A neural network for storing individual patterns in limit cycles

1991 J. Phys. A: Math. Gen. 24 5105

(http://iopscience.iop.org/0305-4470/24/21/022)

J. Phys. A: Math. Gen. 24 (1991) 5105–5119. Printed in the UK

A neural network for storing individual patterns in limit cycles

Varsha Deshpande and Chandan Dasgupta

Department of Physics, Indian Institute of Science, Bangalore 560 012, India

Received 15 March 1991

Abstract. A neural network model in which individual memories are stored in limit cycles is studied analytically and numerically. In this model there are two kinds of interactions: a Hopfield-like term that tends to stabilize the system in a memorized state and a second term with a time delay that acts to induce transitions between a memorized state and its complement state. For a proper choice of the values of the parameters, this model exhibits limit cycle behaviour in which the overlap with a target pattern oscillates in time. An asymmetrically diluted version of the model is studied analytically in the limit of extreme dilution. We find that the model with cycles performs better than a similarly diluted version of the Hopfield model. The performance of the fully connected model is studied by numerical simulations. We find a behaviour qualitatively similar to that of the dilute model. The model with cycles is found to perform better than the Hopfield model as a pattern classifier if the memory loading level and the degree of corruption of the input patterns are high.

1. Introduction

During recent years, neural networks performing computations through attractors have received a lot of attention (for a review see Amit 1989). These networks consist of a large number of simple computing elements (neurons). The computation to be performed is coded in the synaptic interconnections among the neurons. The time evolution of the network is governed by an assumed dynamics of individual neurons. Starting from an initial state, the network evolves in time until a time-persistent state (an attractor of the underlying dynamics) is reached. This stationary state represents the result of the computation. In most of the networks studied so far, the attractors used in the computation are fixed points. In a large class of neural network models of associative memory, the interactions are chosen so as to make the states representing the stored memories fixed points of the dynamics. Any initial state close to one of the memories, representing partial knowledge of the stored information, converges under the dynamics to the memory state. Thus, retrieval of the full information is achieved. These models also act as pattern classifiers. All initial states within the basin of attraction of a particular memory are classified as belonging to the same group. The Hopfield model (Hopfield 1982, 1984) of associative memory is the simplest model of this type.

A general dynamical system involving a large number of interacting variables may, of course, exhibit attractors other than simple fixed points. In this paper, we consider a neural network in which computations are performed with limit cycles. The idea of using time-dependent states in neural computing has received some attention in the past (Peretto and Niez 1986, Sompolinsky and Kanter 1986, Kleinfeld 1986, Dehaene *et al* 1987, Buhman and Schulten 1987, Mori *et al* 1989). The model we consider here

0305-4470/91/215105+15$03.50

is similar to those constructed by Sompolinsky and Kanter (1986) and by Kleinfeld (1986) for the recognition and associative recall of temporal sequences and cycles of patterns. In these models the neurons are assumed to be two-state variables, represented by Ising spins. Two different sets of synaptic interactions among the neurons are assumed to exist. The first set is the usual Hopfield type, tending to stabilize the system in a memory state. The second set of interactions tends to induce transitions from one memory state to another. A time delay associated with the second set of interactions causes the system to stay in a memory state for some time before making a transition to another one. The model we study here is one in which the second set of interactions tends to induce transitions from a memory state to its complement state which is obtained by reversing the signs of all the Ising spins in the memory state. The model and its dynamics are described in detail in section 2. Both sets of interactions in this model are determined by the generalized Hebb rule (Hopfield 1982) and are, therefore, symmetric. An interaction parameter belonging to the second set differs from the corresponding one in the first set in its sign, its magnitude and in the assumed time delay τ associated with its action. If the strength of the second set of interactions is sufficiently large, then the network, when started off from a state close to one of the memory states, goes into a limit cycle in which the overlap of the instantaneous spin configuration with the chosen memory state oscillates in time with a period $\sim 2\tau$, and the overlaps with all other memory states remain close to zero. This way of embedding individual patterns in limit cycles was briefly discussed by Sompolinsky and Kanter (1986), but they did not present any analysis of the behaviour of this model.

The synaptic interactions in the model we consider are not instantaneous. For this reason, the methods of equilibrium statistical mechanics, which have been used extensively to analyse the behaviour of many Hopfield-type models (Amit 1989), cannot be used to study its properties. An asymmetrically diluted version of this model can be studied analytically in the limit of extreme dilution by using a method developed by Derrida *et al* (1987). This calculation is described in section 3. We consider a 'zero-temperature' (deterministic) dynamics and randomly chosen memories. The behaviour of the model is then characterized by two parameters, λ and α. The first parameter λ measures the strength of the second set of interactions relative to the first, Hopfield-like, set. The parameter α is the ratio between the total number of stored memories and the average number of neurons connected to a particular one. The calculated 'phase diagramme' in the λ-α plane exhibits all the qualitative features found by Gutfreund and Mezard (1988) in a similar calculation on networks generating temporal sequences and cycles of patterns. For any $\lambda > 0$, the network is found to exhibit the desired cyclic behaviour if $\alpha_{\min}(\lambda) < \alpha < \alpha_{\max}(\lambda)$, with $\alpha_{\min} = 0$ for $\lambda > 1$. The values of $\alpha_{\max}$ for λ close to unity are found to be larger than α_c, the maximum storage capacity of a similarly diluted version of the Hopfield model. This calculation, thus, shows that the maximum storage capacity of a network in which the memories are stored in limit cycles is higher than that of one in which the memories are stored as fixed points.

In section 4, we describe the results obtained from numerical simulations of the behaviour of the fully connected model. We find results which are qualitatively similar to the predictions of the analytic calculation on the dilute model. The numerical values of quantities such as $\alpha_{min}(\lambda)$ and $\alpha_{max}(\lambda)$, however, turn out to be quite different from the values obtained from the analytic calculation. We also find that this model performs better than the Hopfield model in correctly classifying corrupted input patterns if the memory loading level α and the degree of corruption of the input patterns are high. The retrieval performance of this network is, however, not very good. For large values of α, the amplitude of oscillations of the overlap of the instantaneous spin configuration with a target pattern is rather small. This model, thus, would be useful as a pattern classifier where the primary objective is to identify correctly the class to which the input pattern belongs, rather than as an associative memory where the emphasis is on a correct retrieval of the complete information from partial knowledge.

In section 5, we summarize the main results obtained from this study.

2. The model

We consider a network which consists of N neurons, represented by the Ising variables σ_i, $i=1,2,\ldots,N$. A number, p, of random binary patterns, ξ_i^μ, $i=1,2,\ldots,N$; $\mu=1,2,\ldots,p$, are stored in the network. Two different sets of synaptic interactions among the neurons are assumed to exist. The first set is the usual Hopfield type, which stabilizes the system in a memory state. The second set of interactions tends to induce transitions from a memory state to its complement state which is obtained by reversing the signs of all the Ising spins in the memory state. A time delay associated with the second set of interactions causes the system to stay in a memory state for some time before making a transition to the complement. We assume a deterministic (zero-temperature) dynamics defined by the following update rule:

$$\sigma_i(t+1)=\mathrm{sign}(h_i(t)) \tag{1}$$

where t represents a discrete 'time' label and $h_i(t)$ is the local field acting on the ith spin at time t. The updates may be synchronous or asynchronous. The local field $h_i(t)$ is defined as

$$h_i(t)=\sum_j J_{ij}\sigma_j(t)-\lambda\sum_j J_{ij}\sigma_j(t-\tau) \tag{2}$$

where $\lambda>0$ is a control parameter, τ represents a time delay associated with the second set of interactions and the interaction matrix J_{ij} has the Hebb rule form

$$J_{ij}=\frac{1}{N}\sum_\mu \xi_i^\mu\xi_j^\mu. \tag{3}$$

The first set of interactions acts to stabilize the system in a memory state whereas the second set tends to induce transitions between a memory state and its complement state. The behaviour of this model may be analysed easily in the limit $N\to\infty$, $\alpha=p/N\to 0$. Let us assume that the system has settled down into a memory state, $\sigma_i=\xi_i^{\mu_0}$ at time $t=0$. Configurations at earlier times are assumed to be uncorrelated with any memory state. Thus, for $t<\tau$, the second term in equation (2) does not have much of an effect and the system remains in the memory state $\sigma_i=\xi_i^{\mu_0}$. At time $t=\tau$, the local field acting on the ith spin is given by

$$h_i(\tau)=\xi_i^{\mu_0}(1-\lambda)+\mathrm{O}(1/\sqrt{N}). \tag{4}$$

Thus, if $\lambda>1$, the system makes a transition to the complement state, $\sigma_i=-\xi_i^{\mu_0}$, at time $t=\tau$. This state again flips back to the original memory state at time $t=2\tau$ and, thus, the system continues to oscillate in a limit cycle in which the overlap of the spin configuration with the memory state $\{\xi_i^{\mu_0}\}$ varies periodically between $+1$ and -1 with a time period $T\simeq 2\tau$ and overlaps with other memories remain close to zero. There are p such cycles, corresponding to the p different memories. If the system is started off in a configuration close to one of these memories or its complement, then it goes into the corresponding limit cycle. This network, therefore, functions as a pattern classifier and as an associative memory in which each memorized pattern is stored in a limit cycle. Note that in the present model, a memorized pattern and its complement correspond to the same attractor. In contrast, a pattern and its complement represent two different fixed points in the original Hopfield model.

If the parameter α measuring the loading level of the network is of order unity, then the 'noise' term arising from the interference of non-condensed memory states cannot be neglected and the simple analysis described above breaks down. Analytic and numerical studies of the behaviour of such models are described in the next two sections.

3. Analytic results in the limit of extreme dilution

In this section, we consider an asymmetrically diluted version of the model defined in the preceding section. This model can be solved exactly in the limit of extreme dilution by using methods developed by Derrida *et al* (1987) and by Gutfreund and Mezard (1988). In the model we consider, each interaction J_{ij} is multiplied by a quenched random parameter C_{ij}, so that the local field at site i at time t is given by

$$h_i(t) = \sum_j C_{ij} J_{ij}(\sigma_j(t) - \lambda\sigma_j(t-\tau)). \tag{5}$$

The probability distribution for each independent random variables C_{ij} is assumed to have the form

$$P(C_{ij}) = \frac{c}{N}\,\delta(C_{ij}-1) + \left(1-\frac{c}{N}\right)\delta(C_{ij}). \tag{6}$$

We consider C_{ij} and C_{ji} to be independent variables, so that the dilution is asymmetric. In this calculation, we assume a different normalization for the interaction parameters J_{ij} which are now defined as

$$J_{ij} = \frac{1}{c}\sum_\mu \xi_i^\mu \xi_j^\mu. \tag{7}$$

The memory loading level α for this model is defined as $\alpha = p/c$ where p is the number of patterns stored in the network and c, by definition, is the average number of spins connected to a particular one. We assume a parallel synchronous dynamics governed by the update rule

$$\sigma_i(t+1) = \pm 1 \text{ with probability } [1+\exp(\mp 2h_i(t)/T)]^{-1} \tag{8}$$

where T is the 'temperature'. This update rule reduces to the deterministic one defined in equation (1) in the limit $T \to 0$.

It was pointed out by Derrida *et al* that the dynamics of models of this type can be solved exactly in the extreme dilution limit $N \to \infty$, $C \to \infty$, $C/N \to 0$, $C \ll \ln N$. Similarly diluted versions of a number of models for generation and recognition of temporal sequences and cycles of patterns were solved by Gutfreund and Mezard (1988). Following these calculations, we consider a situation where the system has a macroscopic overlap, $m^{\mu_0}(t)$, with a particular pattern, $\{\xi_i^{\mu_0}\}$, and microscopic overlaps with all other patterns at time t. Let $m^{\mu_0}(t-\tau)$ be the overlap with the same pattern at the earlier time, $(t-\tau)$. These overlaps are defined as

$$m^{\mu_0}(t)=\langle \xi_i^{\mu_0}\sigma_i(t)\rangle/N \tag{9}$$

where the brackets $\langle\ldots\rangle$ denote a thermal as well as configurational average. The local field at the ith spin at time t is given by

$$h_i(t)=\sum_{\alpha=1}^{K_i} J_{ij_\alpha}(\sigma_{j_\alpha}(t)-\lambda\sigma_{j_\alpha}(t-\tau)) \tag{10}$$

where j_α, $\alpha=1,2,\ldots,K_i$ denote the spins connected to the ith one with non-zero interactions. As pointed out by Derrida *et al* (1987), the correlations among the spins $\{\sigma_{j_\alpha}\}$ may be neglected in the limit of extreme dilution considered here. The probability distribution for $h_i(t)$ may then be calculated from the known probability distributions for individual spins:

$$\begin{aligned}
P(\sigma_{j_\alpha}(t))&=\frac{1+m(t)}{2}\,\delta_{\sigma_{j_\alpha}(t),\xi_{j_\alpha}^{\mu_0}}+\frac{1-m(t)}{2}\,\delta_{\sigma_{j_\alpha}(t),-\xi_{j_\alpha}^{\mu_0}}\\
P(\sigma_{j_\alpha}(t-\tau))&=\frac{1+m(t-\tau)}{2}\,\delta_{\sigma_{j_\alpha}(t-\tau),\xi_{j_\alpha}^{\mu_0}}+\frac{1-m(t-\tau)}{2}\,\delta_{\sigma_{j_\alpha}(t-\tau),-\xi_{j_\alpha}^{\mu_0}}.
\end{aligned} \tag{11}$$

In equations (11), we have dropped the superscript μ_0 of m for notational convenience. From these equations, it follows that

$$\begin{aligned}
P(\sigma_{j_\alpha}(t)=\sigma_{j_\alpha}(t-\tau))&=(1+m(t)m(t-\tau))/2\\
P(\sigma_{j_\alpha}(t)=-\sigma_{j_\alpha}(t-\tau))&=(1-m(t)m(t-\tau))/2.
\end{aligned} \tag{12}$$

In defining these probabilities, we have assumed that correlations existing between the spin configurations $\{\sigma_i(t-\tau)\}$ and $\{\sigma_i(t)\}$ are only those arising from their assumed overlaps, $m(t-\tau)$ and $m(t)$, with the memory state $\{\xi_i^{\mu_0}\}$. The validity of this assumption will be discussed later.

Using the definition, equation (7), of the interaction matrix J_{ij} and the probabilities defined in equations (11) and (12), it is straightforward to show that the local field $h_i(t)$ can be written as

$$h_i(t)=\xi_i^{\mu_0}(X1-\lambda X2+X3) \tag{13}$$

where the average values of $X1$ and $X2$ are $m(t)$ and $m(t-\tau)$, respectively, and $X3$ may be treated as a Gaussian random variable with zero average and a variance equal to

$$w=\alpha(1+\lambda^2-2\lambda m(t)m(t-\tau)). \tag{14}$$

Combining this information with the update rule, equation (8), we obtain the following equation describing the time evolution of the overlap m:

$$m(t+1)=\int\frac{\mathrm{d}y}{\sqrt{2\pi}}\exp(-y^2/2)\tanh[(m(t)-\lambda m(t-\tau)+\sqrt{w}y)/T]. \tag{15}$$

The random Gaussian variable y in equation (15) represents the 'noise' produced by the interference of the uncondensed patterns $\{\xi_i^\mu\}$, $\mu\neq\mu_0$. In the deterministic, $T\to 0$

limit, equation (15) becomes

$$m(t+1) = \mathrm{Erf}[(m(t) - \lambda m(t-\tau))/\sqrt{(2w)}]. \qquad (16)$$

These equations, together with the initial condition, $m(t) = Q$, $-\tau < t \leqslant 0$, determine the values of $m(t)$ at all times $t > 0$.

We have studied numerically the time evolution of $m(t)$ given by equation (16). We find that for $0 < \lambda < 1$, there exists a critical value, $\alpha_{\min}(\lambda)$, of the loading level α such that for $\alpha > \alpha_{\min}(\lambda)$, the overlap $m(t)$ shows oscillations with a period $\sim 2\tau$. Typical results for the time evolution of $m(t)$ for $\lambda = 0.8$, $\tau = 10$, $Q = 0.97$ and two different values, 0.25 and 0.75, of α are shown in figure 1. These results show that the internal noise generated by the microscopic overlaps with the non-condensed patterns may cause transitions between a memory state and its complement even if the parameter λ is less than the critical value (=1) needed for cyclic behaviour in the $\alpha \to 0$ limit. The variation of $\alpha_{\min}$ with λ for $\lambda < 1$, calculated with $Q = 0.98$, is shown in the 'phase diagram' of figure 2. The value of $\alpha_{\min}$ is equal to zero for all $\lambda > 1$. As can be seen in figure 1, the amplitude of oscillations of $m(t)$ decreases initially and then settles down to a constant value which decreases as α increases. The asymptotic value of this amplitude, m_0, may be obtained from the self-consistent equation

$$m_0 = \mathrm{Erf}\{(1+\lambda)m_0/\sqrt{[2\alpha(1+\lambda^2+2\lambda m_0^2)]}\}. \qquad (17)$$

This equation is derived from equation (16) by noting that in the steady state, $m(t_0) = -m(t_0 - \tau) = m_0$, where t_0 refers to a peak of $m(t)$. The amplitude m_0 asymptotically approaches a non-zero value if the $m_0 = 0$ solution of equation (17) is unstable. This happens if $\alpha < \alpha_{\max}(\lambda)$ where the maximum storage capacity $\alpha_{\max}$ is given by

$$\alpha_{\max}(\lambda) = \frac{2}{\pi}\frac{(1+\lambda)^2}{1+\lambda^2}. \qquad (18)$$

This result for $\alpha_{\max}$ is the same as that found by Gutfreund and Mezard (1988) for a model in which all the patterns are embedded in a single-limit cycle. The evolution

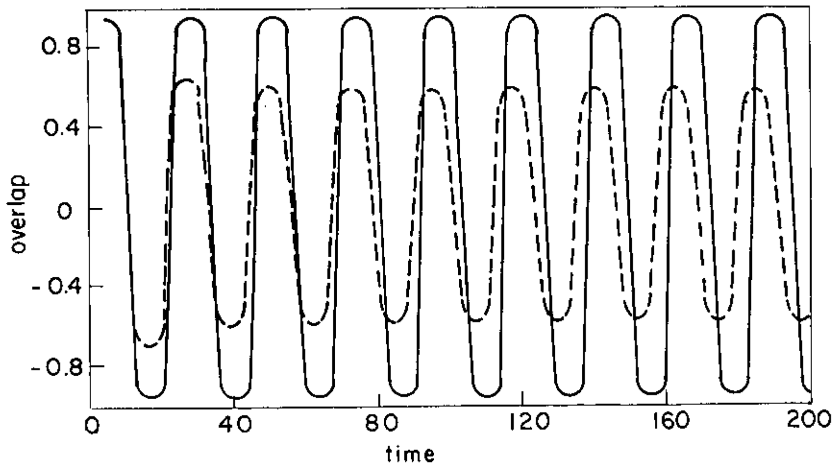

Figure 1. Numerical solution of equation (16) showing the time evolution of $m(t)$ for the diluted model with $\lambda = 0.8$, $\tau = 10$ and $Q = 0.97$ for $\alpha = 0.25$ (full curve) and 0.75 (broken curve).

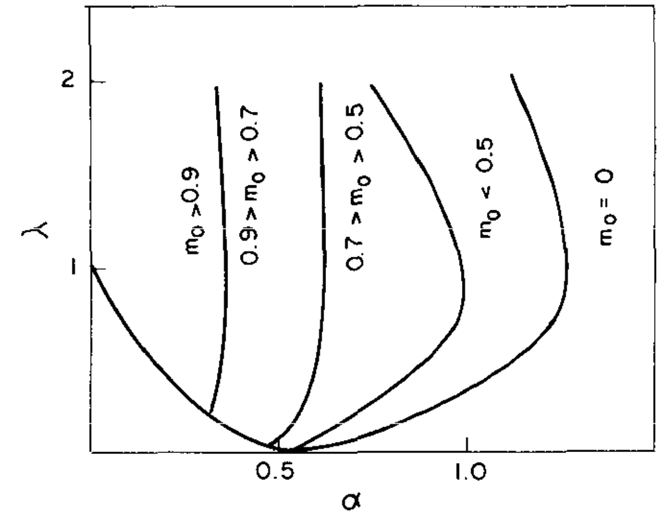

Figure 2. Phase diagram in the (λ, α) plane indicating the region where limit cycle behaviour with a finite amplitude is possible. The initial overlap Q was taken to be 0.98. The decrease in the amplitude of oscillation of $m(t)$ with the increase in loading level α of the network can be observed.

equations of the overlaps and the values of $\alpha_{\min}(\lambda)$ are, however, different in the two models. The maximum storage capacity of the present model near $\lambda = 1$ is two times that of a similarly diluted version of the Hopfield model ($\lambda = 0$). Thus, we find that the storage capacity of the network is increased if the patterns are stored as limit cycles instead of as fixed points.

In figure 3 we have shown the variation of the amplitude m_0 with the memory-loading parameter α for $\lambda = 1.2$. The α dependence of the self-consistent value of the memory overlap obtained for the dilute Hopfield model ($\lambda = 0$) is also shown in the same figure

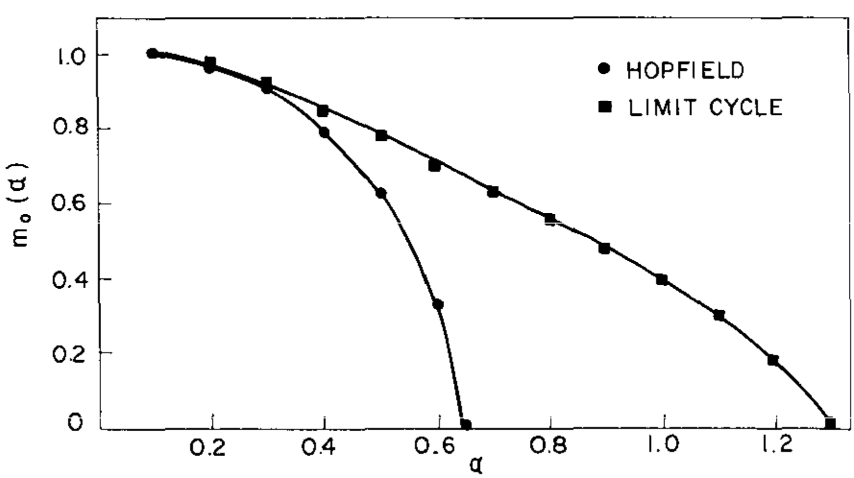

Figure 3. Variation of the amplitude m_0 with the loading level α for the dilute model with cycles (squares) for $\lambda = 1.2$. The dependence of the retrieval overlap m on α for the dilute Hopfield model is also shown (dots) for comparison.

for comparison. It is clear that the model with cycles performs better than the Hopfield model for all values of α. It should, however, be noted that since m_0 decreases continuously with α, the retrieval performance of the model with cycles is not very good for large values of α. Thus, this model would be more useful as a pattern classifier than as a pattern retriever.

The reason for the increased storage capacity of the limit cycle model is not difficult to understand. It is closely related to the assumption made earlier (in deriving equation (12)) that correlations between the spin configurations at time t and time $(t-\tau)$ are only those arising from their overlaps, $m(t)$ and $m(t-\tau)$, with the selected memory state $\{\xi_i^{\mu_0}\}$. If this assumption is correct then the fact that $m(t_0)=-m(t_0-\tau)$ in the steady state does not mean that $\sigma_i(t_0)=-\sigma_i(t_0-\tau)$ for all i. As a result, the noise terms (arising from the interference of non-condensed patterns) associated with the instantaneous and the time-delayed interactions do not always add in phase. This effect causes a reduction of the variance w of the noise from the maximum possible value, $\alpha(1+\lambda)^2$, to the value, $(1+\lambda^2+2\lambda m^2(t_0))$, used in equation (17). The signal term, $(1+\lambda)m(t_0)$, appearing in the argument of the error function in equation (17) is independent of the nature of correlations between the states at time t_0 and $(t_0-\tau)$. Thus, in the limit cycle model, the 'signal term' is increased by the factor $(1+\lambda)$, but the noise is increased by a smaller factor, the difference between the two factors being large if $m(t_0)$ is small. The net effect is a relative suppression of the noise, resulting in an increase in the memory storage capacity. If, instead, we make the assumption of maximal correlation between the spin states at times t_0 and $(t_0-\tau)$, i.e. if we assume that $\sigma_i(t_0)=-\sigma_i(t_0-\tau)$ for all i, then the variance of the noise distribution would have the maximum possible value, $w=(1+\lambda)^2$. The self-consistent equation that determines $\alpha_{\max}(\lambda)$ would then be identical to that for the diluted Hopfield model, which is obtained by setting $\lambda=0$ in equation (17). Thus, the storage capacity of the dilute limit cycle model depends crucially on the nature of correlation between the states separated by the time delay τ. If these two states are perfectly correlated, i.e. if one is the complement of the other, then the values of $\alpha_{\max}(\lambda)$ and $m_0(\alpha,\lambda)$ would be the same as those for the dilute Hopfield model ($\lambda=0$) for all values of λ. Any reduction between the correlations between these two states would cause an increase in the values of $\alpha_{\max}(\lambda)$ and $m_0(\alpha,\lambda)$. It is difficult to determine exactly the amount of correlation actually present. We have tried to derive an evolution equation, similar to equation (15), for the overlap $q(t,\tau)=(1/N)\sum_i\sigma_i(t)\sigma_i(t-\tau)$. We find that the equation for $q(t,\tau)$ involves the overlap, $q(t,2\tau)$, between the spin configurations at times t and $(t-2\tau)$. This leads to a hierarchy of equations which does not close. In the absence of any exact result, the assumption of minimal correlation is plausible if the time delay τ is chosen to be large. We note here that the continuous time dynamics of the diluted Hopfield model ($\lambda=0$) has been analysed exactly by Kree and Zippelius (1987). They find that the time-persistent part of the spin autocorrelation function in the steady state is less than unity for all values of α and T (including $T=0$) in the retrieval 'phase', and that it approaches the value m_0^2 as $\alpha\to\alpha_c=2/\pi$ at $T=0$. They also find that the relaxation time characterizing the decay of the autocorrelation remains finite for all values of α and T. These results lend support to the assumption we have made in deriving equation (12). Numerical simulations on the fully connected model described below indicate that the correlation measured by the quantity $|q(t_0,\tau)|$ lies somewhere in between the maximal value, +1, and the minimal value, m_0^2. Thus, we expect the limit cycle model to show some increase in the storage capacity and the retrieval overlaps over the dilute Hopfield model, the improvement being somewhat smaller than that shown in figure 3. We note that the result, equation (18), for $\alpha_{max}(\lambda)$ remains valid as long as the overlap $q(t_0, \tau)$ goes to zero as $m_0 \to 0$.

4. Simulation results for the fully connected model

The analytic calculation on the dilute model provides an indication of what kind of behaviour is expected from the fully connected model exhibiting limit cycles. In order to determine whether the qualitative behaviour of the fully connected model is similar to that of the dilute model, we have carried out a number of numerical simulations on the model defined in equation (2). The main results obtained from these simulations are described below.

The first question we addressed in the simulation study is whether the fully connected model exhibits stable cyclic behaviour for sufficiently large values of α if the control parameter λ is less than unity. In these simulations, we used the asynchronous deterministic update rule defined in equation (1). As large finite-size effects were found to be present, the behaviour of the network was simulated for three different sizes, $N = 40$, 100 and 200. The value of λ was set at 0.8 and four different values of α ($\alpha = 0.05$, 0.1, 0.15 and 0.2) were studied. The initial configuration of the network at time $t = 0$ was taken to be one of the memory states $\{\xi_i^{\mu_0}\}$ and configurations at earlier times were assumed to be random. The time evolution of the network for $t > 0$ was simulated for a few hundred time units where a time unit corresponds to one attempted update per spin. The number of such runs made for each value of α was 375 for $N = 40$, 225 for $N = 100$ and 200 for $N = 200$. The results obtained from these simulations are summarized in table 1. The first row shows the percentage of failures where a particular run is considered to be a failure if it does not show the expected periodic behaviour. It is clear from the numbers that for $\alpha < 0.05$, the network does not exhibit any cyclic behaviour. For $\alpha = 0.15$ and 0.2, the percentage of failures decreases rapidly as N increases, indicating that these values of α are larger than α_{min}, the minimum value of α needed for a periodic time evolution. For $\alpha = 0.1$, the percentage of failures

Table 1. Percentage of failure (first row) and peak positions of the distributions of the absolute values of the largest Fourier components of $m^{\mu_0}(t)$ (second row) and $m^{\mu}(t)$, $\mu \neq \mu_0$ (third row) are shown for $N = 200$, 100 and 40 with values of $\alpha = 0.2$, 0.15, 0.1 and 0.05.

		Memory loading level (α)			
N		0.2	0.15	0.1	0.05
200	Failure (%)	6	48	91.5	100
	g^{μ_0}	70	80	100	—
	g^{μ}	20	20	20	—
100	Failure (%)	44	70	96.5	100
	g^{μ_0}	65	70	90	—
	g^{μ}	30	30	30	—
40	Failure (%)	74	95	100	—
	g^{μ_0}	60	55	—	—
	g^{μ}	60	60	—	—

decreases slowly with increasing N, but remains large for $N = 200$. From these observations we conclude that the fully connected network with $\lambda = 0.8$ exhibits a cyclic behaviour if $\alpha > \alpha_{min} \simeq 0.1$. This value of α_{min} is close to the analytic result obtained for the dilute model. In figure 4, we have shown the time evolution of the overlap, $m^{\mu_0}(t)$, of the spin configuration at time t with the original memory state for $N = 200$, $\lambda = 0.8$, $\tau = 10$ time units and two different values, 0.1 and 0.2, of α. The curves shown were obtained by averaging over all the runs exhibiting cycles. The qualitative behaviour of $m^{\mu_0}(t)$ is found to be similar to that seen in the dilute model (see figure 1). For both values of α, the amplitude of oscillations of $m^{\mu_0}(t)$ initially decreases with time, eventually settling down to a constant value (this has been checked by performing longer runs). The asymptotic value $m_0^{\mu_0}$ of the amplitude decreases as α increases. The values of $m_0^{\mu_0}$ obtained for the fully connected model are considerably smaller than those for the dilute model at the same values of λ and α. This observation suggests that the value, $\alpha_{max}(\lambda)$, of α at which $m_0^{\mu_0}$ goes to zero in the fully connected model is much smaller than that calculated for the dilute model. We did not attempt to determine $\alpha_{max}(\lambda)$ precisely from the simulations because such a calculation would require very long runs. We believe that the question of whether the amplitude of oscillations goes to zero at very long times is not a particularly important one because most of the computations to be performed by the network (such as the ones described below) require persistence of oscillations with a substantial amplitude for a few (~ 10) time periods.

The fact that the network exhibits a cyclic behaviour for $\alpha > \alpha_{min}(\lambda)$ does not necessarily imply that the cycle the network gets into is a true 'memory cycle'. By a memory cycle, we mean one in which the amplitude, $m_0^{\mu_0}$, of oscillations of the overlap with the original memory state μ_0 is large and the amplitudes m_0^{μ}, $\mu \neq \mu_0$ associated with overlaps with all other memory states are substantially smaller. This is to be

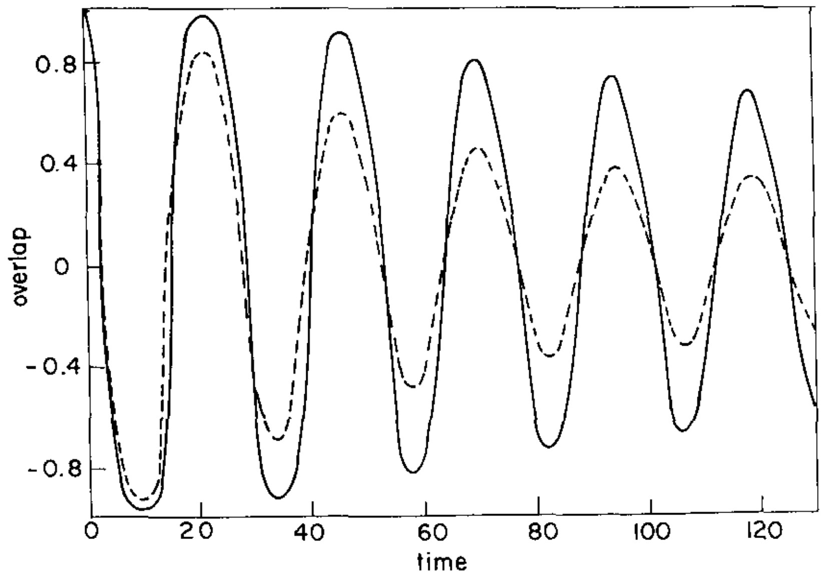

Figure 4. Time evolution of the overlap, $m^{\mu_0}(t)$, between the state of the network at instant t and the original memory state for a fully connected network with $N = 200$, $\lambda = 0.8$ and $\alpha = 0.1$ (full curve) and 0.2 (broken curve). Each time unit corresponds to one attempted spin flip per spin.

contrasted with a 'spurious cycle' in which the network oscillates in a state which is not strongly correlated with the initial memory state. In such a spurious cycle, the amplitude of the overlap with the original pattern may be smaller than that of the overlap with some other pattern. So, we operationally defined a spurious cyclic state as one in which the amplitude $m_0^{\mu_0}$ is not the largest one among the set $\{m_0^\mu\}$. In order to obtain a quantitative measure of the probability of occurrence of spurious cycles, we calculated the Fourier transforms, $\{f^\mu(\omega)\}$, of the overlaps $\{m^\mu(t)\}$ for each run. The Fourier amplitudes were defined as

$$f^\mu(\omega) = \sum_{n=1}^{M} m^\mu(n\Delta t)\exp(\mathrm{i}\omega n\Delta t) \tag{19}$$

where the frequency ω takes the discrete values $\omega = 2\pi j/M\Delta t$, $j = 1, 2, \ldots, M$. In the calculation, we took $M = 256$ and set the value of Δt at 0.5 time unit. For each memory μ, the distribution of $f^\mu(\omega)$ was found to have a sharp peak at a frequency $\omega = \omega_{\max} \simeq \pi/\tau$. We denote by g^μ the value of $f^\mu(\omega_{\max})$ and consider a cycle to be spurious if one or more g^μs with $\mu \neq \mu_0$ have values larger than g^{μ_0}. We find that for $N = 200$, $\lambda = 0.8$ and $\alpha = 0.1 - 0.20$, the fraction of runs in which the cycle the system gets into turns out to be a spurious one according to the criterion described above is less than 10%. The fraction of runs showing spurious cycles decreases rapidly as N increases, indicating that the probability of getting into a spurious cycle would go to zero in the large-N limit for these values of α if the starting configuration coincides with one of the memory states. This network, thus, is able to retain the 'memory' of the initial state if it is started off in one of the stored configurations.

In order to obtain a quantitative measure of the degree of discrimination between the original memory state and other stored memories, we separately calculated the distributions of g^{μ_0} and g^μ, $\mu \neq \mu_0$. In figure 5, we have shown these two distributions obtained for $N = 200$, $\alpha = 0.15$ and $\lambda = 0.8$. The two distributions are found to be well separated with little overlap, showing that the network is able to discriminate quite effectively between the original and other memories. In the second and third rows of

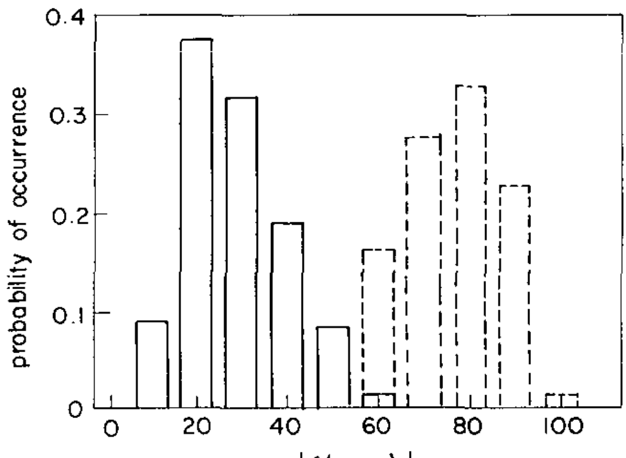

Figure 5. Distributions of the absolute values of the largest Fourier components of the overlaps $m^{\mu_0}(t)$ (dotted lines) and $m^\mu(t)$, $\mu \neq \mu_0$ (full lines) for $N = 200$, $\alpha = 0.15$ and $\lambda = 0.8$. The well-separated distributions indicate that the network cycles in the original pattern most of the time.

table 1, we have given the values of the peak positions of the distributions of g^{μ_0} and g^{μ}, $\mu \neq \mu_0$ for $\lambda = 0.8$ and different values of N and α. For a fixed value of α in the interval $0.1 < \alpha < 0.20$, the position of the peak of the g^{μ_0} distribution moves to higher values as N increases. This observation tells us that the oscillations of the overlap with the original memory state would become more pronounced in the large-N limit. The peak position of the g^{μ} distribution, on the other hand, is found to move to lower values with increasing N. Thus, the separation between the peaks of the two distributions, which is a measure of the degree of discrimination achieved by the network, increases with N. From these observations, we conclude that this network would show the desired memory cycle behaviour in the large-N limit for $\lambda = 0.8$ and $0.1 < \alpha < 0.20$. The results shown in table 1 also indicate that the degree of discrimination between the initial and other memory states decreases, as expected, with increasing α for fixed values of N and λ.

We have also carried out similar simulations for values of λ greater than one. We find $\alpha_{\min} = 0$ for $\lambda > 1$, as expected. The time evolution of the network is found to be qualitatively similar to that described above. Quantitatively, we find that the performance of the network improves as λ is increased above one. For example, for $\lambda = 1.5$, $N = 200$ and $0.05 < \alpha < 0.20$, true memory cycle behaviour is found in all runs. For a given value of α, both the peak position of the g^{μ_0} distribution (which is a measure of the amplitude of oscillations of the overlap with the original memory state) and the separation between the peaks of the g^{μ_0} and g^{μ} distributions (which is a measure of the ability of the network to discriminate between memorized patterns) are found to increase with λ. The analytic results obtained in the preceding section suggest that the maximum storage capacity of the network begins to decrease as λ is increased beyond one. Thus, a value of λ close to but greater than unity should be optimal for this network.

The results described above suggest that the maximum storage capacity of the present network for $\lambda > 1$ is higher than that of the Hopfield network which is known (Amit *et al* 1987) to function as an associative memory if the value of α is less than $\alpha_c \simeq 0.15$. As discussed earlier in this section, an increase in the storage capacity is to be expected if the configurations at times t_0 and $(t_0 - \tau)$ are not very strongly correlated. Here, as before, the time t_0 corresponds to the peak of $m^{\mu_0}(t)$ in the steady state. In order to determine if this is indeed the case, we calculated the correlation $q(t_0, \tau)$ defined earlier in the steady state from the spin configurations generated in the simulations. The observed values of $|q(t_0, \tau)|$ were found to lie between the maximal and the minimal values, 1 and $(m_0^{\mu_0})^2$, respectively. For example, for $N = 200$, $\lambda = 1.2$, and $\alpha = 0.15$, the steady-state value of $|q|$ is 0.7, whereas the value of $(m_0^{\mu_0})^2$ is close to 0.35. Thus, the configurations at times t_0 and $(t_0 - \tau)$ are not maximally correlated, but they are more correlated than two random configuration having overlaps $\pm m_0^{\mu_0}$ with the target pattern μ_0. An increase in the storage capacity is, therefore, plausible. However, we should note that the numerical evidence indicating that $\alpha_{\max}(\lambda) > 0.2$ for λ close to unity is not quite conclusive. Longer runs with larger samples would be necessary for an accurate determination of $\alpha_{\max}(\lambda)$.

The observation that $\alpha_{\max}(\lambda)$ may be higher than the maximum storage capacity, α_c, of the Hopfield model does not mean that the model with individual memories stored in limit cycles performs better than the Hopfield model as an associative memory. This is because the quality of retrieval in the present model, measured by the value of the amplitude $m_0^{\mu_0}$, is not very good for values of α close to or higher than α_c. For example, the value of $m_0^{\mu_0}$ for $\lambda = 1.5$, $\alpha = 0.15$ is found to be $\simeq 0.75$ in our simulations whereas in the Hopfield model, the overlap with the target pattern is known to remain close to unity all the way up to α_c. The fact that the model with cycles is able to discriminate between the target pattern and other stored patterns at values of α greater than α_c suggests that this model may be more useful than the Hopfield model as a pattern classifier at relatively high levels of memory loading. In order to investigate this aspect, we carried out a number of simulations in which the state of the network at $t=0$ was taken to be a randomly corrupted version of one of the memory states. The degree of corruption of the input state was measured by the fraction, C, of bits which are different from the corresponding ones in the target memory state. Thus, the overlap, $m^{\mu_0}(t=0)$, of the corrupted input with the target pattern was equal to $(1-2C)$. Simulations were performed for $N=200$, $\lambda=1.5$, $\alpha=0.15$ and 0.2, and several values of C ranging from 0 to 0.3. For each run, the largest Fourier amplitudes $\{g^\mu\}$ were calculated. The input pattern was considered to be classified correctly if the amplitude g^{μ_0} for the overlap with the target pattern was found to be larger than all other g^μs. The results obtained from 200 runs for each set of values of the parameters are shown in table 2 where the fractions of successful runs (i.e. runs in which the input was classified correctly) for various values of the degree of corruption C are given. For comparison, we have also shown in the same table similar results obtained for an $N=200$ Hopfield model. In the Hopfield model simulations, a corrupted input was considered to be classified correctly if the overlap of the final stable state reached by the network with the target pattern was found to be larger than the overlaps with all other stored patterns. It is readily seen from table 2 that the model with cycles performs better than the Hopfield model as a pattern classifier. The improvement is more substantial for $\alpha=0.2$. This is partly due to the fact that this value of α is above the critical value, $\alpha_c \simeq 0.15$, for the Hopfield model, but below the apparent maximum storage capacity of the limit cycle model. The difference between the performance of the two models is most pronounced at somewhat high degrees of corruption, $C \sim 20\%$. This observation suggests that the problem of getting stuck in a spurious attractor, which is present in varying degrees in most neural network models of memory, is less severe in the limit cycle model than in the original Hopfield model. A similar reduction in the probability of the system getting trapped into a spurious attractor was recently observed by Mori *et al* (1989) in a model in which groups of memories are stored in

Table 2. Fraction of runs showing correct classification for different values of the degree of corruption C for a network of 200 neurons embedded with cycles for $\alpha=0.15$ and 0.2. Similar results for a $N=200$ Hopfield model for the same values of C and α are also shown for comparison.

Memory loading level (α)	Degree of corruption C (%)	Fraction of runs showing correct classification (%)	
		Limit cycle	Hopfield model
0.15	10	100	96
	20	97	84
	25	90	76
	30	79	70
0.20	10	100	87
	20	86	69
	25	74	62
	30	62	52

limit cycles. Thus this property appears to be common to all neural network models in which the relevant attractors are limit cycles rather than fixed points. This, we believe, is an important result because the presence of spurious attractors is one of the major difficulties one encounters in neural network modelling. Our results suggest that this problem may be alleviated to some extent by using limit cycles instead of fixed points as relevant attractors. The origin of this effect is not fully understood. According to Mori *et al* (1989), the reduced tendency for trapping in spurious attractors in their model results from a 'dynamic annealing' effect of interchange of patterns belonging to the same cycle. In the present model, the reduction of the probability of getting stuck in spurious attractors is probably caused by the fact that typical spin configurations separated by the time delay are not perfectly correlated. Let us consider the local field acting on the ith spin at time t_0. It is obvious from equation (2) that this local field $h_i(t_0)$ would be simply $(1+\lambda)$ times the local field in the Hopfield model ($\lambda=0$) if the spin configurations $\{\sigma_i(t_0)\}$ and $\{\sigma_i(t_0-\tau)\}$ are maximally correlated (i.e. $\sigma_i(t_0)=-\sigma_i(t_0-\tau)$ for all i). Therefore, if the configuration $\{\sigma_i(t_0)\}$ happens to be a spurious stable state of the Hopfield model (i.e. $\sigma_i(t_0)\sum_j J_{ij}\sigma_j(t_0)>0$ for all i), then it would be a stable configuration of the limit cycle model also. The system would then stay in this configuration for some time $\sim\tau$, and then make a transition to the complement state, thus producing a spurious cycle. If, on the other hand, the configurations $\{\sigma_i(t_0)\}$ and $\{\sigma_i(t_0-\tau)\}$ are, as observed in the simulations, not perfectly correlated, then some of the $\sigma_i(t_0-\tau)$s would be different from the corresponding $-\sigma_i(t_0)$s. As a result, there would be some probability of the state $\{\sigma_i(t_0)\}$ not being a stable one for the $\lambda\neq 0$ model, especially if λ is large and the spurious state being considered is a shallow one. This effect, which is somewhat similar to that of thermal noise in aiding the escape from shallow spurious energy minima in Hopfield-like models, is probably the reason behind the reduction in the number of spurious attractors in the limit cycle model.

5. Summary and conclusions

In this paper, we have used analytic and numerical methods to study the properties of a class of neural network models of associative memory in which a time delay mechanism is used to store individual memories in limit cycles. The main objective was to analyse the general behaviour of memory models which use limit cycle attractors, and to compare the behaviour of such models with that of Hopfield-type models in which the relevant attractors are fixed points. Analytic studies of an asymmetrically diluted version of the limit cycle model shows some improvements in performance over a similarly diluted version of the Hopfield model. Both the storage capacity and the retrieval overlap are found to be higher in the model with limit cycles. Numerical simulations of the fully connected model also indicate an increase in the storage capacity over the Hopfield model. However, the quality of retrieval in this model for values of the memory loading parameter α close to the critical value α_c for the Hopfield model is worse than that in the Hopfield model. The limit cycle model is found to perform better than the Hopfield model as a pattern classifier if the value of α and the degree of corruption of the input pattern are high. We have found some evidence indicating that the problem of convergence to spurious attractors is less severe in the limit cycle model.

Neural networks exhibiting limit cycles have generated a lot of interest among neurobiologists as possible models for central pattern generators, which are groups of specialized neurons responsible for rythmic motor activities such as respiration, beating of the heart and locomotion (Kristan 1980). The models studied in this paper may be relevant in this context (Kleinfeld and Sompolinsky 1988).

References

Amit D J 1989 *Modelling Brain Functions* (Cambridge: Cambridge University Press)

Amit D J, Gutfreund H and Sompolinsky H 1987 *Ann. Phys., NY* **173** 30

Buhmann J and Schulten K 1987 *Europhys. Lett.* **4** 1205

Dehaene S, Changeaux J P and Nadal J P 1987 *Proc. Natl Acad. Sci., USA* **84** 2727

Derrida B, Gardner E and Zippelius A 1987 *Europhys. Lett.* **4** 167

Gutfreund H and Mezard M 1988 *Phys. Rev. Lett.* **61** 235

Hopfield J J 1982 *Proc. Natl Acad. Sci., USA* **79** 2554

—— 1984 *Proc. Natl Acad. Sci., USA* **81** 3088

Kleinfeld D 1986 *Proc. Natl Acad. Sci., USA* **83** 9469

Kleinfeld D and Sompolinsky H 1988 *Biophys. J.* **54** 1039

Kree R and Zippelius A 1987 *Phys. Rev.* A **36** 4421

Kristan W B 1980 *Information Processing in Nervous Systems* ed H M Pinsker and W D Wills (New York: Raven)

Mori Y, Davis P and Nara S 1989 *J. Phys. A: Math. Gen.* **22** L525

Peretto P and Niez J J 1986 *Disordered Systems and Biological Organisation* ed E Bienenstock, F Fogelman-Soulie and G Weisbuch (Berlin: Springer)

Sompolinsky H and Kanter I 1986 *Phys. Rev. Lett.* **57** 2861