

Homeostatic plasticity improves signal propagation in continuous-time recurrent neural networks

Hywel Williams*and Jason Noble

{hywelw|jasonn}@comp.leeds.ac.uk

School of Computing, University of Leeds, Leeds LS2 9JT, UK

February 28, 2005

Abstract

Continuous-time recurrent neural networks (CTRNNs) are potentially an excellent substrate for the generation of adaptive behaviour in artificial autonomous agents. However, node saturation effects in these networks can leave them insensitive to input and stop signals from propagating. Node saturation is similar to the problems of hyper-excitation and quiescence in biological nervous systems, which are thought to be avoided through the existence of homeostatic plastic mechanisms. Analogous mechanisms are implemented in CTRNNs and are shown to increase node sensitivity and improve signal propagation. These results are found in a variety of network architectures and lend support to the view that homeostatic plasticity may prevent quiescence and hyper-excitation in biological nervous systems.

Keywords: continuous-time recurrent neural network, homeostatic plasticity, signal propagation

1 Introduction

The readily apparent success of neural control systems in the biological world has led to a significant amount of research into the use of neural networks for controlling artificial autonomous agents. Neural networks can offer a number of tempting properties for robot control, such as robustness and nonlinear dynamics, but these benefits can only be realised if the network also provides the

*Corresponding author. Tel. +44 (0) 113 243 1751 ext 36804

basic functionality necessary for all robot controllers. One key function that every controller must offer is that of signal propagation; any autonomous agent acting in a real-world environment must turn sensory input into motor output if it is to behave usefully, and this means that signals must propagate from sensors to effectors.

The most common types of neural controller used in the evolutionary robotics community are variants of the continuous-time recurrent neural network (CTRNN), the basic form of which is exemplified by Beer (1995). These networks are intended to be loosely analogous to biological neural networks and involve neurons that fire at a rate determined by a sigmoid function of their potential. The sigmoidal transfer function is biologically plausible, but leads to the problem of node saturation. Node saturation occurs when the range of inputs a neuron receives is either too high or too low, resulting in hyper-excitation or quiescence. Saturated nodes do not change their activity in response to stimulus, and thus are useless for performing any form of computation. Saturated nodes make for inert robot control networks, where signals do not propagate and no computation is performed.

The focus of this paper is improving signal propagation in CTRNNs by using a novel mechanism, homeostatic plasticity, to solve the saturation problem. Homeostatic plasticity is widely observed in biological nervous systems, where one of its functions is thought to be the avoidance of quiescence and hyper-excitation at the level of individual neurons, but is relatively unexplored in artificial systems. The inclusion of homeostatic plasticity in CTRNNs should allow the saturation problem to be overcome and therefore allow signals to propagate more effectively.

2 Background

2.1 Continuous-time recurrent neural networks for robot control

CTRNNs are popular within the neural robotics community because they offer not only the general benefits of neural networks (such as distributed information processing and robustness) but also more specific benefits such as smooth nonlinear dynamics, the capability to maintain autonomous oscillations, and the ability to approximate the output of any dynamical system when correctly parameterised (Funahashi & Nakamura, 1993). However, parameterisation of these networks is difficult; CTRNNs have a large number of sensitive parameters and are not amenable to traditional

neural network training methods such as back-propagation of error because of their highly recurrent nature. The current best method of parameterisation is the use of genetic algorithms to evolve good networks against some task-based fitness criterion. Harvey, Husbands, Cliff, Thompson, and Jakobi (1997) and Beer (1996) are good examples of this methodology.

2.2 Signal propagation

Reacting to the environment requires that an agent be possessed of sensors and effectors with an effective link between them. Signals must be transduced in some form from input to output. In neural network terms, sensor nodes must communicate with motor nodes. If they do not, then there is no link between sensory information and action, and the agent cannot respond to a changing environment in any meaningful way. Any action taken would therefore be solely a function of internal state and would make no reference to the outside world. Behaviour of this sort is unlikely to be adaptive, since it will not allow opportunities to be exploited or threats avoided. Such a robot will not be able to perform any useful tasks.

At a very basic level, a signal can be said to propagate if a change in state at the transmission end of a channel causes a change in state at the receiving end. In CTRNN terms, a pertinent change in sensory input should cause a change in motor output. For this to occur, each individual node along the path between sensors and effectors must display some change in state as a result of input. This conception of signal propagation is rudimentary and does not consider the information content of a signal or transmission accuracy. However, the possibility of a change in state at the receiving end of a channel is a pre-requisite for any information transmission whatsoever, and we will adopt this primitive metric as it is suitable for the current purpose.

2.3 Node saturation

The firing rate of a CTRNN node is a sigmoidal function of its potential (which is in turn a function of itself and any input the node receives). Figure 1 shows the relationship between firing rate z and potential y for a neuron. For any neuron, there will be some range of inputs that are typically received, and this means that the potential of that neuron will fluctuate within a given range. The ranges A, B and C in figure 1 show three possible ranges of variation of potential. Consider what happens to the firing rate of a neuron whose potential varies within each range. In range A, the





Figure 1: Sigmoidal transfer function: firing rate as a function of potential, showing different possible ranges of habitual fluctuation.

firing rate hardly changes because the slope of the sigmoid in this range is negligible. The neuron fires at a near-constant rate close to zero. Similarly, in range C the neuron fires at a near-constant rate close to 1, for the same reason. Only in range B does a change in potential cause a noticeable change in firing rate.

Figure 1 illustrates the *saturation problem*: neurons saturate on or off if their inputs are too high or too low. When neurons are saturated, their firing rate does not change in response to a change in input. In the context of a network, this means that they cannot propagate a signal and do not play a part in network dynamics. Consider what would happen if node saturation occurs at sensory or motor nodes, or at nodes on the path in between: signals do not propagate and the connection between sensory input and motor output is broken. The network, and hence the agent, cannot respond to environmental stimuli.

2.4 Homeostatic neural plasticity

It is useful to consider biological neural networks and to observe that saturation effects (i.e., hyperexcitation and quiescence) are not a common problem. One postulated reason for this is the existence of homeostatic plastic mechanisms that serve to regulate neural activity (Turrigiano, 1999; Davis & Bezprozvanny, 2001). While the precise feature of neural activity that is regulated is not known (it may be mean firing rate, mean calcium concentration or some other feature) it is clear that neural activity tends towards a constant level in the long term. It is also clear that there are a variety of mechanisms by which this homeostasis is accomplished, amongst which are mechanisms affecting the strength of synaptic connections (Turrigiano, 1999; Abbott & Nelson, 2000; Davis & Bezprozvanny, 2001; Burrone & Murthy, 2003) and the intrinsic excitability of individual neurons



(a) Synaptic scaling affects the gain of a neuron, effectively changing the slope of the input-output function.

(b) Internal plasticity can affect the intrinsic excitability of a neuron, effectively translating the input-output function.

Figure 2: Different kinds of homeostatic neural plasticity.

(Desai, Rutherford, & Turrigiano, 1999; Zhang & Linden, 2003). Figure 2 illustrates the effect of these two different types of mechanism on the transfer function of a hypothetical neuron.

Analogous mechanisms have been adapted for use in CTRNNs and have been shown to make CTRNNs more sensitive and give rise to more interesting behaviours when used to control autonomous agents (Williams, 2004). The effect of homeostatic plasticity on CTRNNs as a substrate for behaviour has also been examined by Williams, Buckley, and Noble (2005).

3 Method

3.1 CTRNN

The CTRNNs used in these experiments are of the standard form exemplified by Beer (1995). The neuron state equation is given by equation 1 with a sigmoidal transfer function given by equation 2.

$$\tau_y \dot{y} = -y + \sum_{i=1}^N w_i z_i + I \quad (1)$$

$$z = \frac{1}{1 + e^{-(y+b)}} \quad (2)$$

By analogy with biological neurons, equations 1 and 2 represent the state of a node with N afferent synapses (a self-connection and connections from $N - 1$ other nodes), where y represents

neuron potential, w_i is the strength of the synapse from the i^{th} afferent neuron, z_i is the firing rate of the i^{th} afferent neuron, I is any external input the neuron receives, and b is the bias term for the neuron. Equation 1 defines the rate of change of potential with respect to time (\dot{y}) moderated by a neuron specific time constant (τ_y). Equation 2 specifies neuron firing rate z as a sigmoid function of neuron potential and bias. Weights can take positive or negative values, representing excitatory and inhibitory synapses. Biases can also be positive or negative, reflecting a neuron’s inherent tendency towards quiescence or excitation. Ranges for different parameter values will be given below where appropriate.

3.2 Homeostatic plastic mechanisms

Homeostatic plasticity can be incorporated into CTRNNs by defining a target range for the firing rates of neurons, corresponding to the postulated set level of activity about which homeostasis is maintained, and then triggering plasticity whenever the firing rate of a neuron is too high or too low. This notion is captured by the use of a plastic facilitation function that varies with firing rate, adapted from Di Paolo (2000); plastic facilitation ρ is zero when the firing rate is within the target range and rises or falls linearly to ± 1 outside this range, as shown in figure 3(b). It is not clear what the optimal bounds for the target range would be for any given situation; in the results reported below the upper and lower bounds were arbitrarily set to 0.75 and 0.25 respectively.¹

Synaptic scaling (figure 2(a)) can be implemented in CTRNNs as multiplicative scaling of synaptic weights when the firing rate of a neuron goes outside the prescribed range. The scaling is directional; it acts so that weights are changed in the direction most likely to bring the neuron firing rate back into bounds. Scaling is applied to both inhibitory (negative weight) and excitatory (positive weight) synapses, and refers to the absolute value of the synaptic weight. The size of the change is determined by the plastic facilitation ρ , by a time constant τ_w , and by the current magnitude of the weight. The plasticity rule for synaptic scaling is expressed by equation 3.

Plasticity of the intrinsic excitability of neurons can be implemented in CTRNNs as an adaptive bias term. When a neuron’s firing rate goes outside the prescribed range, the bias term of the neuron is shifted to make the neuron more or less likely to fire depending on what is required to bring the

¹Sensitivity tests on these parameters showed that other values could have been used without changing the qualitative nature of the results achieved.




Figure 3: Plastic facilitation depends on firing rate. Plasticity occurs when firing is outside a designated target range; the size and direction of the excursion from the range determine the rate and direction of plastic change.

firing rate back into bounds. The size of the change depends on the plastic facilitation ρ and a time constant τ_b . The plasticity rule for intrinsic plasticity is given by equation 4.

$$\tau_w \dot{w} = \rho |w| \quad (3)$$

$$\tau_b \dot{b} = \rho \quad (4)$$

3.3 Measuring signal propagation

Signal propagation will be measured as the change in firing rate for each node in a network caused by a change in input. In order to gain a representative measure of signal propagation in a particular type of CTRNN topology we will look at the mean change in firing caused by many different input changes, for large numbers of randomly parameterised networks of that type. In order to get a picture of signal propagation in CTRNNs in general we will look at several different types of topology, though we will not attempt to cover the whole space of all possible network architectures, since such an undertaking lies beyond the scope of this paper. The topologies chosen are shown in figure 4. To measure the impact of homeostatic plasticity, we will look at signal propagation in networks before and after a period of homeostatic plasticity. Data will always be gathered on




Figure 4: Different network topologies used to test signal propagation: (a) Fully connected. Example shown has 6 nodes. Each node has a self-connection and an afferent connection from every other node. (b) Random. Based on Erdos-Renyi scheme. Example shown has 6 nodes, with 50% probability of connection between each pair of nodes. (c) Feedforward. Example shown has 5 layers of 3 nodes. There are no backwards or lateral connections.

non-plastic networks in order to avoid any interference from the plastic mechanisms and allow a fair comparison; the period of plasticity might be loosely compared with a developmental phase in a biological organism.

4 Results

Each type of network topology was tested in a similar way. Ensembles of randomly parameterised networks were generated. Signal propagation was examined in each network in its original state, and then again after a period in which homeostatic plasticity was applied. All connection weights and biases were drawn from the range $[-10.00, 10.00]$ and all neuron potential decay constants were drawn from the range $[1.00, 4.00]$. Plasticity was applied with $\tau_w = 40$ and $\tau_b = 20$. Networks were updated using Euler integration with a step size of 0.2 timesteps.

Network input was randomly chosen from a uniform distribution. Input was held constant for a period during which network firing rates were measured and then a new input was chosen. The mean size of the change in output caused by each change in input was calculated over a large number of input presentations to give a representative measure of the change in network activity that might typically be expected from a change in input. This measure was used to compare signal propagation between different network topologies before and after plasticity was applied.

The change in output was assessed in different (though similar) ways depending on the network topology. The mathematical formulation of CTRNNs means that stable oscillatory dynamics can occur if two or more nodes are connected in a loop. In networks where cyclic paths are possible (that is, fully and randomly connected networks but not feed-forward networks), there is a good chance that oscillatory dynamics will occur. In these networks, the mean firing rate was measured for each node over the full period for which input was held constant. The period was chosen to be long enough so that any transient dynamics while the network settled to a new stable state following the change in input would have an insignificant on the recorded mean. The use of the mean allowed a rough comparison between oscillatory and fixed point stable states; if the oscillation changed then its mean value would most likely change also. For the feedforward networks, where oscillations cannot occur, the firing rate of each node was measured at the end of each period prior to the presentation of a new input vector. This allowed the network to settle to a new fixed point before measurement.

4.1 Fully connected networks

This is the architecture most commonly used in the evolutionary robotics literature. There is a connection in both directions between each pair of nodes and each node has a self-connection (see figure 4(a)). A key point to note is that there is a path length of one link between any pair of nodes, meaning that in a fully connected architecture there is a direct connection between input and output nodes. The direct influence of the input node on the output node is modulated by the activity of the other nodes.

Networks were created with 1, 3, 5 and 10 nodes; 200, 600, 1000 and 2000 networks of each respective size were created to reflect the combinatorial expansion in the number of parameters.² A single node in each network was designated the input node and received input randomly drawn from the range $[-5.00, 5.00]$, held fixed for 200 timesteps. Signal propagation was measured over 1000 input presentations. Homeostatic plasticity was applied for 500 timesteps and then signal propagation was measured again.

²There are $N(N + 2)$ parameters in a fully connected N-node CTRNN; in 1-node, 3-node, 5-node and 10-node networks there are 3, 15, 35 and 120 parameters respectively. The authors realise that the numbers of networks used to generate the data do not therefore reflect an even sampling rate but feel that sufficient data was produced to suffice for the intended demonstration.




Figure 5: Signal propagation in fully connected CTRNNs after homeostatic plasticity has been applied. Mean change in node firing rate in response to a random change in network input is shown for N -node networks for $N \in \{1, 3, 5, 10\}$. Mean changes in node firing rates in response to stimuli are increased by homeostatic plasticity: dark grey represents pre-plasticity level, light grey is post-plasticity increase.

Results are plotted in figure 5, which shows the mean changes in output caused by a change in network input for each node in the network, before and after the homeostatic plasticity is applied. The input node in each network is marked on the plots; the other nodes shown are the hidden nodes. The input node shows a significantly greater response than the other nodes. However, before plasticity is applied, even the input node does not demonstrate a large mean change in state, and the other nodes show negligible change. The mean change of state in any node is inversely related to the size of the network; this is because as the size of the network grows, the external network input becomes less significant compared to the influence of other nodes. After plasticity, the same pattern is repeated, with the input node showing a much greater change in state in response to a change in input. However, the overall level of response is much greater than the pre-plasticity networks.

4.2 Randomly connected networks

Biological neural networks are not fully connected, but are much more sparsely connected. For this reason we studied the effects of homeostatic plasticity on signal propagation in 10-node CTRNNs where connectivity was based on the random graph scheme devised by Erdos-Renyi (Newman, 2003). In these graphs edges between vertices are assigned at random with a fixed probability (see figure 4(b)). Here we created networks by assigning afferent connections between each pair of nodes with fixed probability, generating a random weight value for each connection created. These




Figure 6: Signal propagation in CTRNNs based on Erdos-Renyi random graphs. Networks are created by assigning afferent connections between each pair of nodes with fixed probability. Mean change in node firing rate in response to a random change in network input is shown for 10-node networks for $P(\text{Connection}) \in [0.0, 1.0]$. Mean changes in node firing rates in response to stimuli are increased by homeostatic plasticity: circle and square markers represent pre-plasticity and post-plasticity levels respectively.

networks are not intended to mimic the structure of biological neuronal networks (which in any case varies in different species and in different regions of the brain), but simply to give an idea of signal propagation in more sparsely connected networks.

Ensembles of 1000 networks were generated for $P(\text{Connection}) \in \{0.0, 0.1, \dots, 1.0\}$. A single node in each network was designated the input node and received input drawn from the range $[-5.00, 5.00]$, held fixed for 200 timesteps. Signal propagation was measured over 1000 input presentations. Homeostatic plasticity was applied for 500 timesteps and then signal propagation was measured again.

The effect of changes in input on randomly connected networks is shown in figure 6. Input nodes are most affected for all P values, as would be expected, but as connectivity increases the effect of input decreases due to the increased influence of input from other nodes. Hidden nodes are most affected at intermediate connectivity rates of around 20-30%; below this connectivity rate there are insufficient connections for signals to be able to propagate, above this rate the large number of inputs each nodes receives reduces the effective influence of the input node. Homeostatic plasticity significantly increases the effect of input in all cases.

4.3 Feed-forward networks

An example of the feedforward architecture used is given in figure 4(c). Each node has a self-connection and receives input from every node in the preceding layer; there are no lateral connections within a layer or return connections to the preceding layer. This architecture is included to look at signal propagation through multiple network layers.

For this architecture input was applied to all nodes in the first layer of the network, and signal propagation was measured as the magnitude of the change in the firing rate vector at each layer subsequent to a change in input. Each node in the input layer received input from the range $[-1.00, 1.00]$ which was held constant for 100 timesteps. The mean size of change was measured over 100 input presentations, for ensembles of 200, 1000 and 2000 networks with 1, 3 and 5 nodes per layer respectively.³

Figure 7 shows the mean change in state vector at each layer for the ensembles generated, before and after the application of homeostatic plasticity. It can clearly be seen that prior to plasticity, the change in input typically does not affect many downstream layers of the network for any of the network sizes. For the $N = 1$ networks (effectively chains of individual neurons) only neurons in the first 3 layers are affected, and by small amounts. As N rises the signal travels further, but even for $N = 5$ the signal does not get further than the 10th layer. The signal travels further when N is larger because not only is the change in the input vector more significant (incorporating a change in N component dimensions), but each neuron receives input from more neurons in the previous layer (recall that in the feedforward architecture each neuron receives input from all neurons in the previous layer as well as its own self-stimulation; there are no lateral or backwards connections). This has the effect of amplifying the change in state at each layer and thus allows the signal to travel further. It is worth pointing out that if any of these feedforward controllers were used for robot control they would produce a robot that never did anything. Changes in input never cause a change in output, meaning that an agent controlled by the network would never change its behaviour in response to input.

After the homeostatic plasticity has been applied, the change in input clearly causes a change

³The number of parameters in a feedforward network with L layers of N nodes is $LN(N + 3)$; for networks with 25 layers of 1, 3 and 5 nodes there are 100, 450 and 1000 parameters respectively. Obviously the sampling rate is low, but again the authors hope that sufficient data has been generated to support the argument here.




Figure 7: Signal propagation in feedforward networks with 25 layers of 1, 3 and 5 nodes, shown as mean change in firing rate vector for each layer caused by a random change in network input. Layer 0 in the plots represents the input vector. Plots shown for networks before (dashed line) and after (solid line) homeostatic plasticity has been applied.

in state at layers much further downstream from the input layer than beforehand. This trend is seen in all the network sizes tested, although it is more prominent in the networks with larger N since these are inherently more conducive to signal propagation because of their greater level of connectivity.

5 Discussion

The results presented clearly show that the effects of homeostatic plasticity on CTRNNs are increased sensitivity and improved signal propagation. The reason for this is that the homeostatic plasticity offers a directed mechanism that avoids node saturation and maintains each node in the network in its most responsive region of parameter space. At the network level this offers improved signal propagation. Sensitive nodes make sensitive networks.

While the usefulness of homeostatic plasticity in CTRNNs for robot control is yet to be demonstrated, it is reasonable to expect that the increased sensitivity and improved signal propagation characteristics that this type of plasticity introduces will result in control networks that are a better substrate for the generation of behaviour than their non-plastic counterparts. Input should be effectively translated into output, allowing the agent to respond to environmental stimuli (a basic pre-requisite for behaviour that is not always offered by standard-form CTRNN controllers). Homeostatic plasticity creates *poised* networks that are ready to behave, as compared to the *inert* networks that result from node saturation.

The impact of homeostatic plasticity on the evolvability of CTRNNs will be examined in future work, but we feel optimistic that the improvement it creates in CTRNNs as a substrate for the generation of behaviour will be transferred into their improvement as a substrate for evolution. Mathayomchan and Beer (2002) showed that the evolution of good CTRNNs for a rhythmic motor control task was made significantly quicker when the evolutionary search was seeded with a population of centre-crossing networks. Centre-crossing networks are strongly related to the types of network created by homeostatic plasticity. In centre-crossing networks the bias term for each node is calculated from the weights so that the node's equilibrium output is in the centre of its range, while homeostatic plasticity dynamically maintains nodes close to this position. Although there is the possibility that homeostatic plasticity will prove to be a disruptive force on the evolutionary process, there is still good reason to believe that it may deliver evolvability pay-offs similar to those offered by centre-crossing networks.

A more speculative discussion concerns the relevance of the results presented here for biology. Might homeostatic plasticity cause similar effects in living neural systems to those that it does in artificial ones? Caution must be exercised when attempting to compare the behaviours of such simple model networks to the behaviours of biological brains and it would be unwise to suggest that such models can generate facts about real neural systems. However, the simplicity of CTRNNs compared to real brains means that they are much more amenable to analysis, and this tractability also allows for the generation and testing of hypotheses about the behaviour of neural systems in general (cf. the idea of the 'frictionless brain' (Beer, 2003)). Current opinion in neuroscience is that homeostatic plasticity in biological nervous systems prevents the occurrence of quiescence and hyper-excitation, but this is extremely difficult to test in real brains. Here we have demonstrated that homeostatic plasticity prevents node saturation in an artificial neural network, and while we do not claim to have answered the neuroscience question we feel that the current work at least offers a proof of concept.

Criticisms that might be raised include the lack of noise in the systems tested here; we would maintain that the introduction of noise would not make any significant difference to the results presented, and also that the increased sensitivity caused by homeostatic plasticity may offer a good means for overcoming the noise inherent in any real-world system. Another criticism might be that we make no attempt to differentiate the different types of homeostatic plastic mechanism, but

instead lump them in together. Our reply to this would be to point out that there is no space for separate consideration of each mechanism in the present work and to anecdotally mention that we have run all the experiments described above on each mechanisms individually and found no qualitative difference in the results.

To conclude, we have shown that the saturation problem in artificial neural networks can be solved with homeostatic plastic mechanisms. These artificial plastic mechanisms are inspired by the biological mechanisms which are thought to lead to the avoidance of the analogous problem of hyperexcitation and quiescence in real neuronal networks. The results of preventing node saturation are increased sensitivity and improved signal propagation, the functional benefits of which will be the subject of future work.

References

- Abbott, L., & Nelson, S. (2000). Synaptic plasticity: Taming the beast. *Nature Neuroscience*, *3*, 1178–1183.
- Beer, R. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M., Meyer, J., Pollack, J., & Wilson, S. (Eds.), *From Animals to Animats 4: Proceedings of the 4th Intl. Conf. on Simulation of Adaptive Behavior*, pp. 421–429 Cambridge, MA. MIT Press.
- Beer, R. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, *11*(4), 209–243.
- Beer, R. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, *3*, 469–509.
- Burrone, J., & Murthy, V. (2003). Synaptic gain control and homeostasis. *Current Opinion in Neurobiology*, *13*, 560–567.
- Davis, G., & Bezprozvanny, I. (2001). Maintaining the stability of neural function: A homeostatic hypothesis. *Annual Review of Physiology*, *63*, 847–869.
- Desai, N., Rutherford, L., & Turrigiano, G. (1999). Plasticity in the intrinsic excitability of neocortical pyramidal neurons. *Nature Neuroscience*, *2*, 515–520.
- Di Paolo, E. (2000). Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions. In Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. (Eds.), *From Animals to Animats 6: Proceedings of the 6th Intl. Conf. on Simulation of Adaptive Behavior*, pp. 440–449 Cambridge, MA. MIT Press.
- Funahashi, K., & Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, *6*, 801–806.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., & Jakobi, N. (1997). Evolutionary robotics: the sussex approach. *Robotics and Autonomous Systems*, *20*(2-4), 205–224.
- Mathayomchan, B., & Beer, R. D. (2002). Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation*, *14*, 2043–2051.
- Newman, M. (2003). The structure and function of complex networks. *SIAM Review*, *45*, 167–256.
- Turrigiano, G. (1999). Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends in Neuroscience*, *22*, 221–228.
- Williams, H. (2004). Homeostatic plasticity in recurrent neural networks. In Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., & Meyer, J.-A. (Eds.), *From Animals to Animats 8: Proceedings of the 8th Intl. Conf. on Simulation of Adaptive Behavior* Cambridge, MA. MIT Press.
- Williams, H., Buckley, C., & Noble, J. (2005). Continuous-time recurrent neural networks as a substrate for the evolution of behaviour.. Submitted to *Adaptive Behavior*.
- Zhang, W., & Linden, D. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nature Reviews Neuroscience*, *4*, 885–900.