

# Homeostatic Plasticity in Recurrent Neural Networks

Hywel Williams

Biosystems, Informatics Research Institute, University of Leeds, Leeds LS2 9JT, UK  
hywelw@comp.leeds.ac.uk

## Abstract

Recent work in the neurosciences has highlighted the existence of homeostatic plasticity mechanisms that regulate the activity in neuronal networks to a constant level. Functionally similar mechanisms are applied to continuous-time recurrent neural networks. Homeostatic plasticity is found to increase the sensitivity of neural networks to changes in input, giving richer dynamics and robustness to severe changes in the mean level of input. Simulated robots controlled by continuous-time recurrent neural networks using homeostatic plasticity display more complex behaviours than those controlled by non-plastic networks. Synaptic scaling in randomly parameterised networks is observed to generate phototactic behaviour without further adaptation.

## 1. Introduction

When we are cold, we shiver. When we are too warm, we start to sweat. Both of these mechanisms help us to maintain our body temperature at a constant level. Alternatively, we may move closer to the fire or shed a layer of clothing, or on a longer timescale we may learn new behaviours; jumping up and down to keep warm, a dip in the pool to cool off. Homeostasis is maintained by a combination of physiological response, behaviour and learning.

Recent work in the neurosciences has identified a number of plastic mechanisms by which homeostasis of neural activity is maintained. This paper draws inspiration from these mechanisms to explore the effect of homeostatic plasticity in dynamic recurrent neural networks, such as those that are often used in evolutionary robotics. The evolutionary optimisation of these networks for dynamic control tasks is a significant challenge, and networks using homeostatic plasticity may provide an evolutionary substrate offering improved evolvability and greater robustness of evolved solutions.

Section 2 gives some background to the paper, describing biological neural homeostasis and relevant work in evolutionary robotics. Section 3 covers the experimental method used, the results from which are presented in section 4. Sections 5 and 6 attempt to interpret the results and offer some thoughts on their implications.

## 2. Homeostatic Adaptation

### 2.1 *Ashby's Homeostat*

The connection between homeostasis of internal variables and adaptive behaviour was first recognised by Ashby in the 1950s (Ashby, 1952). Ashby talked about two kinds of homeostatic adaptation: adaptive behaviour that maintains homeostasis of internal variables, and adaptation by which such homeostasis-maintaining behaviour is created. Ashby created an analogue electrical device called the Homeostat to demonstrate the idea of *ultrastability*; plastic change triggered by loss of homeostasis and continuing until homeostasis is regained.

### 2.2 *Neural homeostasis*

It is now accepted in neuroscience that many neural circuits adapt to changing patterns of activity and appear to dynamically maintain their overall level of activity around some set point (Turrigiano, 1999, Turrigiano and Nelson, 2000). 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 and the intrinsic excitability of individual neurons.

#### 2.2.1 *Synaptic scaling*

One way of maintaining a relatively constant level of activity in the face of a prolonged increase or decrease in firing rates is by altering the strengths of excitatory synaptic connections. If the overall level of firing of a neuron falls too low, an increase in the strength of all its excitatory connections should help to raise it; similarly, if the level of firing gets too high a decrease in the strength of excitatory connections should help to reduce it. This form of homeostatic plasticity is known as *synaptic scaling* and specific mechanisms that enable it have been identified (Turrigiano et al., 1998, O'Brien et al., 1998). It is thought that the trigger signal is some function of post-synaptic firing rate (Leslie et al., 2001), which




Figure 1: Homeostatic plasticity: (a) synaptic scaling increases/decreases the effect of input, (b) intrinsic plasticity translates the activation curve.

causes a multiplicative scaling of the strengths of all the neuron’s afferent synaptic connections. Contrast this with the synapse-specific changes associated with Hebbian correlation-based plastic mechanisms such as long-term potentiation or depression. Synaptic scaling affects the function of a neuron by altering its rate of firing for a given amount of stimulus; in effect it moves the neuron along its input/output curve (see figure 1).

### 2.2.2 Intrinsic plasticity

Another way of regulating activity around a constant level is alteration of the intrinsic excitability of neurons. If firing rate is consistently too high, a reduction in the excitability of the neuron should reduce it; if firing rate is consistently too low, increased excitability should raise it. Alteration of excitability assumes some intrinsic plasticity of neurons; such plasticity has been reported (Turrigiano et al., 1994). The effect of this plasticity on intrinsic excitability has also been demonstrated in cultured pyramidal cortical neurons, where externally imposed damping of activity led to lowered spiking thresholds (Desai et al., 1999). Intrinsic plasticity of the excitability of neurons affects the neuron’s function by changing the firing rate response to a given input; effectively this translates the neuron’s input/output curve along the output axis (see figure 1).

### 2.3 Homeostatic adaptation in neural networks for robot control

To the best of the author’s knowledge there is only one previously published study of homeostatic adaptation in artificial neural networks, which is the work presented by Di Paolo (2000). Using ideas from Ashby and Turrigiano, Di Paolo developed a system that incorporated a form of homeostatic adaptation into a continuous-time

recurrent neural network used to control a phototactic robot. In Di Paolo’s scheme, each synapse in the network had a Hebbian learning rule associated with it. The learning rules were only activated when the firing rate of the post-synaptic neuron passed out of a prescribed range, becoming either too high or too low; when post-synaptic firing was inside the range there was no plasticity.

Di Paolo’s work was an exploration of the utility of Ashby’s concept of homeostatic adaptation for providing robustness to sensorimotor disruptions. Networks were evolved using a genetic algorithm to give phototactic behaviour in a simulated robot, with fitness awarded for phototaxis and for maintenance of neural activation within the target range. The successful controllers were then tested for their robustness to inversion of their visual field, performed by swapping the positions of the two light sensors. Di Paolo found that around half of the evolved controllers that gave long-term stable phototaxis were able to adapt to this disruption.

While Di Paolo’s work produced some very interesting results, the types of homeostatic plastic mechanism implemented in his experiment were qualitatively different to those observed in neuroscience. Di Paolo’s use of plasticity triggered when neural firing is too high or too low fits in with the neuroscientific data, but the form of the plasticity does not. Di Paolo used Hebbian learning, which acts on the strengths of individual synapses. In contrast, the synaptic scaling reported by Turrigiano acts upon groups of synapses afferent to a node. Intrinsic plasticity does not operate on synaptic strengths at all but on the excitability of the neuron. In addition, the homeostatic plasticity observed in biological neural networks gives negative feedback, whereas Hebbian rules are seen by neuroscientists as likely to provide positive feedback on both activity levels and synaptic strengths (Turrigiano and Nelson, 2000).

## 3. Method

While this study is clearly influenced by Di Paolo’s work and makes use of a similar experimental set-up for studying robot behaviour, its aims and objectives are significantly different. The intention here is to apply the kinds of homeostatic plasticity observed in neuroscience to continuous-time recurrent neural networks (CTRNNs), of the type commonly described in the evolutionary robotics literature (Beer and Gallagher, 1992, Beer, 1996, Harvey et al., 1997, Slocum et al., 2000, Nolfi and Floreano, 2000), and to examine the effect of homeostatic plasticity on network dynamics and on the behaviour of simulated robots controlled by these networks.

This section describes the networks used, the way in which the homeostatic plasticity mechanisms identified by Turrigiano are implemented, and the robot simula-

tion. Networks using the two types of homeostatic plasticity (synaptic scaling and intrinsic plasticity) will be compared with non-plastic networks throughout.

### 3.1 CTRNN

The CTRNNs used here are made up of artificial neurons defined by the equations 1 and 2 given below. Neuron state is governed by equation 1, in which  $y$  is the neuron potential,  $\tau$  is a constant governing the decay of potential with time,  $w_i$  and  $z_i$  are the connecting synapse strength and firing rate respectively of the  $i$ th afferent neuron, and  $I$  is any external input the neuron receives.

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


The firing rate  $z$  of a neuron is given by equation 2, in which  $b$  is the bias term for the neuron. Firing rate is thus given by a sigmoid function of potential returning values in the range  $[0, 1]$  as shown in figure 2(a).

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


The networks used are all fully connected, including self-connections. Synapses can be inhibitory or excitatory depending on their sign, and may have varying strengths; this is represented by an associated signed real weight value  $w$ . Values are initialised from the following ranges:  $\tau \in [0.4, 4.00]$ ,  $b \in [-3.00, 3.00]$ ,  $w \in [-8.00, 8.00]$ . Networks are updated using Euler integration with a time step of 0.2. Certain nodes are designated as input or output (sensor or motor) nodes as appropriate and are the only nodes that directly interact with the external environment; all other nodes are hidden nodes and only interact with the external environment via input/output nodes.

### 3.2 Homeostatic plasticity mechanisms

Homeostatic plasticity is incorporated into the CTRNN by defining a permitted range for the firing rates of neurons, corresponding to Turrigiano’s postulated set level of activity about which homeostasis is maintained. Plasticity is triggered whenever the firing rate of a neuron is too high or too low and is proportional to the amount by which the firing rate has exceeded its range. This notion is captured by the use of a plastic facilitation function that varies with firing rate (Di Paolo, 2000); plastic facilitation  $\rho$  is zero when the firing rate is within the target range and rises or falls to  $\pm 1$  outside this range, as shown in figure 2(b). The target range is arbitrarily set to  $[0.3, 0.7]$ .



(a) Sigmoid activation curve showing upper and lower bounds of target range



(b) Plastic facilitation as a function of firing rate

Figure 2: Plastic facilitation depends on firing rate

#### 3.2.1 Synaptic scaling

When the firing rate of a neuron goes outside the prescribed range all afferent synapses to that neuron are multiplicatively scaled. 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. If the firing rate is too high, then excitatory synapse strengths are scaled down and inhibitory synapse strengths are scaled up. If the firing rate is too low, excitatory inputs are scaled up and inhibitory inputs are scaled down. Scaling up or down here refers to the absolute value of the synaptic weight, so that scaling down a negative weight makes it less negative. The size of the change is determined by the plastic facilitation  $\rho$ , by a node learning rate parameter  $\mu$ , and by the current magnitude of the weight. The plasticity rule for synaptic scaling is therefore expressed by equation 3, with  $\mu \in [0.00, 0.10]$ . Weight values are restricted to the range  $[-8.00, 8.00]$  and are clipped where they go outside this range.

$$\Delta w = \rho \mu |w| \quad (3)$$

### 3.2.2 Adaptive bias

This mechanism is an implementation of the intrinsic plasticity described by Turrigiano (Turrigiano, 1999, Turrigiano and Nelson, 2000). 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 firing rate back into bounds. If firing rate is too low the bias is increased, effectively translating the sigmoid activation function so that the neuron is more excitable and hence more likely to fire. If firing rate is too high the bias is decreased, making the neuron less excitable. The size of the change depends on the plastic facilitation  $\rho$  and the node learning rate  $\mu$ . The plasticity rule for intrinsic plasticity is therefore given by equation 4, with  $\mu \in [0.0, 0.10]$

$$\Delta b = \rho\mu \quad (4)$$

### 3.3 Robot simulator

The behavioural effect of homeostatic plasticity at an agent level is studied by incorporating the plastic mechanisms into CTRNNs used to control a simple simulated robot. The robot is placed in a featureless infinite plane along with a single light source. The robot's light sensors and control network are then sequentially updated, with network output being used to calculate the motion of the robot relative to the light source.

#### 3.3.1 Simulated robot

The robot is modelled as a circular body with two motors mounted at either ends of an axle along its diameter and two light sensors mounted at angles of  $\pi/3$  radians from the forward direction, as shown in diagram 3.




Figure 3: Diagram of photosensitive robot

#### 3.3.2 Robot kinematics

The simulated robot kinematics are adapted from Dudek and Jenkin (2000). Differential drive steering is accomplished by the two motors, which may give thrust both forwards and backwards, allowing the robot to spin while stationary or to move forwards or backwards with any instantaneous angular velocity. Robots are assumed to have negligible mass, so that the motor output can be taken as the tangential velocity of the robot at the motor mount point. The instantaneous angular velocity  $\omega$  of the robot is calculated using equation 5, and this is used to calculate the motion of the robot using the system of equations 6, where  $V_L$  and  $V_R$  are the instantaneous velocities of the left and right motors respectively,  $d$  is the length of the axle,  $ICC$  denotes the instantaneous centre of curvature (the imaginary point about which the robot orbits at any given instant),  $\beta$  is the current heading of the robot, and  $(x, y)$  is the robot's position. Details of how this system is derived can be found in Dudek and Jenkin (2000). The equations are updated using time step  $\delta t = 0.2$ .

$$\omega = (V_R - V_L)/d \quad (5)$$

$$\begin{aligned} x' &= (x - ICC_x)\cos(\omega\delta t) - (y - ICC_y)\sin(\omega\delta t) + ICC_x \\ y' &= (x - ICC_x)\sin(\omega\delta t) + (y - ICC_y)\cos(\omega\delta t) + ICC_y \\ \beta' &= \beta + \omega\delta t \end{aligned} \quad (6)$$

#### 3.3.3 Light sensors

Light sources are of fixed intensity 1000 unless otherwise stated. The light sensors return a signal value proportional to source intensity and the inverse square of the distance between the source and the sensor. Sensors are assumed to be mounted on top of the agent, so that they are never in the shadow of the robot's body.

#### 3.3.4 Control network

The robot is controlled by a fully connected 6-node CTRNN of the type described above. Two nodes are designated sensor nodes and each receive the signal from one of the light sensors modified by an agent-symmetrical gain parameter (drawn from the range  $[0, 10.00]$ ) as external input. Two other nodes are designated motor nodes, and the output from these nodes is mapped to the range  $[-1.00, 1.00]$  and modified by an agent-symmetrical gain (again drawn from the range  $[0, 10.00]$ ) to give the instantaneous velocity created by that motor. Both sensory input and motor output have noise added before the gain is applied, drawn from a uniform distribution in the range  $[-0.25, 0.25]$ .

## 4. Results

### 4.1 Network dynamics

The first experiments were performed on stand-alone CTRNNs, that is, without embedding the networks in the robot simulation environment. These tests were to get a feel of how network dynamics varied with the different forms of homeostatic plasticity. Random sets of network parameters (weights, bias terms, decay constants, learning rates) were generated and networks formed from them. Network behaviour was observed with random input fluctuating around a mean level of zero and with random input fluctuating around a steadily increasing base level. The results reported are for fully-connected 3-node networks with a single input node and two hidden nodes.

#### 4.1.1 Random input with zero mean

Randomly generated networks were updated for 400 timesteps (2000 Euler integration steps). The input signal to the input node was reset every 20 time steps to a value randomly drawn from a uniform distribution in the range  $[-1.00, 1.00]$  and remained constant between each change. Samples of network behaviour are given in figures 4, 5 and 6, which show the dynamics of three typical examples of each network type. As can be seen from these plots, the behaviour of the plastic and non-plastic networks was significantly different. The plots show the firing rates of the the three nodes over time, as well as plastic changes where appropriate. The dashed horizontal lines on the plots mark the upper and lower boundaries of the target range; above and below these lines plasticity is triggered.

Non-plastic CTRNNs typically settle to a steady equilibrium with each neuron giving a constant output. The input node is sometimes affected by changes in input signal, but the other nodes typically do not change their activity in response to a change in input to the network; even when the firing rate of the input node changes in response to a change in input signal these changes are rarely transmitted to the other nodes. This is because in a randomly generated fixed-weight CTRNN there is a strong chance that the nodes will reach an equilibrium with firing rate close to the maximum or minimum level, which because of the sigmoid shape of the activation function means that changes in input have little effect on output.

The dynamics of networks with both forms of homeostatic plasticity are more complex than those of non-plastic networks. The homeostatic plasticity functions to keep the firing rates of all the nodes within the target range, as can be seen by inspection of the changes in synaptic weights or biases which occur at times when the firing rate of a node goes outside the target range. This




Figure 4: Behaviour of three typical non-plastic 3-node CTRNNs. Input signal changed every 20 timesteps, randomly drawn from the range  $[-1.00, 1.00]$ . There is a tendency towards saturated/constant firing rates, and changes in the firing rate of the input node do not affect the firing rates of the hidden nodes.

restriction means that the inputs to each node are fluctuating around the region of the sigmoid activation curve where the gradient is greatest; at this point the nodes are most sensitive to changes in input. All nodes are sensitive to changes in each other's firing rate, and the input node is additionally sensitive to changes in the input signal. Thus a change in the input signal is almost always transmitted via the input node to the hidden nodes.

The increased sensitivity of the homeostatic networks is shown in figure 7, which shows the mean change in the firing rate of each node in the network caused by a change in the input signal. This statistic is calculated by measuring the change in the firing of each node following each change in the input signal. Readings are taken just before the next change in input signal so that the network has time to settle on a steady activation level following the previous change. The mean change in firing rate was calculated for 1000 randomly generated networks of each type to give an overall mean value plotted in figure 7. The plots show that randomly generated non-plastic CTRNNs typically have a very low sensitivity to changes in external input, whereas the CTRNNs with homeostatic plasticity have much greater sensitivity.

CTRNNs with synaptic scaling showed higher sensitiv-

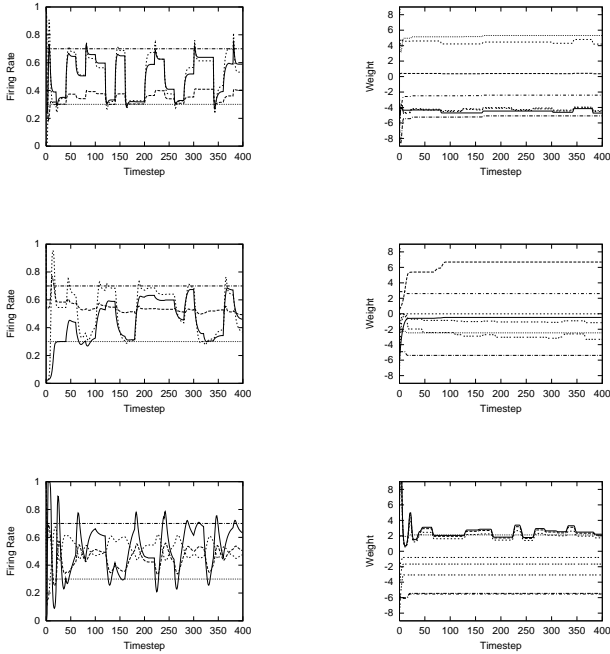


Figure 5: Behaviour of three typical 3-node CTRNNs with synaptic scaling. Input signal changed every 20 timesteps, randomly drawn from the range  $[-1.00, 1.00]$ . Plots in the left-hand column show node firing rates, plots in the right-hand column show the connection weights for the network. Horizontal dotted lines in the plots of firing rates show the upper and lower bounds of the target range. Firing rates are kept within the target range by the synaptic scaling; changes in connection weights are correlated with periods when firing rates leave the target range. Dynamics are richer than those of non-plastic CTRNNs.

ity to input fluctuations than the non-plastic CTRNNs on all nodes. The input node is more sensitive than the hidden nodes, which is explained by its direct exposure to the changing input signal; the hidden nodes are only affected by the changing input indirectly through their interaction with the input node. CTRNNs with adaptive bias are the most sensitive of all the network types, with all the nodes showing relatively large changes in firing rate in response to fluctuating network input. This is because the adaptive bias translates the sigmoid activation curve for each node so that its point of steepest gradient is positioned roughly at the mean level of input to the node, giving the largest possible change in firing rate in response to a given change in input. The increased sensitivity of the CTRNNs with adaptive bias compared to those with synaptic scaling perhaps suggests that adaptive bias is more effective in adjusting the nodal response than synaptic scaling.




Figure 6: Behaviour of three typical 3-node CTRNNs with adaptive bias. Input signal changed every 20 timesteps, randomly drawn from the range  $[-1.00, 1.00]$ . Plots in the left-hand column show node firing rates, plots in the right-hand column show the bias terms for each node. Horizontal dotted lines in the plots of firing rates show the upper and lower bounds of the target range. Firing rates are kept within the target range by the adaptive bias; changes in node bias terms are correlated with periods when firing rates leave the target area. Dynamics are richer than those of non-plastic CTRNNs.




Figure 7: Mean change in node firing rates caused by a change in external input signal. Input signal randomly drawn from range  $[-1.00, 1.00]$ , applied to node 1 (left-most column in each graph). Homeostatic plasticity increased network sensitivity. Sensitivity of networks using synaptic scaling was greater than that of non-plastic networks, while networks using adaptive bias were most sensitive of all.

#### 4.1.2 Random input with increasing mean

Random networks were generated as before and tested for their sensitivity to fluctuations in input around an


incrementally increasing base input level. Every run started with input base level  $base = -20$ . During the run the base level was incremented through the series  $\{-20, -19.5, -19, \dots, 19, 19.5, 20\}$ . After each base level increment, network sensitivity was tested as before, but in this case with input drawn from a uniform distribution with range  $[base - 1.00, base + 1.00]$ . Base level was incremented every 100 time steps and input signal re-assigned every 10 timesteps. Figure 8 shows mean change in node firing rates against input base level, measured across 1000 randomly generated networks of each type.

The non-plastic CTRNNs were relatively insensitive to changes in input signal; they were only sensitive to changes in input signal for input base levels close to zero, and for more extreme levels of input they were completely insensitive. The explanation is that at very high or very low levels of external input the input node of the non-plastic CTRNNs gives a saturated response and the small fluctuations in the input signal have no effect on its firing rate, so the change in input signal is not communicated to the hidden nodes. Even when the input node is not saturated, these networks are relatively insensitive for the reasons outlined above.


CTRNNs with synaptic scaling were more sensitive than the non-plastic CTRNNs for the whole range of input base levels, but were still significantly less sensitive at extreme levels of network input. At very high or low levels of input the input node is actually less sensitive than the other nodes to input fluctuations; this initially seems counter-intuitive, but can be explained by the saturation of its response. The hidden nodes are not directly affected by the external input and so do not saturate, retaining their sensitivity. Synaptic scaling acts to make them more sensitive, so that even a small change in the signal they receive from the input node can cause a significant change in their activation.

The filtering of the input signal through the input node helps the hidden nodes to remain sensitive when the base level is extreme, but also stops them from becoming as sensitive as the input node when the input signal is moderate. The skew of the peak in figure 8(b) to the left of the origin is hard to account for (it might be expected to be centred on the origin). This may be a feature of the way the synaptic scaling was implemented. Possible explanations involving the direction of the incremental change in input base level (from low to high) can be ruled out as the same skew appears when input base level is changed in the other direction.


CTRNNs with adaptive bias remained highly sensitive across the whole range of input base level. This is because the bias term can directly cancel out the external input, as can be seen from equations 1 and 2; the external input translates the activation curve in one direction and the adaptive bias term translates it back



(a) Non-plastic CTRNN



(b) CTRNN with synaptic scaling



(c) CTRNN with adaptive bias

Figure 8: Mean change in node firing rates caused by a change in external input signal. Input signal randomly drawn from range  $[base - 1.00, base + 1.00]$  and applied to node 1, for  $base$  incremented from -20 to 20 (step size 0.5). Compared to non-plastic networks, networks with synaptic scaling showed greater sensitivity across the whole range of base input level. Networks with adaptive bias showed the greatest sensitivity, with a high level of sensitivity consistent across the whole range (the early lower level is due to the time needed for adaptation).

again. Thus the input node does not saturate and the network retains its sensitivity. The dip in sensitivity at the leading edge of the plots in figure 8(c) is accounted for by the time taken for the plastic mechanism to adapt to the initial low level of input.

## 4.2 Robot behaviours

The next set of experiments were performed with the simulated robot described previously. Parameter sets for 6-node CTRNNs and sensor/motor gains were randomly generated and instantiated as robot controllers. It is hard to quantify robot behaviour; instead this section will try and give a descriptive account of the typical behaviours of robots controlled by each kind of CTRNN. Figures 9, 10, 11 and 12 show examples of the robot behaviours described; these plots show motion for 4000 time steps (20000 Euler integration steps) where the light source was randomly repositioned every 800 time steps. The symbol  $\times$  marks the position of light sources and the symbol  $+$  marks the start position of the robot. Note the different scales of the plots.

Robots controlled by randomly generated non-plastic CTRNNs almost always rotated on the spot and ignored the light source. Sometimes they moved in circles, but their motion was consistently a steady rotation. This is explained by the observations of the behaviours of fixed random CTRNNs; they are not generally sensitive to external input and tend to reach an equilibrium state with each neuron firing at a constant rate. If the motor neurons fire at a constant rate the differential drive kinematics will produce a constant angular velocity for the robot, which gives the types of rotational movement observed. An example of this is shown in figure 9. The rotation cannot be seen in the plot as the robot was spinning about its axis almost without any movement relative to the light source. The net distance travelled by the robot is very small, and is due to gradual drift as the robots rotation is very slightly influenced by the different light intensities its sensors experience as it rotates.




Figure 9: Behaviour of robot controlled by random non-plastic CTRNN.

In contrast to the non-plastic CTRNNs, robots controlled by the CTRNNs with synaptic scaling displayed quite complex behaviours. Since the synaptic scaling ensured that they were sensitive to external input (light intensity being at a low enough level not to force input node saturation except when very close to the source), these agents were strongly influenced by the light source. Their motion, as seen in figure 10, is complex; the easiest

description is to liken it to that of a moth near a flame. The robot approaches the source, then veers away, then approaches, then spends some time moving erratically away from the light source. Periods of phototactic behaviour occur, but then cease, as do periods of photo-aversion. The difference in behaviour for different randomly generated networks is considerable, but the plot shown is representative of the general class of behaviours observed.




Figure 10: Behaviour of robot controlled by random CTRNN with synaptic scaling.

Interestingly, while many of these agents showed occasional periods of phototaxis, some (roughly 10% by visual inspection) of the agents controlled by CTRNNs with synaptic scaling displayed stable phototaxis. In these cases the typical pattern was an early period of erratic behaviour, while the synaptic scaling configured the network to be sensitive to external input, followed by consistent phototaxis. An example of this is shown in figure 11, where the robot approaches and then circles the light sources.




Figure 11: Phototaxis by robot controlled by random CTRNN with synaptic scaling.

Robots controlled by CTRNNs with adaptive bias also displayed interesting behaviours, which were generally some form of cycloidal motion (see figure 12). The cycloidal motion is caused by the interaction of the adaptive bias with the random set of fixed weights. Without the adaptive bias term the network is non-plastic and would tend to produce a rotational movement as

described above, but with the adaptive bias the rotation is modulated by the changing bias term to become cycloidal. In most cases the bias term actually becomes a part of the network dynamics, oscillating in phase with the oscillations of neural firing rates. The cycloidal motion means that the robot moves relative to the light source, in contrast to the near-stationary rotation of the non-plastic CTRNN controllers. The cycloidal motion displayed was generally the same irrespective of how far the agent was from the light source (figure 12), which can be explained by reference to the dynamics of the CTRNNs with adaptive bias described earlier. In a similar manner to how it was able to negate the large changes in external input base level in the previous section, the adaptive bias negates the change in the base level of light intensity so that the motion caused by the fixed set of weights remains unchanged.




Figure 12: Behaviour of robot controlled by random CTRNN with adaptive bias.

## 5. Discussion

We have seen that homeostatic plasticity in CTRNNs makes them much more sensitive to changes in their input signal. Synaptic scaling significantly increases the sensitivity of all nodes. Plasticity of intrinsic excitability (implemented here as an adaptive bias term) makes a still bigger increase in sensitivity while also allowing the network to cope with extreme changes in the level of external input received. When these networks are randomly parameterised and used as controllers for a photosensitive robot, they produce more interesting behaviours than randomly parameterised non-plastic CTRNNs. Random non-plastic networks make robots rotate, random networks with adaptive bias cause robots to display cycloidal motion, and random networks with synaptic scaling give complex behaviours including consistent phototaxis.

It is tempting to think that the networks with synaptic scaling would be a good substrate for the evolution of robot controllers using genetic algorithms,

since randomly generated networks of this kind produce interesting behaviours before any evolutionary optimisation has taken place. The networks with adaptive bias might also be seen as a good evolutionary substrate because of their sensitivity, especially given the improved evolvability of centre-crossing recurrent neural networks for the generation of rhythmic behaviours (Mathayomchan and Beer, 2002). The adaptive bias serves to make the CTRNN a dynamically centre-crossing network. However, a cautionary note is required.

In some preliminary work that there is not space to report on fully here, CTRNNs were successfully evolved for phototaxis in the same simulated robot scenario used above, achieving good results in few ( $< 10$ ) generations with a population size of 50. It was found that the networks with homeostatic plasticity were actually less evolvable than the networks without plasticity, giving consistently lower fitness scores. There are several possible reasons for this, many of which are concerned with the design of fitness functions that allow a fair comparison of evolvability, which is not straightforward. Plastic networks have more parameters than non-plastic networks, and require longer trials to allow the plasticity to configure the network, amongst other problems.

A more general reason why homeostatic plastic networks may not necessarily be more evolvable is to do with the utility of sensitivity. Intuitively, it seems that more sensitive networks should be more useful; networks must change their behaviour in response to changing stimuli if they are to perform any non-trivial function. However, the sensitivity given by homeostatic plasticity as implemented here comes at a cost; it precludes the networks from making use of the beneficial effects of node saturation. An example of such a beneficial effect is given by the non-plastic CTRNNs evolved to perform phototaxis, which all made use of a simple strategy: have one motor node constantly saturated (thus giving a fixed contribution to the angular velocity of the agent) and the other motor node sensitive to changes in the activation of one of the sensor nodes. Put another way, use one motor node to give forward motion and the other to steer. This strategy gives the robot a cycloidal motion that slowly spirals towards the light source, the angular velocity increasing as it draws near until it rotates on the spot when close to the source.<sup>1</sup> This cheap and robust phototaxis strategy is not available to the networks with homeostatic plasticity, since their plasticity acts to stop any of their nodes becoming saturated.

Another point of interest is that the evolved controllers incorporating homeostatic plasticity were more robust

<sup>1</sup>As an aside, it was found that agents with only one sensor could reliably perform phototaxis by this method. The agents that had two sensors only made use of one of them, meaning that sensor inversion was a trivial disruption (since either sensor could supply the required stimulus).

than the non-plastic controllers. The networks with adaptive bias were able to cope with dramatic changes in the intensity of the light sources, whereas the other networks were much less able to do so. The adaptive bias networks continued to perform phototaxis when light source intensity was scaled up or down by as much as 2 orders of magnitude, whereas the synaptic scaling and non-plastic networks were able to cope only with changes up to around  $\pm 50\%$  in most cases. This fits in with the results presented above showing the ability of the adaptive bias networks to cope with large changes in the level of external input. Another example of the robustness of the homeostatic plastic networks is that the synaptic scaling networks were often able to recover phototaxis after sensor inversion, despite having phototaxis strategies that relied on the use of both sensors. This is reminiscent of the results achieved by Di Paolo (2000).

One significant area that demands further exploration is that of the timescales of the plasticity mechanisms. In the experiments reported upon here the plastic mechanisms operate at timescales that are much slower than the timescales of neural activation, but they are still very fast compared to the biological mechanisms that are their inspiration. The implications of different timescales for network dynamics will be examined in future work.

## 6. Conclusion

The results described in this paper are modest, but are enough to suggest that homeostatic plasticity is a worthwhile avenue of exploration. Both synaptic scaling and adaptive bias terms make CTRNNs more sensitive to changes in input, and adaptive bias also gives robustness to large changes in the mean level of input. Simulated robotic agents controlled by CTRNNs with homeostatic plasticity are more sensitive to their environment and display more complex behaviours than those controlled by non-plastic CTRNNs. Randomly generated CTRNN controllers using synaptic scaling to maintain homeostasis of neural firing were repeatedly observed to perform stable phototaxis after an initial period of adaptation. Also, evolved solutions based on homeostatic plastic CTRNNs display greater robustness to some sorts of perturbation.

## Acknowledgements

This paper has benefited greatly from discussions with Jason Noble, Chris Buckley, and the rest of the Biosystems group at the University of Leeds.

## References

- Ashby, W. (1952). *Design for a Brain*. Chapman and Hall, London.
- Balaam, A. (2001). *Analysing homeostatic adaptation to sensorimotor disruptions*. Unpublished MSc thesis, University of Sussex, UK.

- Beer, R. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M., Meyer, J., Pollack, J., and Wilson, S., (Eds.), *From Animals to Animats 4: Proceedings of the 4th Intl. Conf. on Simulation of Adaptive Behavior*, pages 421–429, Cambridge, MA. MIT Press.
- Beer, R. and Gallagher, J. (1992). Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1:91–122.
- Desai, N., Rutherford, L., and 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., and Wilson, S., (Eds.), *From Animals to Animats 6: Proceedings of the 6th Intl. Conf. on Simulation of Adaptive Behavior*, pages 440–449, Cambridge, MA. MIT Press.
- Di Paolo, E. (2002). Evolving robust robots using homeostatic oscillators. Technical Report CSRP 548, School of Cognitive and Computing Sciences, University of Sussex, UK.
- Di Paolo, E. (2003). Organismically-inspired robotics: Homeostatic adaptation and natural teleology beyond the closed sensorimotor loop. In Murase, K. and Asakura, T., (Eds.), *Dynamical Systems Approach to Embodiment and Sociality*, pages 19–42, Adelaide. Advanced Knowledge International.
- Dudek, G. and Jenkin, M. (2000). *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge.
- Floreano, D. and Mondada, F. (1996). Evolution of plastic neurocontrollers for situated agents. In Maes, P., Mataric, M., Meyer, J., Pollack, J., and Wilson, S., (Eds.), *From Animals to Animats 4: Proceedings of the 4th Intl. Conf. on Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the sussex approach. *Robotics and Autonomous Systems*, 20(2-4):205–224.
- Leslie, K., Nelson, S., and Turrigiano, G. (2001). Postsynaptic depolarization scales quantal amplitude in neocortical pyramidal neurons. *J. Neurosci.*, 21:RC170.
- Mathayomchan, B. and Beer, R. D. (2002). Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation*, 14:2043–2051.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA.
- O'Brien, R., Kamboj, S., Ehlers, M., Rosen, K., Fischbach, G., and Hagan, R. (1998). Activity-dependent modulation of synaptic ampa receptor accumulation. *Neuron*, 21(5):933–935.
- Slocum, A., Downey, D., and Beer, R. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S., (Eds.), *From Animals to Animats 6: Proceedings of the 6th Intl. Conf. on Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.
- Turrigiano, G. (1999). Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends in Neuroscience*, 22:221–228.
- Turrigiano, G., Abbott, L., and Marder, E. (1994). Activity-dependent changes in the intrinsic properties of cultured neurons. *Science*, 264:974–977.
- Turrigiano, G., Leslie, K., Desai, N., Rutherford, L., and Nelson, S. (1998). Activity-dependent scaling of quantal amplitude in neocortical pyramidal neurons. *Nature*, 391:892–895.
- Turrigiano, G. and Nelson, S. (2000). Hebb and homeostasis in neuronal plasticity. *Current Opinion in Neurobiology*, 10:358–364.