

Homeostatic plasticity improves continuous-time recurrent neural networks as a behavioural substrate

Hywel Williams

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

hywelw@comp.leeds.ac.uk

Abstract—Homeostatic plasticity is applied to continuous-time recurrent neural networks. It is observed to make networks more sensitive, improve signal propagation and increase the likelihood of autonomous oscillations. Evolutionary experiments with a simulated robot show that in some circumstances homeostatic plasticity improves evolvability of good control networks, but in others it makes good controllers less easy to evolve.

I. INTRODUCTION

Animals are by far the most sophisticated and robust autonomous agents known. This has led to a significant amount of research into how biological principles can be applied to improve the design of artificial autonomous agents such as robots. In particular, the use of neural control systems continues to be the focus of widespread effort. One approach in this area is evolutionary robotics, which typically uses artificial evolution to generate neural robot controllers. While this approach has met with some limited success (which together with the tempting possibility of scalable automated design has encouraged continued research effort) it is not yet a consistent and reliable methodology. This paper seeks to improve on current practice by looking at how a commonly used neural architecture (continuous-time recurrent neural networks) might be turned into a better substrate for the evolution of robot behaviours by the addition of homeostatic plasticity. Homeostatic plasticity is a neural mechanism that acts on network parameters to keep neuron activation levels within pre-specified bounds. It stops node saturation and promotes useful dynamics, which is shown to aid network evolvability in some circumstances.

II. BACKGROUND

A. Continuous-time recurrent neural networks

Continuous-time recurrent neural networks (CTRNNs) [1] are widely used within the evolutionary robotics and agent-based modelling communities because they are perceived to be a good substrate for the evolution of behaviour. They offer smooth nonlinear dynamics and have been shown to be able to approximate the output of any dynamical system when correctly parameterised [2]. Their parameters can be easily mapped to a genotype suitable for use with evolutionary algorithms and many examples exist in the literature of the successful evolution of CTRNN-based robot controllers for a variety of tasks. However, successful evolution of good controllers is not necessarily easily achieved and while attempts have been made to improve the method, the general

problem of how to reliably evolve good CTRNNs has not been solved. Even where the evolvability of neural controllers has been improved, the difficulty inherent in the post-hoc analysis of such complicated systems means that the reason for the improvement may not be fully understood.

CTRNNs can be described by a set of equations for each neuron describing the change of potential over time (1) and the relation between potential and firing rate (2):

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

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

where by analogy with biological neurons y represents neuron potential, w_i is the strength of the afferent synapse from the i^{th} neuron in the network, z_i is the firing rate of the i^{th} neuron, I is any external (e.g., sensory) input the neuron receives, b is the bias and τ_y is a neuron-specific decay constant. 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.

B. Node saturation in CTRNNs

The sigmoidal shape of the transfer function (2) means that there is a tendency for node firing rate to saturate when potential is very high or very low. Figure 1 shows three examples of different ranges in which the potential of a neuron might vary as a result of network activity. Neurons with potential fluctuating in range B are not saturated, since range B lies on the steepest part of the sigmoid; a change in input will cause a change in potential and a corresponding change in firing rate. However, neurons with potential fluctuating in ranges A and C are saturated, since ranges A and C lie on the flat tails of the sigmoid; a change in input may cause a change in potential but there will be a negligible change in firing rate. Saturated neurons play no part in network dynamics since they give constant output and ignore changes in input. They do not oscillate and act as barriers to the propagation of signals.

C. Homeostatic plasticity in biological nervous systems

When considering node saturation in artificial neural networks, it is useful to look at biological neural networks and to observe that saturation effects (i.e., hyper-excitation and





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

quiescence) typically do not occur. One postulated reason for this is the existence of homeostatic plastic mechanisms that serve to regulate neural activity [3], [4]. 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 regulation of neural activity exists. It is also clear that this homeostatic regulation is accomplished by a variety of plastic mechanisms, amongst which are mechanisms affecting the strength of synaptic connections [3], [5], [4], [6] and mechanisms affecting the intrinsic excitability of individual neurons [7], [8]. Figure 2 illustrates the effect of these two different types of mechanism on the transfer function of a hypothetical neuron. Homeostatic plasticity prevents persistent hyper-excitation or quiescence and is thought to counter-balance the destabilising effects of Hebbian plasticity [9] and other developmental processes [10]. Analogous mechanisms have been adapted for use in CTRNNs and shown to increase network sensitivity [11], improve signal propagation [12] and give rise to more interesting behaviours when used to control autonomous agents [11].


D. Behavioural substrates

Adaptive behaviour in autonomous agents is always founded in a behavioural substrate, which is the ‘raw material’ from which useful behaviour can be formed. Here we are concerned with CTRNNs as a behavioural substrate. We want to know how capable and how likely they are to produce desirable behaviours.

Reacting appropriately to the environment is essential and requires that the agent be possessed of sensors and actuators with an effective link between them (i.e., signals must be transduced from input to output). This notion is in keeping with the concepts of the sensorimotor loop and embeddedness [13], [14]. Internal dynamics are also likely to be important. Without internal dynamics the agent is restricted to actions that are purely reactive; it can only respond to instantaneous stimuli and cannot integrate stimuli over time. In particular, autonomous oscillations are a kind of internal dynamic that is widely thought to be important for the generation of behaviour. A huge number of biological systems depend on oscillatory dynamics in some form [15], from legged locomotion to digestion, and it is believed that many of these phenomena depend on rhythmic patterns generated in the nervous system. Central pattern generation has thus been the focus of much research effort in the adaptive behaviour community (e.g, [16],



(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.

Fig. 2. Different kinds of homeostatic neural plasticity.


[17], [18]).

As a behavioural substrate, CTRNNs are capable of providing all the properties outlined above. They hold internal state, can propagate a signal from sensors to effectors, and can maintain autonomous oscillatory dynamics. However, while CTRNNs can display these properties, they are not necessarily likely to do so. It seems likely that if CTRNNs can be improved as a behavioural substrate by making such properties more prevalent, they may become more evolvable and thus more useful for robotics.

III. METHOD

A. Aim

We want to know if homeostatic plasticity improves CTRNNs as a behavioural substrate. For comparison, we will look at CTRNNs before and after a period of homeostatic plasticity and consider some basic network performance metrics that measure signal propagation and the likelihood of autonomous oscillations. Because we are interested in the typical effects of homeostatic plasticity on CTRNNs as a general class of objects, we will look at mean effects calculated over ensembles of randomly parameterised networks. To see whether homeostatic plasticity makes CTRNNs more evolvable, we will then look at the performance of networks with and without homeostatic plasticity in evolutionary experiments. In these experiments, controllers for a simple simulated agent are evolved to perform ball catching and shape discrimination tasks.



(a) Sigmoidal transfer function showing upper (H_U) and lower (H_L) bounds of target range for firing rate z .

(b) Plastic facilitation ρ as a function of firing rate z .

Fig. 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.

B. CTRNN

The CTRNNs used here are as described above. Parameters are drawn from ranges such that $\{w, b\} \in [-10.00, 10.00]$ and $\tau_y \in [1.00, 4.00]$. CTRNN dynamics are calculated by integrating the neuron state equation (1) over time using Euler's forward method with a step size of 0.2 timesteps.

C. Homeostatic plasticity

Homeostatic plasticity can be added to CTRNNs by defining a target range for the firing rates of neurons (corresponding to the set level of activity that is homeostatically maintained) and triggering plasticity whenever the firing rate of a neuron goes outside this range. This notion is captured by the use of a plastic facilitation function that varies with firing rate (adapted from [19]). This function is given by equation (3) and plotted in Figure 3(b), where ρ is the level of plastic facilitation and H_L and H_U are the lower and upper bounds of the target range. For these experiments $H_U = 0.75$ and $H_L = 0.25$.¹

$$\rho = \begin{cases} \frac{H_L - z}{H_L} & : & 0 \leq z < H_L \\ 0 & : & H_L \leq z \leq H_U \\ \frac{H_U - z}{1 - H_U} & : & H_U < z \leq 1 \end{cases} \quad (3)$$

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

¹It is not clear what the optimal bounds for the target range would be for any given situation, but sensitivity tests on these parameters showed that other values could have been used without changing the qualitative nature of the results achieved.

constant τ_w , and by the current magnitude of the weight. The plasticity rule for synaptic scaling is expressed by equation (4).

$$\tau_w \dot{w} = \rho |w| \quad (4)$$

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 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 (5).

$$\tau_b \dot{b} = \rho \quad (5)$$

D. Network performance metrics

As a rough measure of how well signals propagate through a network, we will measure the change in activity of all nodes in the network that is caused by a change in the network input signal. To detect autonomous oscillations simply and cheaply, we will look for the combination in each node of non-zero variance in firing rates (indicating some perceptible dynamic activity) and a regular periodic change in sign of the first differential of node potential (\dot{y}), since the co-occurrence of these two traits necessarily implies an oscillation in the firing rate of the node.

E. Evolvability experiments

The first steps in an ongoing exploration of the evolvability of homeostatic plastic CTRNNs will be taken here by comparing the evolutionary performance of non-plastic CTRNNs with that of networks incorporating homeostatic plasticity. Two simple agent-based tasks are used, derived from [20]. They involve a CTRNN-controlled simulated agent orienting itself using visual information, so as to catch falling objects or discriminate between different object types. The experimental set-up is briefly described below; space precludes a more complete description and interested readers are referred to [20] or to the current author for more detail.

The agent (see Figure 4) has 3 ray sensors arranged in a forward-facing fan that return a signal proportional to the proximity of any object with which they intersect, and 2 motors (acting in opposite directions) that give a horizontal velocity proportional to the difference in their activity. The agent is controlled by a fully connected 5-node CTRNN, with each sensor being mapped to a unique node and motor output being read from the remaining nodes. There are no interneurons.

CTRNN parameters (weights, biases and decay constants) are encoded onto a real-valued genotype and evolved for 500 generations with a population size of 50. Each population is formed from mutants of the most successful controllers of the previous generation. The initial population is made up of randomly generated genotypes. Elitism and point mutation are used with no crossover, approximating an asexual population. 10 evolutionary runs are performed (with different random




Fig. 4. Agent has 3 ray sensors and 2 motors that allow it to move horizontally in response to sensory input.

number seeds) for each type of control network on each of the two tasks. For each task, the fitness of a controller is measured as the mean value achieved over 10 trials. All fitness values are normalised to the range $[0.0, 1.0]$.

1) *Ball catching*: In each trial a series of 20 circular objects are dropped from 10 different start positions with randomly chosen velocities. The agent is awarded fitness for moving so that it ‘catches’ all the falling objects.

2) *Shape discrimination*: In each trial a series of 20 circular or diamond-shaped objects (10 of each) are dropped with the same velocity. The agent is awarded fitness for catching circles but avoiding diamonds.

IV. RESULTS

Homeostatic plasticity prevents node saturation by moving each node towards the centre-crossing condition [21], where the steepest part of its transfer function is centred on its typical level of activation. This makes the node more sensitive to input. Sensitive nodes make for sensitive networks, and this is reflected in improved signal propagation and an increased likelihood of autonomous oscillations in networks where homeostatic plasticity has been applied.

The results presented here are taken from simulations where both synaptic scaling and adaptive bias were applied simultaneously (with $\tau_w = 40$ and $\tau_b = 20$), but similar results are observed when either mechanism is applied on its own.

Figures 5 and 6 show that homeostatic plasticity significantly improves signal propagation in fully connected and randomly connected CTRNNs of various sizes. Ensembles of randomly parameterised networks were generated, and then signal propagation in each network was tested before and after a period of 500 timesteps when homeostatic plasticity was active. Signal propagation was tested by designating a single node in each network as the input node and presenting to that node an input signal drawn from a uniform distribution on the range $[-5.00, 5.00]$. The input was held fixed for 200 timesteps and then replaced with a new value drawn from the same range. The mean difference in equilibrium firing rate caused by the change in network input was measured for each node over 1000 input presentations, with the mean value used to compare networks before and after plasticity was applied. Further details of this method can be found in [12].

Figure 7 shows that homeostatic plasticity increases the likelihood of autonomous oscillations in fully connected CTRNNs.




Fig. 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 $N \in \{1, 3, 5\}$, calculated from ensembles of $\{200, 600, 1000\}$ networks of each size respectively. 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.




Fig. 6. Signal propagation in randomly connected CTRNNs (based on Erdos-Renyi random graphs [22]). 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, measured over ensembles of 1000 networks for each $P(\text{Connection}) \in \{0.0, 0.1, \dots, 1.0\}$. 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.

We measured the statistical likelihood of an oscillatory dynamic occurring in a randomly parameterised network by creating large ensembles of random CTRNNs, starting each network in a variety of different initial conditions, and observing its subsequent behaviour to see if an oscillation occurred in any part of the network. Each network was initialised from 100 randomly sampled positions in its state space. After initialisation the network was updated for 100 timesteps to allow it to settle into a stable equilibrium, after which it was observed for another 100 timesteps to determine whether this equilibrium was a fixed point or an oscillatory dynamic,² with oscillations detected as described previously. Each network was tested before and after a period of 400 timesteps in which homeostatic plasticity was active. Homeostatic plasticity was observed to significantly increase the likelihood of oscillatory dynamics. This is because it not only makes the network more sensitive but also forces all nodes in the network to play an

²No distinction was made between limit cycles and chaotic dynamics.




Fig. 7. Mean proportion of initial conditions that lead to oscillatory autonomous dynamics for fully connected N -node CTRNNs before and after homeostatic plasticity has been applied for 400 timesteps. Calculated for $N \in \{3, 5, 8, 10\}$ over ensembles of 500 networks of each size. There is a significant increase after the plasticity has been applied: dark grey shows the pre-plasticity level, light grey shows the post-plasticity level.

active part in network dynamics, increasing the number of possible oscillatory subcircuits.


Figure 8 displays the results from the evolutionary experiments. A comparison was made between the evolutionary performance of fixed CTRNNs against that of homeostatic plastic CTRNNs. Also, the effect of a developmental phase was tested, because it was noted that homeostatic plastic networks displayed long transient dynamics before settling to a stable equilibrium behaviour, and it was felt that this might affect fitness testing and preclude a fair comparison. The developmental phase consisted of a period of 6000 timesteps when homeostatic plasticity was applied to the network before it was used to control the agent.

The results differed on the two tasks (see Figure 8). The most successful network type on the ball catching task was the fixed CTRNN with a development period, while on the discrimination task the fixed CTRNN without a development period was most successful. The homeostatic plastic CTRNNs were the worst performers on both tasks; some reasons for this will be discussed in the next section.


V. DISCUSSION

We have seen that homeostatic plasticity makes CTRNNs more sensitive, improves signal propagation and makes them more likely to display oscillatory dynamics. In effect, homeostatic plasticity creates networks that are poised to behave, and we might be forgiven for expecting that this will make them more evolvable. However, when we look at the results of the evolutionary experiments we see that the situation is not so straightforward. While in one task the use of a developmental period makes fixed CTRNNs much quicker to evolve, in the other task development makes fixed CTRNNs slower to evolve. When the homeostatic plasticity is left active in the network, good controllers are always much slower to evolve.

First of all, let us consider the case where homeostatic plasticity is used as a developmental mechanism and then switched off again. The good performance of networks evolved under this scheme (i.e., fixed CTRNNs with development)



(a) Ball catching task. Fixed CTRNNs with a development phase (dotted crossed line) give the best performance, followed by fixed CTRNNs without development (solid crossed line). Homeostatic plastic CTRNNs perform poorly by comparison and have similar fitness with (dotted line) or without (solid line) a development period.



(b) Discrimination task. Best performance is by fixed CTRNNs without development (solid crossed line), followed by fixed CTRNNs with a developmental phase (dotted crossed line). Homeostatic plastic CTRNNs perform poorly, with a slight advantage offered by networks with a development period (dotted line) over those without (solid line).

Fig. 8. Fitness of different network types on ball catching and shape discrimination tasks. Mean best fitness calculated at each generation over 10 evolutionary runs of 500 generations for each network type. Maximum fitness level achievable is 1. On the ball catching task the best agent in the initial generation always had a score greater than 0.5, while on the discrimination task a fitness score of 0.5 was achieved simply by ignoring all the objects (and hence avoiding diamonds), so the fitness axis is plotted only above 0.5.

on the ball catching task can be explained by considering how the homeostatic plasticity changes network parameters. The results on improved signal propagation and increased likelihood of oscillatory dynamics show that the plasticity moves the fixed CTRNN into a region of parameter space that is populated by behaviour-rich networks. This reduces the burden placed on the genetic algorithm by providing it with a population of networks that are poised to behave. Fixed CTRNNs with development are also the best performers (in early generations) when tested on the discrimination task, lending partial support to this hypothesis. It is possible that the subsequent worse performance of the developmental fixed CTRNNs on this task is a result of phenotypic noise introduced by stochasticity in the developmental process, but this requires further clarification.

Now consider the homeostatic plastic CTRNNs, i.e., those networks where the plasticity is left switched on while the agent is performing its task. These networks consistently evolved more slowly and achieved lower fitness than the fixed CTRNNs on both tasks. This would appear to be strong evidence against the usefulness of homeostatic plasticity as an online mechanism, but there are some complications that make this conclusion less certain.

Fair comparison of the different types of control network is not straightforward. A fixed CTRNN is a significantly different dynamical system to a homeostatic plastic CTRNN. The plastic networks have more variables than the fixed networks, and these variables interact in different ways and on different timescales to those in a fixed CTRNN. This makes the dynamics of the two types of network significantly different, with the dynamics of the plastic networks typically being significantly more complex and having vastly longer transient trajectories. This makes a fair comparison between the two network types difficult, since it is likely that a fixed CTRNN will require a larger number of nodes to be of equivalent complexity to a homeostatic plastic CTRNN. If different-sized networks are compared the genotypes will be of different lengths, adding a new area of uncertainty to the comparison. These reasons, together with more commonly faced problems with evolvability testing (such as task-dependence or reliance on a particular flavour of genetic algorithm), mean that any final conclusions concerning the evolvability of homeostatic plastic networks would be premature.

Finally, it should be noted that the current view in neuroscience is that the function of homeostatic plasticity in biological nervous systems is to counteract the destabilising positive feedbacks associated with other forms of plasticity such as Hebbian learning [9]. Given this observation, it may be wrong-headed to look for an evolvability benefit from homeostatic plasticity in artificial neural systems, and perhaps research effort should instead be focussed on whether it promotes stability in the face of perturbation. These issues, together with further exploration of the evolvability of homeostatic plastic networks, will be the subject of future work.

ACKNOWLEDGMENT

Thanks to Jason Noble, Chris Buckley and the Biosystems group at the University of Leeds for useful discussions.

REFERENCES

- [1] R. Beer, "On the dynamics of small continuous-time recurrent neural networks," *Adaptive Behavior*, vol. 3, pp. 469–509, 1995.
- [2] K. Funahashi and Y. Nakamura, "Approximation of dynamical systems by continuous time recurrent neural networks," *Neural Networks*, vol. 6, pp. 801–806, 1993.
- [3] G. Turrigiano, "Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same," *Trends in Neuroscience*, vol. 22, pp. 221–228, 1999.
- [4] G. Davis and I. Bezprozvanny, "Maintaining the stability of neural function: A homeostatic hypothesis," *Annual Review of Physiology*, vol. 63, pp. 847–869, 2001.
- [5] L. Abbott and S. Nelson, "Synaptic plasticity: Taming the beast," *Nature Neuroscience*, vol. 3, pp. 1178–1183, 2000.
- [6] J. Burrone and V. Murthy, "Synaptic gain control and homeostasis," *Current Opinion in Neurobiology*, vol. 13, pp. 560–567, 2003.
- [7] N. Desai, L. Rutherford, and G. Turrigiano, "Plasticity in the intrinsic excitability of neocortical pyramidal neurons," *Nature Neuroscience*, vol. 2, pp. 515–520, 1999.
- [8] W. Zhang and D. Linden, "The other side of the engram: experience-driven changes in neuronal intrinsic excitability," *Nature Reviews Neuroscience*, vol. 4, pp. 885–900, 2003.
- [9] G. Turrigiano and S. Nelson, "Hebb and homeostasis in neuronal plasticity," *Current Opinion in Neurobiology*, vol. 10, pp. 358–364, 2000.
- [10] ———, "Homeostatic plasticity in the developing nervous system," *Nature Reviews Neuroscience*, vol. 5, pp. 97–107, 2004.
- [11] H. Williams, "Homeostatic plasticity in recurrent neural networks," in *From Animals to Animats 8: Proceedings of the 8th Intl. Conf. on Simulation of Adaptive Behavior*, S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer, Eds. Cambridge, MA: MIT Press, 2004.
- [12] H. Williams and J. Noble, "Homeostatic plasticity improves signal propagation in continuous-time recurrent neural networks," 2005, to appear in special issue of Biosystems: Proceedings of IPCAT'05.
- [13] H. Chiel and R. Beer, "The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment," *Trends in Neurosciences*, vol. 20, pp. 553–557, 1997.
- [14] A. Clark, *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press, 1997.
- [15] L. Glass and M. Mackey, *From Clocks to Chaos: The Rhythms of Life*. Princeton, NJ: Princeton University Press, 1988.
- [16] A. Ijspeert, J. Hallam, and D. Willshaw, "From lampreys to salamanders: evolving neural controllers for swimming and walking," in *From Animals to Animats 5: Proceedings of the Fifth International Conference of The Society for Adaptive Behavior (SAB98)*, R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. Wilson, Eds. Cambridge, MA: MIT Press, 1998, pp. 390–399.
- [17] H. Chiel, R. Beer, and J. Gallagher, "Evolution and analysis of model cpgs for walking I: Dynamical modules," *Journal of Computational Neuroscience*, vol. 7(2), pp. 99–118, 1999.
- [18] R. Beer, H. Chiel, and J. Gallagher, "Evolution and analysis of model cpgs for walking II: General principles and individual variability," *Journal of Computational Neuroscience*, vol. 7(2), pp. 119–147, 1999.
- [19] E. Di Paolo, "Homeostatic adaptation to inversion in the visual field and other sensorimotor disruptions," in *From Animals to Animats 6: Proceedings of the 6th Intl. Conf. on Simulation of Adaptive Behavior*, J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. Wilson, Eds. Cambridge, MA: MIT Press, 2000, pp. 440–449.
- [20] R. Beer, "Toward the evolution of dynamical neural networks for minimally cognitive behavior," in *From Animals to Animats 4: Proceedings of the 4th Intl. Conf. on Simulation of Adaptive Behavior*, P. Maes, M. Mataric, J. Meyer, J. Pollack, and S. Wilson, Eds. Cambridge, MA: MIT Press, 1996, pp. 421–429.
- [21] B. Mathayomchan and R. D. Beer, "Center-crossing recurrent neural networks for the evolution of rhythmic behavior," *Neural Computation*, vol. 14, pp. 2043–2051, 2002.
- [22] M. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, pp. 167–256, 2003.