International Journal of Bifurcation and Chaos, Vol. 9, No. 4 (1999) 713–722

PERIODIC RESPONSE TO EXTERNAL STIMULATION OF A CHAOTIC NEURAL NETWORK WITH DELAYED FEEDBACK

ALBAN PUI MAN TSUI*
Department of Computing, Imperial College of Science, Technology and Medicine, 180 Queen's Gate, London SW7 2BZ, UK

ANTONIA J. JONES
Department of Computer Science, University of Wales, Cardiff, P.O. Box 916, Cardiff CF2 3XF, Wales, UK

Received June 10, 1998; Revised September 24, 1998

We construct a feedforward neural network so that when the outputs are fed back into the inputs and the system is iterated it behaves chaotically. We call this the "rest state". Suppose now that an input stimulus is added to one or more inputs. Following a biologically inspired model suggested by Freeman [1991], under these conditions we should want the behavior of the network to stabilize into an unstable periodic orbit of the original system. We call this the "retrieval behavior" since it is analogous to the act of recognition. Standard methods of chaos control, such as OGY for example, used to elicit the retrieval behavior would be inappropriate, since such methods involve calculations external to the system being controlled and can be considered unlikely in a biological neural network. Using a chaos control method originally suggested by Pyragas [1992] we show that retrieval behavior can occur as a result of delayed feedback and examine the variety of the responses that arise under different types of stimuli and under noise. This artificial neural system has a strong dynamical parallel to Freeman's observed biological phenomenon.

1. Introduction

On the basis of studies of the olfactory bulb of a rabbit Freeman [1991] suggested that in the rest state the dynamics of this neural cluster is chaotic but that when a familiar scent is presented the neural system rapidly simplifies its behavior and the dynamics becomes more orderly, more nearly periodic than when in the rest state. This suggests an interesting model of recognition in biological neural systems. To realize this in an artificial neural system, some form of control of the chaotic neural behavior is necessary to achieve periodic dynamical behavior, from the normally chaotic behavior, when a stimulus is presented.

There is now an extensive literature demonstrating experiments on controlling chaotic physical systems using the original chaos control techniques, such as the OGY method [Ott *et al.*, 1990] or its similar variants. Many such methods require careful and systematic analysis of the chaotic dynamical behavior, which is usually difficult and computationally expensive, before successful control can be achieved. Moreover, such control techniques are *external* to the system being controlled, whereas for a neural system to behave as described by

*E-mail: a.tsui@ic.ac.uk

Freeman [1991] the control should be *intrinsic* to the neural dynamics.

Therefore for our iterative neural model, we implement a much simpler control method based on delayed feedback and similar to Pyragas' [1992] original continuous delayed feedback technique. One of the attractions of this method is that it has a very low computational overhead and so is extremely easy to implement in hardware. It would also be very easy to implement in biological neural circuitry and so offers one plausible mechanism whereby such stabilization might occur.

We first produce a chaotic neural system by training a feedforward neural network on a chaotic iterative map. The map used here is in fact the Ikeda map. However, the precise nature of the system used to produce the training data is largely immaterial, as is the precise accuracy of the neural network model. All that is required is that when outputs are fed back into inputs and the trained neural network is iterated it exhibits a sufficiently rich chaotic dynamical behavior.

Delayed feedback is then introduced into the model and this provides a mechanism for stabilization onto unstable periodic behaviors. The particular unstable periodic orbit which is stabilized depends quite strongly on the precise character of the applied stimulus. Thus the system can act as an associative memory in which the act of recognition corresponds to stabilizing onto an unstable periodic orbit which is characteristic of the applied stimulus. The entire artificial system therefore exhibits an overall behavior and response to stimulus which precisely parallels the biological neural behavior observed by Freeman.

2. A Trained Chaotic Neural Net

In our experiment, the chaotic neural net is simply a standard feedforward neural network, with 2 inputs, 2 output neurons and two hidden layers with 10 neurons each (i.e. 2-10-10-2) trained on the input/output data from the chaotic Ikeda map [Hammel *et al.*, 1985] as in [Tsui & Jones, 1997]. The map is defined by

$$z_{n+1} = g(z_n) = \gamma + Rze^{\mathbf{i}}\left(\kappa - \frac{\alpha}{1+|z|^2}\right) \qquad (1)$$

where $\alpha = 5.5$, $\gamma = 0.85$, $\kappa = 0.4$, $R = 0.9$ and z is a complex variable of the form $x + y\mathbf{i}$ and z can also be identified with the point (x, y) in the $\mathbb{R}^2$ plane. This chaotic attractor is shown in Fig. 1. With only 4000 input/output data pairs $\{(x_n, y_n), (x_{n+1}, y_{n+1})\}$ (re-scaled into the range $[0, 1]$) we can train the feedforward neural network, with the described architecture and sigmoid neuron

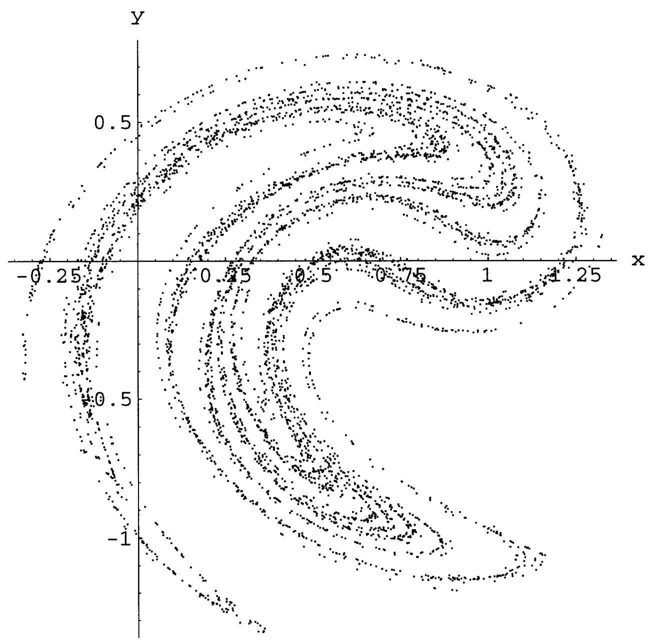

Fig. 1. The Ikeda strange attractor, for parameter values $\alpha = 5.5$, $\gamma = 0.85$, $\kappa = 0.4$ and $R = 0.9$.

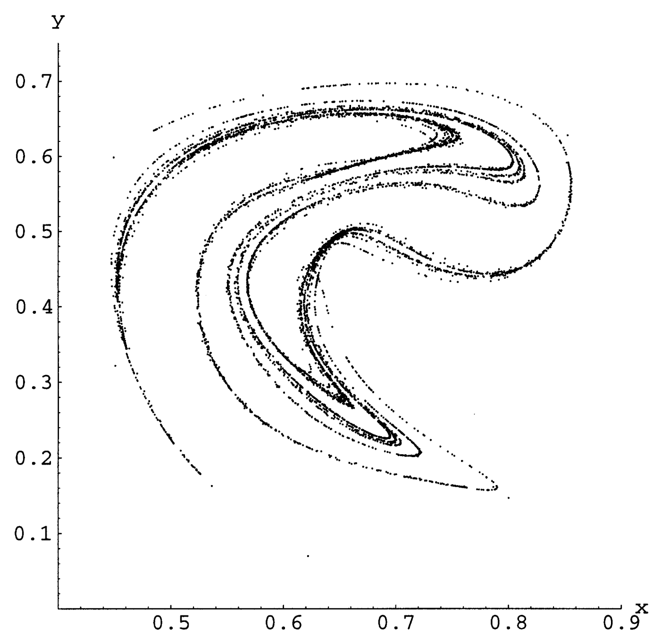

Fig. 2. Attractor for chaotic network with architecture 2-10-10-2 (with inputs and outputs re-scaled to $[0, 1]$).

with transfer function

$$\mathrm{act}(x) = \frac{1}{1 + e^{-x/T}} \qquad (2)$$

where $T = 1$, to capture the essential features of the chaotic attractor of the Ikeda map knowing the fact that a feedforward network can learn dynamical behavior [Welstead, 1991; Dracopoulos & Jones, 1993]. The recursive behavior can simply be achieved by feeding the outputs into the inputs, output x is fed back into input x and output y is fed back into input y. With the training MSE error of about 9.9×10^{-5}, this network already produces a chaotic attractor similar to the original Ikeda map, see Fig. 2. The Lyapunov exponents are $\{0.368973, -0.769616\}$ as opposed to $\{0.403642, -0.614225\}$ for the original Ikeda map from which the training data is drawn.

3. Controlling Chaos with Delayed Feedback

To control our chaotic neural network, we implement a discrete version of a delayed feedback technique not too different from the Pyragas' original continuous form [Pyragas, 1992]. The Pyragas method for continuous-time control applies to a chaotic system which can be represented by a set of ordinary differential equations

$$\frac{dy}{dt} = P(y, \mathbf{x}) + F(t)\,, \quad \frac{d\mathbf{x}}{dt} = \mathbf{Q}(y, \mathbf{x})\,. \qquad (3)$$

Here y is the output variable and the vector $\mathbf{x}$ describes the remaining variables of the dynamic system which are not available. The feedback control signal $F(t)$ disturbs only the first equation, corresponding to the output variable. We suppose that without an input signal the system being considered has a strange attractor.

The idea behind this method is to construct the feedback control signal $F(t)$ in such a way that it vanishes when the system moves along the desired unstable periodic orbit. One approach suggested by Pyragas [1992] is to use

$$F(t) = -k[y(t) - y(t - \tau)] = -kD(t)\,, \qquad (4)$$

where $k > 0$. Here τ is a delay time and $y(t - \tau)$ is the delayed system state. Therefore the magnitude of the control signal is proportional to the difference $D(t) = y(t) - y(t - \tau)$. If this time τ coincides with the period of the unstable periodic orbit $\tau = \tau_0$ then the perturbation becomes zero for the solution of the system (3), i.e. $y(t) = y(t - \tau_0)$.

Stabilization of the system can often be accomplished by choosing an appropriate weight k so that a negative feedback is achieved. Though Qu [Qu *et al.*, 1993] argued that in some cases, a positive feedback is needed. Therefore there are two variables, k and τ that can be adjusted in the experiment. The delay τ is expected to be the period of the stabilized orbit from the controlled chaotic system if the system eventually stabilizes. Some experimental results can be found in [Pyragas, 1993; Pyragas & Tamaševičius, 1993; Celka, 1994; Cooper & Schöll, 1995; Tsui & Jones, 1998].

4. Delayed Feedback Applied to the Chaotic Neural Net

A simple delayed feedback, similar to the Pyragas' delayed feedback, can be added to the chaotic neural net to control the chaotic behavior with a careful choice of the parameters k and τ. The basic control setup of the neural model is shown in Fig. 3. Here the trained chaotic feedforward neural net described earlier is now equipped with extra delayed feedback control circuitry, which is activated on presentation of an external stimulus. The delayed feedback is added to the state variable y_n to effect the control. External stimulation is performed by feeding signals into the input line x_n of the network. Let FF be the feedforward network mapping such that $FF[(x_n, y_n)] = (x_{n+1}, y_{n+1})$ then the controlled system with external stimulation s_n at time n is described by

$$(x_{n+1}, y_{n+1}) = FF[(x_n + s_n, y_n + p_n)] \qquad (5)$$

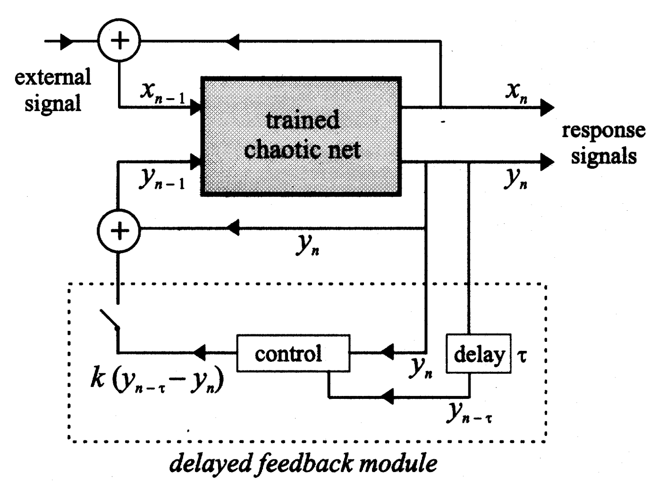

Fig. 3. Delayed feedback on chaotic neural net.

where $p_n = k(y_{n-\tau} - y_n)$ is the delayed feedback control signal.

After some initial investigation we fixed $k = 0.5$ and $\tau = 6$ for the experiments. These values stabilized the system with control switched on but with no external stimulus present. Other values of k and τ can also stabilize the system (x_n, y_n) successfully.

We imagine that the presence of an external stimulus excites (activates) the control circuitry which is otherwise inhibited. Thus to achieve a stabilized dynamical regime in response to a stimulus the control is switched on at the same time as the external signal is fed into the input line x_n. By varying the external signal in small steps and holding the new setting fixed long enough for the system to stabilize we can observe the response of the network to small changes in stimulus.

In Fig. 4 the system is iterated for 100 cycles to eliminate any initial transients. Next an external constant stimulus $s_n = s$ is applied for 400 steps. In Fig. 4 the stimulus s is varied in steps of 0.025 over the interval $[0, 1]$ for every 400 network iterations. We can see that the system exhibits a fairly "smooth" transition of stabilized behavior from one stimulus to the next. For the most part in this case the response is a 1-period behavior but a 2-period behavior is also exhibited after the strength of the external signal crosses a threshold at around 0.8 which is therefore a *bifurcation point*. For a stimulus $s_n = s$ with $s > 0.2$ the delayed feedback control signal quickly becomes small, which indicates that the system has stabilized onto one of its own unstable periodic behaviors. However, for a stimulus $s_n = s$ with $s < 0.2$ a *large* feedback control signal p_n often seems to create some new periodic behavior.

We can study the response of the system as the stimulus 0.2 is applied and removed and as control is turned on and off. This is shown in Fig. 5. In general terms the system stabilizes after about 50 iterations. If the stimulus is applied without control the dynamical regime seems not to correspond to an unstable periodic behavior of the original network, but with control switched on the dynamics quickly stabilizes to a 1-period corresponding to an unstable periodic behavior of the iterated network.

Note that in the transition sc $\to$ s in Fig. 5, in which the control is removed but the stimulus

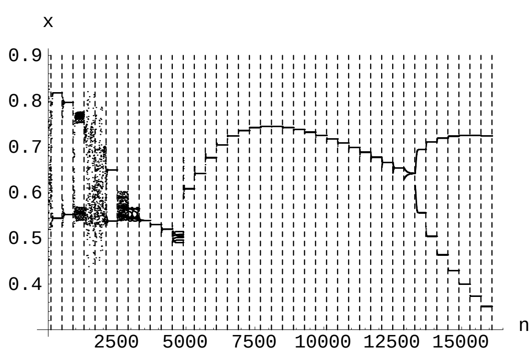

(a) Response of x_n.

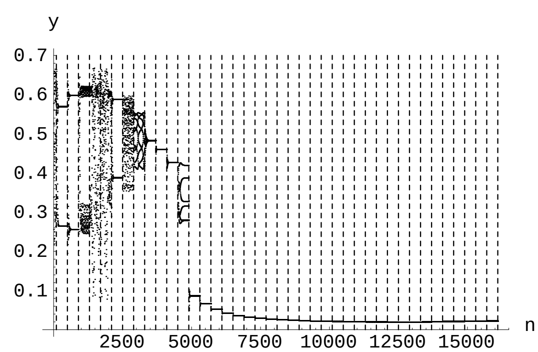

(b) Response of y_n.

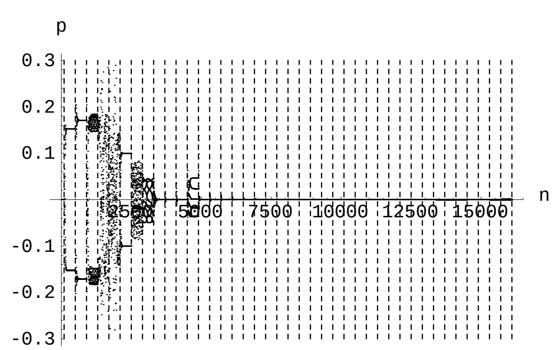

(c) Size of delayed feedback control signal p_n.

Fig. 4. Results of external constant stimulation $s \in [0, 1]$ varying in steps of 0.025. The stimulus changes at 400 iteration steps after an initial 100 iterations to eliminate transients. The control parameters were $k = 0.5$ and $\tau = 6$.

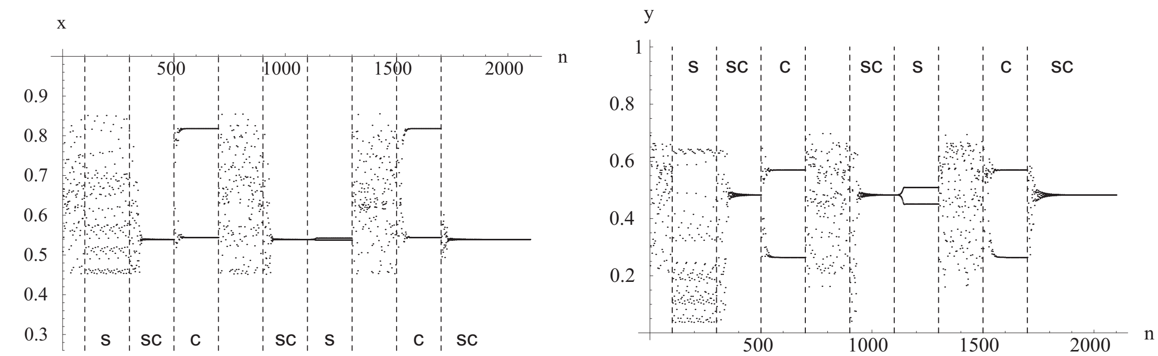

Fig. 5. Responses of x_n and y_n to presentation and removal of stimulus $s = 0.2$ with and without control. Intervals labeled "s" indicate the presence of the stimulus, intervals labelled "c" indicate control is switched on, a label "sc" indicates both, and no label indicates no stimulus and no control. The particular regime is changed every 200 iterations after 100 iterations have been allowed for transient removal. $k = 0.5$ and $\tau = 6$.

remains, surprisingly the system shifts from a 1-period to a 2-period, rather than reverting to the more chaotic regime illustrated in the first 400 step interval, where the same stimulus without control proved unable to stabilize the system.

In some cases, the external stimulation signal is enough to stabilize the system without switching on the control module. The explanation of this might be that when such an external signal is strong enough, or it is a particular kind of signal, it may shift the underlying dynamics from a chaotic region into a periodic region in the bifurcation diagrams, as shown in Fig. 6. This figure originally appeared in [Tsui & Jones, 1997] which studied the same feedforward neural network. A similar shift from a chaotic behavior to a more stabilized retrieval behavior is also observed in a large scale model with structured synaptic connections in [Adachi & Aihara, 1997].

Apart from a constant external stimulation signal applied to one of the inputs other forms of s_n can also be used. Low period square waves can also result in stabilized periodic responses as shown in Fig. 7.

A completely different way of applying a stimulus was suggested in [Hoff, 1994]. The stimulus can be applied directly to the control variable k. In this way different behaviors can be achieved by using the external signal s_n to directly modify k. Some results of this type of control applied to our system are illustrated in Fig. 8.

These experiments are merely illustrative and many variations are possible. For example, delayed feedback could equally be applied to several (or all) of the network outputs. With the same τ and multiple feedbacks it should be easier to achieve stabilization compared to the case where feedback is applied to just one variable. However, if delayed feedback on different network outputs also had differing τ then the outcome is less predictable. There remain many possibilities for exploring this type of neural model.

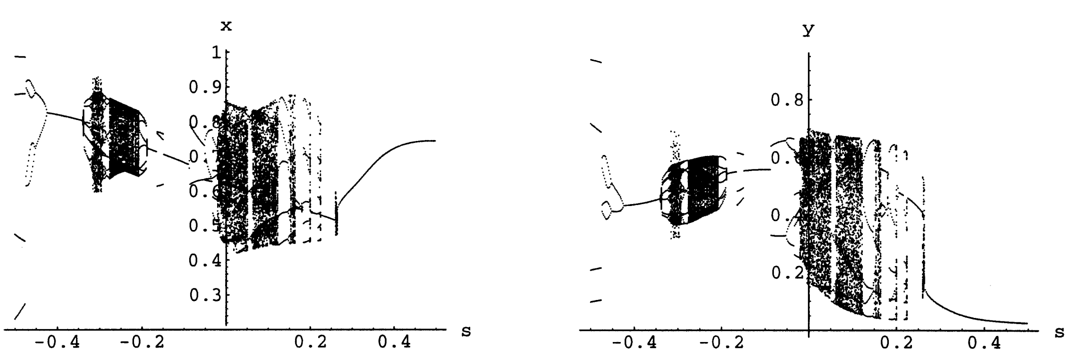

Fig. 6. Bifurcation diagrams for the outputs x_{n+1} and y_{n+1} using an external variable s added to the input x_n.

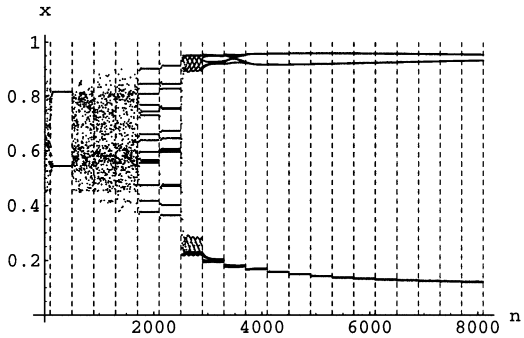

(a) Response of x_n.

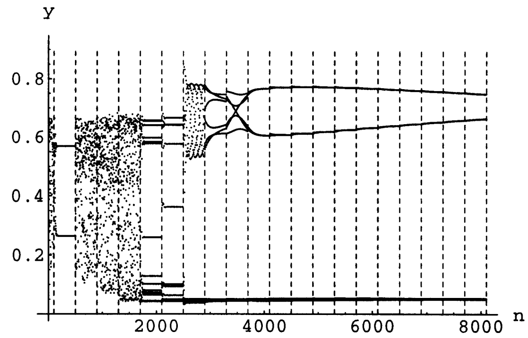

(b) Response of y_n.

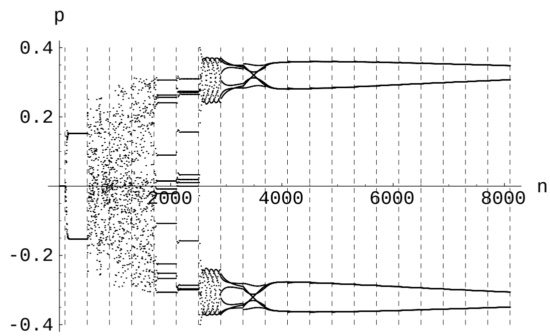

(c) Size of delayed feedback control signal p_n.

Fig. 7. Results of periodic stimulation $s_n \in \{j,\ 0,\ j,\ 0, \ldots\}$ of strength j from 0 to 1 in steps of 0.05. The stimulus changes at 400 iteration steps after an initial 100 iterations to eliminate transients. The control parameters were $k = 0.5$ and $\tau = 6$.

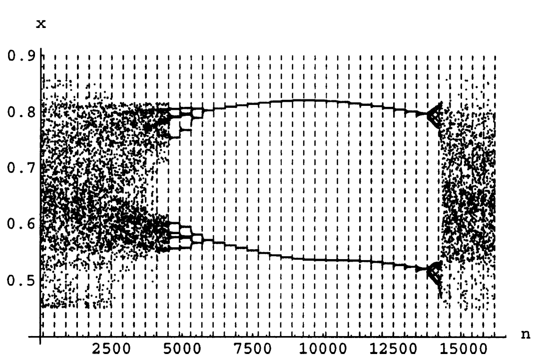

(a) Response of x_n.

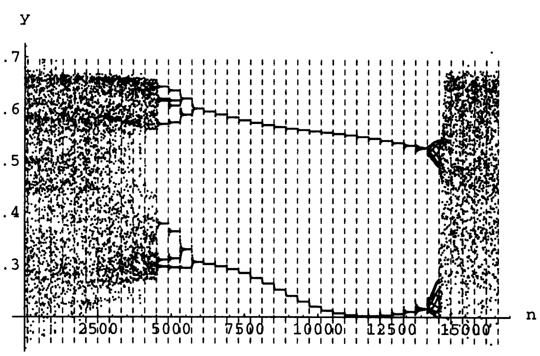

(b) Response of y_n.

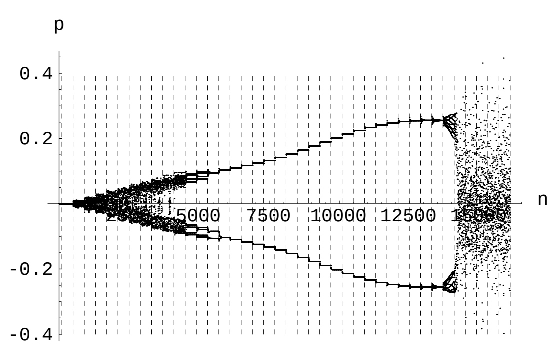

(c) Size of delayed feedback control signal p_n.

Fig. 8. Results of external signal $s_n = s$ added to the value k from -0.5 to 0.5 in steps of 0.025. The stimulus changes every 400 network iterations after 100 initial iterations with no stimulus and no control. $k = 0.5$ and $\tau = 6$.

We also investigated the response of the system when sensory input was perturbed by stochastic noise. The stimulus was perturbed at each iteration step by multiplying it by Gaussian noise with a mean of 1 and a variance σ, where σ varied from $\sigma = 0$ to $\sigma = 0.1$. The response was surprisingly robust as illustrated in Figs. 9–11. These results should be compared with the non-noisy case in Fig. 4. The noisy dynamics remain essentially unchanged, although as one might expect the attractor becomes progressively "blurred" as the noise level increases.

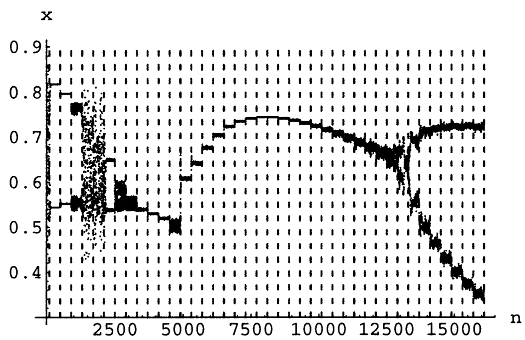

Fig. 9. The response of the system to noise. The stimulus s_n is replaced by $s_n r$ at each iteration step, where r is Gaussian noise with mean 1 and variance $\mathrm{Var}(r) = 0.005$. The stimulus changes at 400 iteration steps after an initial 100 iterations to eliminate transients. The control parameters were $k = 0.5$ and $\tau = 6$.

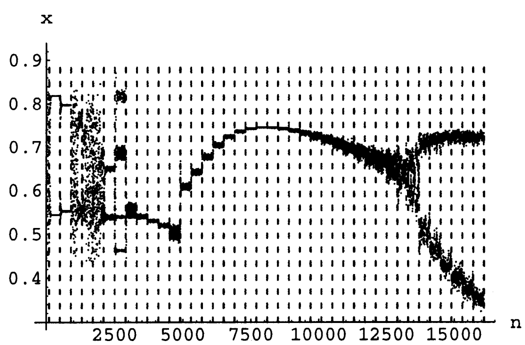

Fig. 10. The response of the system to noise. The stimulus s_n is replaced by $s_n r$ at each iteration step, where r is Gaussian noise with mean 1 and variance $\mathrm{Var}(r) = 0.01$. The stimulus changes at 400 iteration steps after an initial 100 iterations to eliminate transients. The control parameters were $k = 0.5$ and $\tau = 6$.

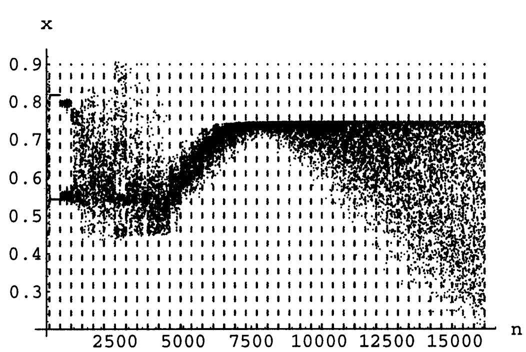

Fig. 11. The response of the system to noise. The stimulus s_n is replaced by $s_n r$ at each iteration step, where r is Gaussian noise with mean 1 and variance $\mathrm{Var}(r) = 0.1$. The stimulus changes at 400 iteration steps after an initial 100 iterations to eliminate transients. The control parameters were $k = 0.5$ and $\tau = 6$.

5. Local Stability Analysis

Little theoretical analysis is available for the Pyragas method of continuous delayed feedback control, let alone for the discrete form of the method used here. However, a discrete version of a variation of Pyragas' method has already been successfully applied to the synchronization of two identical iterative chaotic maps in [Oliveira & Jones, 1998]. The version used there for *synchronization* is similar to but not identical to the method used here for *stabilization*. Oliveira and Jones [1998] also contained a suggestive discussion of the local stability properties of the method used. For both the Hénon map and the chaotic neural network used here it was shown that whilst the synchronization control method used by Oliveira and Jones [1998] was not locally stable it was nevertheless *probabilistically locally stable*.

We next try to provide a similar empirical analysis for the method of stabilization proposed here in the case where no external stimulus is present. First, we note the stabilized state when control is switched on with $k = 0.5$, $\tau = 6$ and no external stimulus is applied. This gives a 2-period controlled behavior $\{\xi_{F1}, \xi_{F2}\} = \{(0.81808, 0.569261), (0.543838, 0.264166)\}$.

If we define a measure of contraction

$$\mu_n = \frac{\min(|\xi_{n+1} - \xi_{F1}|, |\xi_{n+1} - \xi_{F2}|)}{\min(|\xi_n - \xi_{F1}|, |\xi_n - \xi_{F2}|)} \tag{6}$$

towards $\{\xi_{F1}, \xi_{F2}\}$ from step n to step $n+1$ then μ_n depends on the eigenvalues of the Jacobian of the associated four-dimensional system $\{\xi_n, \xi_{n+1}\}$ in the vicinity of $\{\xi_{F1}, \xi_{F2}\}$ and these (although bounded) can be much larger than 1. Thus it is simply not true that with this control method the system will monotonically approach the unstable periodic behavior. However, if we examine the effects of control after several iterations we find that the *probability* that the cumulative net contraction becomes small is very large.

To establish this we generate a random initial point ξ_0 and iterate the controlled system. At the nth iteration we define

$$\rho_n = \frac{\min(|\xi_n - \xi_{F1}|, |\xi_n - \xi_{F2}|)}{\min(|\xi_0 - \xi_{F1}|, |\xi_0 - \xi_{F2}|)}\,. \tag{7}$$

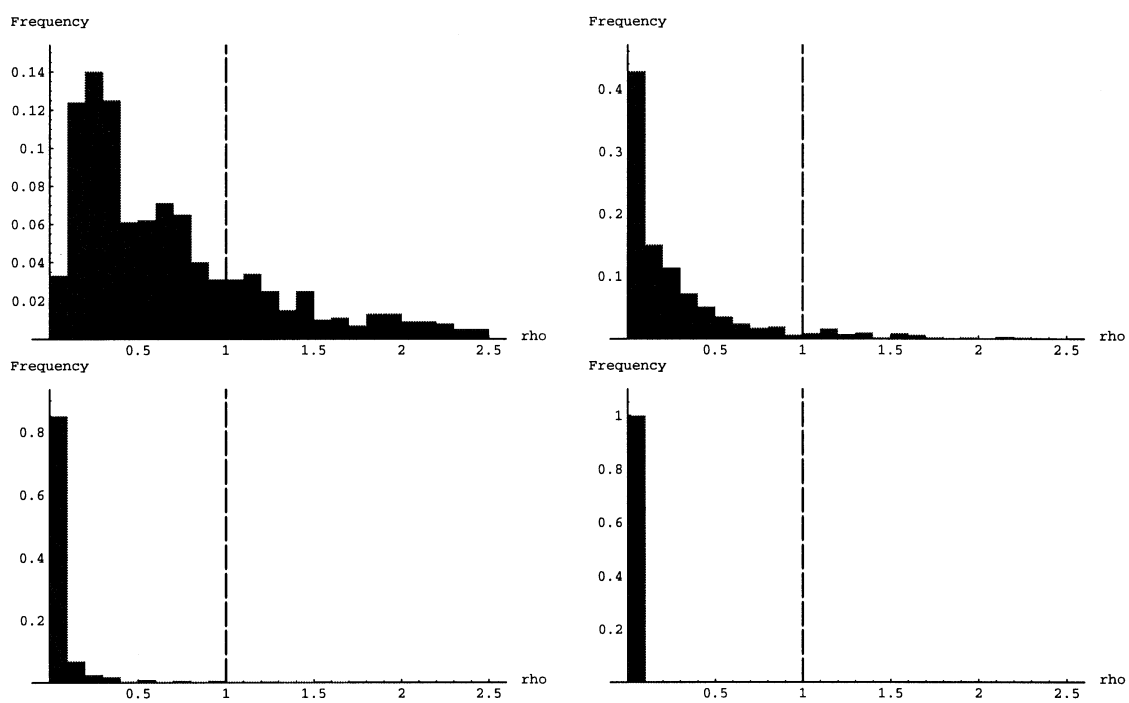

Fig. 12. Histograms of ρ_n at $n = 20$ (top left), 40 (top right), 60 (bottom left), 80 (bottom right) of 1000 random initial starting points for control $k = 0.5$ and $\tau = 6$.

The quantity ρ_n gives us a measure of the extent to which after n iterations with control the system has *contracted* towards the unstable 2-period.

By showing that ρ_n becomes small with high probability, i.e. $\rho_n \to 0$ as $n \to \infty$, where the convergence is in probability, we can establish that the method is *probabilistically locally stable.*

We repeated the calculation of ρ_n for 1000 different initial starting points and $n \leq 80$ and create histograms showing the frequency of ρ_n against the value. These results are shown in Fig. 12. These histograms suggest that $\forall \varepsilon > 0$,

$$P[|\rho_n| < \varepsilon] \to 1 \tag{8}$$

as $n \to \infty$. Thus the system without stimulus is probabilistically locally stable.

The application of an external stimulus basically modifies the system dynamics by *shifting* the dynamic behavior along the bifurcation diagrams as mentioned earlier. Many new chaotic and nonchaotic behaviors are produced by the neural system which are different from its initial built-in dynamics without stimulation. Thus the delayed feedback control seems to act as a supporting tool for stabilizing the system into periodic states. Although there is insufficient theoretical explanation for the dynamical behavior of our neural system, the above heuristic analysis seems to fit very well with the observed simulation results.

6. Conclusions

We have shown how a conventional artificial feedforward neural network equipped with delayed feedback can simulate the type of rest behavior and response to stimuli observed by Freeman in the olfactory bulb of the rabbit. The system is in effect an associative memory in which the act of recognition corresponds to the stabilization of the system onto an unstable periodic orbit characteristic of the applied stimulus.

If the dynamics is chaotic then unstable periodic orbits are dense on the chaotic attractor and there are infinitely many of them. Thus such an associative memory for which the computations are performed to an *arbitrary* precision could in principle accommodate infinitely many memories; at any rate such a system is not subject to the conventional Hopfield upper bound of $0.15n$, where n is the number of neurons [Amit *et al.*, 1987]. Of course, for the Hopfield net the situation is rather different. In the Hopfield model memories are associated with specified (preferably uncorrelated) point attractors, whereas in the present model memories are associated with unstable periodic behaviors which cannot be specified *ab initio.* This introduces the possibility of responding to stimuli over varying *time scales.*

Our experiments were based on high precision digital simulations. In a low arithmetical precision analog implementation it is possible that much of the rich variety of dynamical behavior would be lost.

Nevertheless, the model has a certain compelling simplicity which is suggestive. The responses described are *intrinsic* to the network model and control is not artificially applied from outside the network itself. The method of delayed feedback is simple to apply in hardware and feasible in biological neural circuitry.

As with the many applications of the method of Pyragas to control more conventional chaotic dynamics our approach lacks a full formal analysis. However, we have investigated the local stability properties of the method applied to the particular model described here and have concluded that although control is not stable in the conventional sense it is nevertheless *probabilistic locally stable.*

The experiments described here raise several interesting issues. An investigation of essentially the same model could be performed with delayed differential equations using a more biologically accurate description of the neurons. In [Tsui & Jones, 1998] we describe delayed feedback control applied to the stabilization of a six-dimensional *smooth* dynamical system and this illustrates that the ideas described here could quite probably be applied successfully to a similar model based on differential equations.

Another question which naturally arises is whether "the basin of attraction" of a particular unstable periodic orbit, which has emerged as the response to a specific stimulus, could be "widened" by repeated presentations using some form of weight adjustment based on Hebbian learning. The critical aspect to investigate here would be whether this could be done without destroying the essential underlying chaotic dynamics or other conditioned responses.

The periodic responses exhibited are common in coupled oscillator models (e.g. [Stewart, 1992]) which are very different from the model described here. It is therefore interesting to note that, by incorporating delayed feedback, periodic neural responses can be achieved with an essentially conventional feedforward neural network model without the introduction of an oscillator neuron.

Acknowledgments

The author Alban Pui Man Tsui is supported by EPSRC, UK, Studentship number: 96307552.

References

Adachi, M. & Aihara, K. [1997] "Associative dynamics in a chaotic neural network," *Neural Networks* **10**(1), 83–98.

Amit, D., Gutfreund, H. & Sompolinsky, H. [1987] "Information storage in neural networks with low levels of activity," *Phys. Rev.* **A35**, 2239–2303.

Celka, P. [1994] "Experimental verification of Pyragas's chaos control method applied to Chua's circuit," *Int. J. Bifurcation and Chaos* **4**(6), 1703–1706.

Cooper, D. & Schöll, E. [1995] "Tunable real space transfer oscillator by delayed feedback control of chaos," *Zeitschrift für Naturforschung, Section A – A J. Phys. Sci.* **50a**, 117–124.

Dracopoulos, D. C. & Jones, A. J. [1993] "Neuromodels of analytic dynamic systems," *Neural Comput. Appl.* **1**(4), 268–279.

Freeman, W. [1991] "The physiology of perception," *Sci. Am.* Feb, 34–41.

Hammel, S., Jones, C. & Moloney, J. [1985] "Global dynamical behaviour of the optical field in a ring cavity," *J. Opt. Soc. Am.* **B2**(4), 552–564.

Hoff, A. A. [1994] "Chaos control and neural classification," *Zeitschrift für Naturforschung, A, Physik Physikalische Chemie Komophysik* **49**(a), 589–593.

Oliveira, A. G. d. & Jones, A. J. [1998] "Synchronisation of chaotic maps by feedback control and application to secure communications using neural networks," *Int. J. Bifurcation and Chaos* **8**(11), 2225–2237.

Ott, E., Gerbogi, C. & Yorke, J. [1990] "Controlling chaos," *Phys. Rev. Lett.* **64**(11), 1196–1199.

Pyragas, K. [1992] "Continuous control of chaos by self-controlling feedback," *Phys. Lett.* **A170**, 421–428.

Pyragas, K. [1993] "Predictable chaos in slightly perturbed unpredictable chaotic systems," *Phys. Lett.* **A181**, 203–210.

Pyragas, K. & Tamaševičius, A. [1993] "Experimental control of chaos by delayed self-controlling feedback," *Phys. Lett.* **A180**, 99–102.

Qu, Z., Hu, G. & Ma, B. [1993] "Controlling chaos via continuous feedback," *Phys. Lett.* **A178**, 265–270.

Stewart, I. [1992] "Introduction to nonlinear oscillator," in *Coupled Oscillating Neurons, Perspectives in Neural Computing*, eds. Taylor, J. & Mannion, C. (Springer-Verlag), pp. 1–20.

Tsui, A. P. & Jones, A. J. [1997] "Parameter choices for control of a chaotic neural network," in *Advances in Intelligent Systems, Frontiers in Artificial Intelligence and Applications*, Vol. 41, ed. Morabito, F. (University of Reggio Calabria, Italy. AMSE, SIGEF and BUFSA, IOS Press), pp. 118–123.

Tsui, A. P. & Jones, A. J. [1998] "The control of higher dimensional chaos: Comparative results for the chaotic satellite attitude control problem," *Physica D*, to appear.

Welstead, S. [1991] "Multilayer feedforward networks can learn strange attractors," *Proc. IIJCNN-91-Seattle*, Vol. II, pp. 139–144.