chialvo, Phys. Rev. E 63 (2001) 031912.
[15] J. Wakeling, P. Bak, Phys. Rev. E 64 (2001) 051920.
[16] T. Kohonen, Self-Organizing Maps, Springer-Verlag, Berlin, 2001.
[17] J. J. Hopfield, Proc. Natl. Acad. Sci. USA 79 (1982) 2554–2558.
[18] D. J. Amit, H. Gutfreund, H. Sompolinsky, Phys. Rev. Lett. 55 (1985) 1530.
[19] E. Gardner, J. Phys. A 20 (1987) 3453.
[20] H. Sompolinsky, N. Tishby, H. S. Seung, Phys. Rev. Lett. 65 (1990) 1683.
[21] E. Barkai, D. Hansel, H. Sompolinsky, Phys. Rev. A 45 (1992) 4146.
[22] H. S. Seung, H. Sompolinsky, N. Tishby, Phys. Rev. A 45 (1992) 6056.
[23] T. L. H. Watkin, A. Rau, M. Biehl, Rev. Mod. Phys. 65 (1993) 499.
[24] J. A. Flanagan, Phys. Rev. E 63 (2001) 036130.
[25] G. Parisi, J. Phys. A 19 (1986) L617.
[26] J. Z. Young, Proc. Royal Soc. London B 163 (1965) 285.
[27] R. M. Fitzsimonds, H. J. Song, M. M. Poo, Nature 388 (1997) 439.
[28] J. S. Albus, Brains, Behaviour and Robotics, BYTE Books, McGraw-Hill, Peterborough, NH, 1981.
[29] G. M. Edelman, Neural Darwinism: The Theory of Neuronal Group Selection, Basic Books, New York, 1987.
[30] T. Araujo, R. Vilela Mendes, Complex Syst. 12 (2000) 357.