Incorporating Energy Expenditure into Evolutionary Robotics Fitness Measures

Gary McHale[1] and Phil Husbands[1,2]
[1]Department of Informatics, University of Sussex, Brighton BN1 9QH, UK
[2]Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton BN1 9QH, UK
cglife@hotmail.com

Abstract

Evolutionary Robotics seeks to use Evolutionary Algorithms for the purpose of creating real and simulated robots. The choice of fitness functions is a key determinant in the evolved behavior exhibited by the robot. The paper introduces energy constraints into the fitness function as a preliminary investigation into seeking to influence *how* evolved robots resolve tasks rather than *what* tasks they accomplish. An experiment is described where an artificial Sensor-Motor control system (based on the GasNet model of neuromodulated neural networks), is evolved for a robot whose task is to seek and acquire balls in a physically simulated environment. The results indicate that at least for this simple task a neural network can be evolved that achieves an energy efficient solution that is at least equal in performance to a control study where energy expenditure is not included in the fitness measure. This paper seeks to make out a case for the inclusion of energy expenditures in fitness measures, as agents are under greater selection pressure to make use of available sensory systems. It is hoped that the approach outlined in this paper will be useful in helping to develop energy efficient robots, particularly in applications such as legged locomotion.

Introduction

One of the goals of Evolutionary Robotics is to find ways in which we can build intelligent robots through using methodologies and algorithms that are inspired by natural systems. In the process we also gain some insight into how natural systems operate, since a model or approach that proves suitable for solving a given task may give us some clues as to how natural systems solve similar problems.

Typically in Evolutionary Robotics simple experiments involving sensor/motor tasks are solved through the evolution of artificial neural networks (Beer, 1995; Urzelai and Floreano, 2000; Nolfi and Floreano, 2001; Bongard and Paul, 2000). However, more complex locomotive problems have also been addressed (Reil and Husbands, 2002; Vaughan et al., 2004). Other researchers have sought to compare the suitability of alternative neural network models in evolutionary simulations (Tuci and Quinn, 2003; McHale and Husbands, 2004b).

Often in such simulations the energy constraints placed on a simulated organism are implicit, either due to the physical design of the robot and its power-supply, or through parameters chosen for actuators in physical simulation. Typically there is no cost penalty associated with higher energy use in the completion of a task, nor benefit in engaging in energy efficient activity. As a consequence the behavior exhibited by the evolved robots is often not typical of behavior exhibited by living organisms, which clearly do have energy budgets determined by their activities and metabolism. This problem is of interest to Evolutionary Roboticists for two reasons; clearly energy efficient robots have greater economic value than those that are wasteful, and it may be more difficult to evolve intelligent behavior where energy supplies are relatively unconstrained.

Energy and Life

At the macro-level Energy processes are of major importance in driving Ecological Systems (Jorgensen and Bendoricchio, 2001) whilst at the micro-level, energy governs the biological reactions that support life (Haynie, 2001). Within this broad spectrum of work, Biophysics and Biomechanics are amongst the most relevant areas to scientists seeking to evolve life-like robots or artificially simulated creatures. Interesting examples of the application of Biophysics to predict animal behavior include that by D. M. Gates, who uses energy based models to predict lizard activity patterns, and predator-prey relationships (Gates, 2003). In the field of Biomechanics the extensive works of R. McNeill Alexander, are of particular relevance to scientists considering locomotion in living organisms from an energy perspective (Alexander, 2003).

There are three key questions that are of pragmatic interest to researchers in Evolutionary Robotics, these are:

1. How do energy constraints influence behavior?

2. How can we generate appropriate locomotor control systems that are energy efficient?

3. Does the imposition of energy constraints make it easier or more difficult to evolve effective integrated sensor-motor control systems?

Quite clearly the questions raised above are very broad in their scope. This paper describes a starting point and a basic methodology from which it is hoped that we can start to address some of these issues in greater detail. The focus of this paper is on the effect of the imposition of energy constraints on the relative utilization of motor and sensor facilities in solving a simple task.

Experimental Setup

Previous work (Husbands et al., 1998) involved evolving a GasNet based neural network for a real and simulated Robot, with the goal of moving towards a triangle and ignoring a rectangle. The experiment reported here differs in its focus on energy efficiency rather than the successful execution of a simple sensor-motor task. The following section describes the experimental set-up, together with details on the GasNet implementation and the evolutionary algorithm used. Whilst the following section describes the key experimental parameters, additional parameters that relate to the generation of the GasNet morphology are discussed in more detail in the source cited above.

The Robot and Its Environment

A model was made of a “toy” robot with minimal sensor and motor capacity. The robot exists in a physically simulated 3d world. Motion is achieved through the application of linear and rotational forces to the robot at its center of mass. A GasNet (described in more detail below) comprising 16 nodes is evolved to provide the motor signals for the application of forces to the robot. Each “node” represents a neuronal cell. Four of these nodes act as motor neurons. The rotational torque is determined by the sum of two motor output neurons. Torque is applied to the robot around a vertical axis centered at the robots center of mass. The linear force is the sum of two motor output neurons, oriented in the robots “forward” direction, and passing through the robots center of mass. As a result, the robot is capable of rotating clockwise and counter-clockwise, as well as moving forwards and backwards.

The robot has four sensors. These comprise raycasts into the physically simulated world. This data is minimally preprocessed before passing to two sensor neuron inputs. One sensor input simply registers whether or not an object has been hit. Any of the rays striking an object will result in an input value of +1 applied to the sensor neuron input, and -1 when no objects are detected. The second sensor neuron receives a value that corresponds to the average distance from the ray source to the detected object across all sensor rays when an object is within sensor range. If a sensor ray detects no object, then the distance measured by the robots’s sensor rays is taken to be the maximum value of the raycast’s sensor range. The average distance value for all rays is mapped to value within the range [+1,-1]. Effectively we have two sensor modalities; the first neuron will detect any objects within the robots “visual” range, the second sensor neuron will be activated when the robot is close to a target object. The raycast sensors and robot can be seen in Figure 1.

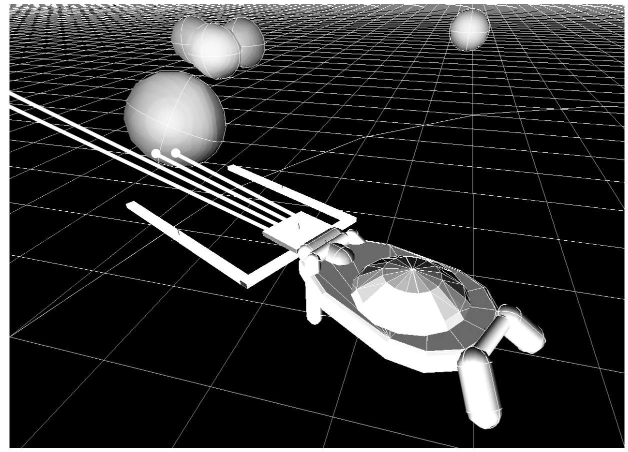

Figure 1: Side View of Robot and Sensor Rays

The robot exists in a planar world within which, in initial experiments, 24 large spheres are placed. The “jaws” of the robot are slightly wider than the width of the spheres. The width of the robot jaws is 16 units, the spheres have a diameter of 15 units. The spheres are randomly distributed within a two-dimensional annulus. An additional “exclusion corridor” (70 units wide) is created, such that a robot traveling straight forward or backwards will not collide with any spheres. The length of the robot is approximately 48 units, the range of sensor rays 100 units, the inner radius of the annulus is 95 units and the outer radius 400 units. The aim of the task is to capture a sphere in the robot’s jaws.

Trial Description and Fitness Function

Each trial consists of 10 sub-trials. At the start of each subtrial the robot is placed at the origin (the center of the annulus), and 24 spheres are randomly distributed within the annulus (with the exception of the exclusion corridor described above). Typically one or two spheres would be within sensor range if the robot were to rotate 360 degrees around its vertical axis. Due to the high variability in the distribution of spheres in each sub-trial, a large number of sub-trials are required to establish a representative fitness value. Figure 2 shows the distribution of spheres for a typical sub-trial.

There are two forms of the fitness function used in this experiment. The first form is used in tests where there is no cost penalty for using the maximum energy available to the robot. In this case the fitness value of the robot is taken to be the relative closeness of the sphere as measured by the robots sensor rays (measured as a fraction of the total ray distance) at the termination of a trial. The trial is terminated if the robot closes to within 10 percent of its total ray distance

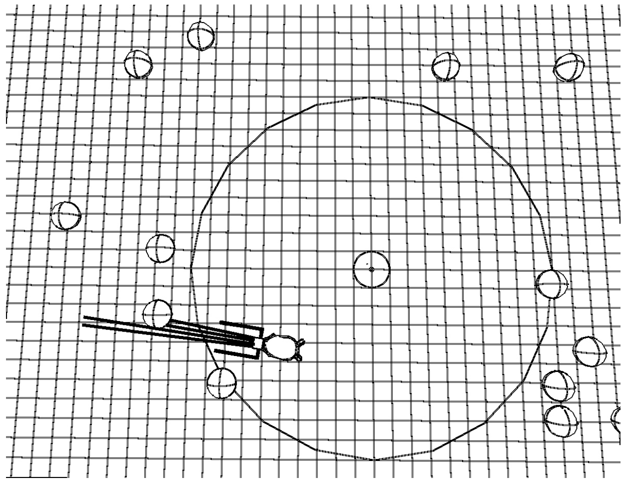

Figure 2: Top View of Robot in its Environment

(i.e. 10 units from the origin of the sensor rays). This corresponds to the sphere entering completely into the robots “jaws” (which constitute a channel that is 20 units deep). Each sub-trial is weighted to be a 10th of the total fitness sum of the trial. The maximum theoretical fitness of the robot is 1.0 corresponding to 10 trials where the robot acquires a sphere target in its jaws. However, due to the early termination condition, practically the maximum fitness attainable is 0.9 plus a small value corresponding to the distance traveled in one physically simulated time-step prior to the early termination condition being detected.

The alternative form of the fitness function penalizes the robot for the excessive use of energy. Energy expenditure of the robot is measured indirectly by considering the energy of the entire system. At each time step the kinetic and potential energy of the system is calculated for all objects. The total energy of the system at the prior time-step is saved, so that we can observe any increases in the total energy of the system (E_c in Equation 1 below). The total energy of the system can decline due to energy dissipation occurring in non-elastic collisions, or in work done against friction. Since the only source of energy in the system is that provided by the application of forces to the robot, all positive increases in the energy of the system are attributed to energy expenditure by the robot. Since there are no springs in the robot, we do not need to consider the transference of kinetic energy to that of potential energy in joints.

This approach yields only an approximate estimate of energy used by the simulated robot, as it neglects work done by the robot in decelerating bodies in the system. However, it is likely to be sufficient for making estimates of relative energy expenditure between simulated robots when averaged over 10 sub-trials. We use this value (E_c) to calculate an *Energy Efficiency Factor* ρ, that tends to one for low energy usage and zero for high energy usage in accordance with the function;

$$\rho = \frac{e^{-(E_c - k_1)/k_2}}{(1 + e^{-(E_c - k_1)/k_2})} \qquad (1)$$

Where:

ρ is the Energy Efficiency Factor.

E_c is the cumulative positive changes in kinetic plus potential energy of the system.

k_1 is a bias term chosen to be 600.

k_2 a scale term that is chosen to be 150.

The Energy Efficiency Factor is multiplied with the fitness term used to calculate fitness in systems that are not energy constrained. For example, in the case where the robot acquires 10 spheres in 10 trials and attains an unadjusted fitness value of approximately 0.9 (due to the early termination condition), this value is multiplied by the Energy Efficiency Factor ρ. If the Energy Efficiency Factor evaluates to 0.5, then the fitness of the robot is taken to be 0.45. The values of k_1 and k_2 were chosen heuristically so that a wide range of sensor-motor “behaviors” were not excessively penalized, but those that appeared to rely largely on exploiting the maximum motor output capacities of the robot were heavily penalized. The values were chosen such that in an initial population most individuals would typically exhibit an Energy Efficiency factor between 1.0 and 0.5. Robots that made continuous use of the maximum linear and rotational forces available to the robot would typically exhibit an Energy Efficiency factor of less than 0.1.

Artificial Neural Network - GasNets

GasNets draw their inspiration from the biological action of Nitric Oxide in neural systems (Husbands et al., 2001), as such they are an abstraction of a neuromodulated neural network. The nodes exist in a 2-dimensional plane, where node position, connectivity, gas emission and sensitivity characteristics are under evolutionary control. Equation 2 shows how the transfer function of each node is affected by local gas concentration.

$$y_i^{t+1} = \tanh[k_i^t(\sum_{j \varepsilon M_i} \omega_{ji}\sigma(y_j^t + I_i)) + b_i] \qquad (2)$$

Where:

$y_i t+1$ is the activation of the i'th node at time $t+1$.

y_i^t is the activation of the i'th node at time t.

k_i^t is a time-varying transfer function modulator. The value of k varies with gas concentrations at the i'th node, see equation 6.

M_i is the set of all nodes that have an input to the i'th node.

I_i a sensor input to the i'th node.

b_i a bias term for the i'th node where $b_i \in [-2, 2]$.

ω_{ji} is the weight of the output from the j'th node to the i'th node where $\omega \in [-4.0, 4.0]$.

σ is the logistic activation function.

In comparative studies GasNets have shown themselves to be more *evolvable* than comparable networks that do not incorporate gas modulation, in simulation and when used in real robots (Smith et al., 2003). Other studies have shown them to excel in the evolution of networks given the task of supporting bipedal (McHale and Husbands, 2004a) and quadrupedal locomotion.

The GasNet diffusion model is controlled by two genetically specified parameters; the radius of influence r and the rate of build up and decay s. Spatially, the gas concentration varies as an inverse exponential of the distance from the emitting node with a spread governed by the radius of influence (r), with the concentration set to zero for all distances greater than r (Equation 3).

$$C(d,t) = \begin{cases} e^{-2d/r} \times T(t) & d < r \\ 0 & \text{else} \end{cases} \tag{3}$$

Here C(d,t) is the concentration at a distance d from the emitting node at time t. t_e is the time at which emission was last turned on, t_s is the time at which emission was last turned off, and s (controlling the slope of the function T) is genetically determined for each node. The total concentration at a node is then determined by summing the contributions from all other emitting nodes (nodes are not affected by their own concentration, to avoid runaway positive feedback).

The maximum concentration at the emitting node is 1.0 and the concentration builds up and decays from this value linearly as defined by Equations 4 and 5 at a rate determined by s.

$$T(t) = \begin{cases} H\left(\frac{t-t_e}{s}\right) & \text{emitting} \\ H\left[H\left(\frac{t_s-t_e}{s}\right) - H\left(\frac{t-t_s}{s}\right)\right] & \text{not emitting} \end{cases} \tag{4}$$

$$H(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x < 1 \\ 1 & \text{else} \end{cases} \tag{5}$$

In the basic GasNet model there are two 'gases', one whose modulatory effect is to increase the transfer function gain parameter (k_i^t from equation 2) and one whose effect is to decrease it. It is genetically determined whether or not any given node will emit one of these two gases (gas 1 and gas 2), and under what circumstances emission will occur. The concentration-dependent modulation is described by Equation 6, with transfer parameters updated on every time step as the network runs.

$$k_i^t = k_i^0 + \alpha c_1^t - \beta c_2^t \tag{6}$$

where k_i^0 is the genetically set default value for k_i, c_1^t and c_2^t are the concentrations of gas 1 and gas 2 respectively at node i at time t, and α and β are constants. Both gas concentrations lie in the range $[0, 1]$. The values of c_1^t and c_2^t are calculated using equation 3.

Genetic Algorithm

A distributed steady-state GA is used, utilizing a 2-dimensional grid of 100 individuals, with a tournament size corresponding to three individuals. A *principal* is selected, followed by two *neighbors*. These neighbors are selected based on a random walk (of length in the range [1,4] grid cells) originating at the principal. If the principal is fitter than both neighbors the weakest individual is replaced by a mutated version of the principal. If not, then the weakest member of the tournament group is replaced by the fitter two individuals genes, using single-point crossover, followed by mutation. The replacement of 100 individuals corresponds to a single pseudo-generation. Each run was allowed to continue for 80 pseudo generations.

Network parameters are stored on a node basis. Each gene comprises a list of real valued and integer parameters (comprising 16 parameters per node). Mutation takes place either after recombination, or after cloning of the principal tournament member (as described earlier). Mutation takes place at 20 percent of the nodes (rounded to 3 in a 16 cell network) selected at random. A single mutation event will result in the mutation of a single real or integer parameter in each of the randomly selected nodes. The magnitude of this mutation corresponds to 4 percent of the real valued parameters range with a probability of 0.2, and 1 percent of the parameters range with a probability of 0.8. In the case of integer parameters we follow a similar strategy of small mutations with a probability of 0.8 and large mutations with a probability of 0.2.

Experimental Results: Initial Results

Initially three scenarios were tested;

1. No Energy Penalty, Ray Sensors.
2. Energy Penalty, Ray Sensors.
3. No Energy Penalty, No Ray-Sensors.

The graphs shown in Figure 3 and Figure 4 show the performance of the robot with and without the fitness function which penalizes excessive energy use (scenarios 2 and 1 respectively). They display the average results of 5 trial runs over 80 generations. The maximum theoretical fitness is 1.0, although this is practically limited to nearer 0.9 due to the early exit condition utilized. In both cases average and peak absolute fitness values were very similar. However, robots

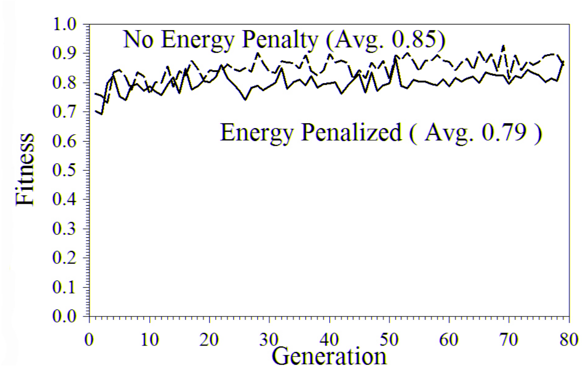

Figure 3: Absolute Fitness

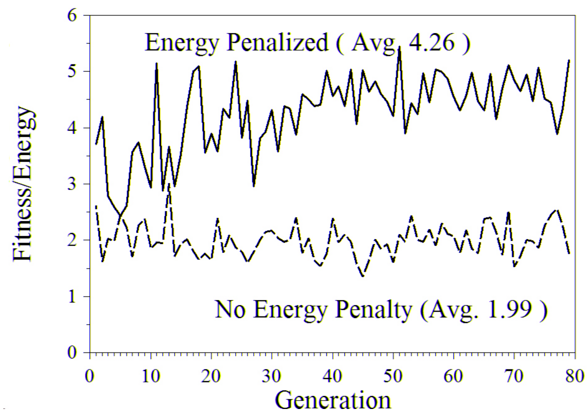

Figure 4: Fitness Per Energy Expended

that were not penalized for the excessive use of energy expended nearly twice the energy of robots that were subject to penalization in the case of excessive energy use. This demonstrates that we were able to achieve results that are considerably more energy efficient by incorporating an energy penalty without sacrificing absolute performance. It is suggested by these results that consideration of energy expenditures in evolving gaits for legged robots may well benefit from the use of fitness function that takes into account energy expenditure.

Robots that were not energy constrained did not appear to be making use of sensory data in achieving their high level of fitness, but appeared to rely upon energy intensive motor activity. The fittest individuals achieved a high fitness score by rotating rapidly whilst spiralling outwards from the origin. We may have expected robots to exhibit angular deceleration and an increase in forward motion towards detected spheres, if they were making use of their sensory data.

In order to test this hypothesis, we carried out further tests where sensors were disabled (in robots that were not subject to energy penalties), listed above as Scenario 3. Over 60 generations an average fitness level of 0.849 (versus 0.842) and an efficiency (fitness over energy expended) of 2.01 (versus 1.94) were achieved. Qualitatively the behaviors exhibited by the fittest individuals with sensors disabled, were indistinguishable from those with sensors enabled. The qualitative and numerical similarity of results in Scenario 1 and 3, suggest that the robots which were not not subject to energy penalties made little use of their sensory data.

Experimental Results: Secondary Results

Initial results suggested that there were two possible routes by which high levels of fitness could be attained; either with energy intensive motor activity or through active use of sensor data. In order to explore this hypothesis further a second set of scenarios were tested. This time the results were calculated over 35 generations with 5 trials for each scenario. It was hypothesized that factors affecting the relative importance to fitness of exploiting sensor data might include the “utility” of the sensory data, and the sparseness of “prey” in the environment (in this case prey being the spheres themselves). In order to investigate this, the value of the sensor data was improved, by inhibiting rotational movement on the detection of a sphere. It was assumed that an optimal low-energy strategy would involve a two stage process, where the robot would initially orient in the direction of a detected sphere and then engage in linear motion towards the sphere. In effect, by coupling the raycast sensors to lateral inhibition, we were attempting to increase the ‘value’ of sensor inputs. Additionally the number of spheres was reduced from 24 to 18. Under these experimental conditions the following scenarios were tested and results obtained:

1. No Energy Penalty, No Sensor Data; Average Fitness 0.77, Fitness/Energy 1.33.

2. Energy Penalty, No Sensor Data; Average Fitness 0.69, Fitness/Energy 2.36.

3. Energy Penalty, Sensor Data; Average Fitness 0.89, Fitness/Energy 6.03.

These results help to corroborate the idea that at least under these experimental circumstances energy intensive motor activity (in Scenario 1 above), can more than compensate for a lack of sensory data (Scenario 2), since we are able to achieve a higher level of fitness without sensor data, when there is no penalty for excessive energy use. In Scenario 3, we see however that the improved value of the sensor data ensures that in absolute terms, even with energy intensive motor activity, the same absolute levels of fitness cannot be achieved. This preliminary result suggests that incorporating energy penalties into fitness measures may be a useful strategy in encouraging the use of sensors in evolved behaviors which in turn may help with the evolution of more complex sensorimotor capabilities.

Future Work

An obvious extension of this work is to incorporate measures of energy efficiency into fitness functions used for the evolution of locomotion in legged robots. Simulations involving the evolution of locomotion in legged robots typically employ fitness functions that assign the highest fitness to solutions that enable the robot to travel the furthest distance in a fixed amount of time. Clearly this biases results to gaits that make use of the maximum available energy. Experiments have shown horses select the most energy efficient gait appropriate to a given speed (Hoyt and Taylor, 1981). Similar results have also been found for humans, and kangaroos (Alexander, 2003). Incorporating energy efficiency into fitness functions should allow us to evolve gaits that are optimized for energy efficiency.

Another challenge is in evolving neural circuitry that is capable of autonomously switching between gaits to minimize energy expenditure over a range of speeds. There are two parts to this problem. This first part relates to a mechanism of switching, so that we can modify neural networks dynamically to produce the required range of gaits. One possible solution is to continue with neural networks that incorporate models of neuromodulation. Neuromodulators such as Dopamine, Octopamine and Serotonin are known to have the capacity to chemically "re-wire" motor circuits (Kiehn and Katz, 1999). A modification of the GasNet model described in this paper may be of value in developing such switching circuits.

The second part of the problem relates to a requirement to provide the robot with a simple metabolism, or at least a method of providing some input into the artificial neural network that reflects energy expenditure. By incorporating the "proprioception" of energy expenditure within the model, we are providing an evolutionary pathway by which energy conservative gaits can be evolved. Neural sensors that detect excessive energy use could be used as switches to trigger alternative motor gaits.

Discussion

We can perhaps imagine a continuum of experimental scenarios with variations in the scarcity of "prey", and the acuity and utility of sensors. In an environment where prey is abundant, and there are readily accessible supplies of energy, then the value of energy intensive motor activity in improving fitness may well diminish the importance of sensors. Conversely in an environment with sparse prey, the relative value of sensors in improving fitness is increased. This has repercussions if our primarily goal is to seek to evolve agents that make full use of sensory data in the solution of a task.

The imposition of an energy constraint changes the fitness landscape such that robots that *do* make use of sensory data, have an evolutionary advantage. A failure to impose energy penalties in evolutionary simulations reduces selection pressure on evolved entities, such that they may not necessarily take full advantage of the sensory data that is accessible to them, but may discover energy intensive motor solutions to achieve the same effective fitness. This is a strong indication that imposing energy penalties may well play a useful role in helping us to achieve more subjectively "intelligent" agent activity.

Acknowledgments

The authors would like to thank the anonymous reviewers for their extensive comments, suggestions and corrections.

References

Alexander, R. M. (2003). *Principles of Animal Locomotion*. Princeton University Press, Princeton (NJ).

Beer, R. (1995). On the Dynamics of Small Continuous-Time Recurrent Neural Networks. *Adaptive Behavior*, 3(4):469–509.

Bongard, J. and Paul, C. (2000). Investigating Morphological Symmetry and Locomotive Efficiency using VirtualEmbodiedEvolution. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S., editors, *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior (SAB 2000)*, pages 420–429. MIT Press.

Gates, D. M. (2003). *Biophysical Ecology*. Dover Publications, London.

Haynie, D. T. (2001). *Biological Thermodynamics*. Cambridge University Press, Cambridge, Englsnd.

Hoyt, D. and Taylor, C. (1981). Gait and the energetics of locomtion in horses. *Nature*, 292:239–240.

Husbands, P., Phillipedes, A., Smith, T., and O'Shea, M. (2001). Volume Signalling in Real and Robot Nervous Systems. *Theory in Biosciences*, 120:253–269.

Husbands, P., T.M.C.Smith, Jakobi, N., and O'Shea, M. (1998). Better Living Through Chemistry: Evolving GasNets for Robot Control. *Connection Science*, 10(3-4):185–210.

Jorgensen, S. and Bendoricchio, G. (2001). *Fundamentals of Ecological Modelling, Third Edition (Developments in Environmental Modelling*. Elsevier Science, Oxford, England.

Kiehn, O. and Katz, P. (1999). Making circuits dance: the neuromodulation of motor systems. In Katz, P., editor, *Beyond Neurotransmission*, pages 275–313. Oxford University Press.

McHale, G. and Husbands, P. (2004a). GasNets and other Evolvable Neural Networks applied to Bipedal Locomotion. In Cliff, D., Husbands, P., Meyer, J., and Wilson, S., editors, *From Animals to Animats 8: Proceedings of the Eighth International Conference on Simulation of Adaptive Behavior (SAB 2004)*, pages 205–224. MIT Press.

McHale, G. and Husbands, P. (2004b). Quadrupedal Locomotion: GasNets, CTRNNs and Hybrid CTRNN/PNNs Compared. In J. Pollack, M.A. Bedau, P. H. T. I. and Watson, R., editors, *Artificial Life IX : Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 106–116. MIT Press.

Nolfi, S. and Floreano, D. (2001). *Evolutionary Robotics: Biology, Intelligence and Technology of Self-OrganizingMachines*. MIT Press, Cambridge, MA.

Reil, T. and Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6:159–168.

Smith, T., Husbands, P., Philippedes, A., and O'Shea, M. (2003). TemporallyAdaptiveNetworks:Analysis of GasNet Robot Controllers. In Standish, R., Bedau, M., and Abbass, H. A., editors, *Artificial Life VIII: Proceedings of the Eight International Conference on Artificial Life*, pages 274–282. MIT Press.

Tuci, E. and Quinn, M. (2003). Behavioural Plasticity in Autonomous Agents: A Comparison Between Two Types of Controller. In Cagnoni, S., Cardalda, J., Corne, D., Gottlieb, J., Guillot, A., Hart, E., Johnson, C., Marchiori, E., Middendorf, J.-A. M. M., and Raidl, G., editors, *Applications of Evolutionary Computing: EvoWorkshops 2003*, pages 661–672, Berlin. Springer-Verlag.

Urzelai, J. and Floreano, D. (2000). Evolutionary Robots with Fast Adaptive Behavior in New Environments. In Thompson, A., editor, *Third International Conference on EvolvableSystems: From Biology to Hardware (ICES2000)*. Springer Verlag.

Vaughan, E., Paolo, E. D., and Harvey, I. (2004). The Evolution of Control and Adaptation in a 3D Powered Passive Dynamic Walker. In J. Pollack, M.A. Bedau, P. H. T. I. and Watson, R., editors, *Artificial Life IX : Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, pages 139–147. MIT Press.