

Quadrupedal Locomotion: GasNets, CTRNNs and Hybrid CTRNN/PNNs Compared

Gary McHale¹ and Phil Husbands²

¹Department of Informatics, University of Sussex, Brighton BN1 9QH, UK

²Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton BN1 9QH, UK
cglife@hotmail.com

Abstract

Evolutionary Robotics seeks to use evolutionary techniques to create both physical and physically simulated robots capable of exhibiting characteristics commonly associated with living organisms. Typically, biologically inspired artificial neural networks are evolved to act as sensorimotor control systems. These networks include; GasNets, Continuous Time Recurrent Neural Networks (CTRNNs) and Plastic Neural Networks (PNNs). This paper seeks to compare the performance of such networks in solving the problem of locomotion in a physically simulated quadruped. The results in this paper, taken together with those of other studies (summarized in this paper) help us to assess the relative strengths and weaknesses of these three different approaches.

Introduction

Evolutionary Robotics has two broad goals; to gain an insight into biological systems through suitable abstractions of these systems, and secondly to seek to discover techniques that are of economic value in the development of robots and physically simulated agents. The underlying methodology involves the use of Genetic Algorithms (GAs) to evolve Artificial Neural Networks (ANNs) that act as sensorimotor control systems in real and simulated robots.

There is an increasing body of work that has successfully applied these techniques to the evolution of ANNs suitable for controlling bipedal and quadrupedal locomotion. This includes earlier work (Reil and Husbands, 2002) which demonstrates that it is possible to evolve a bipedal motor control in a physically simulated agent using a conventional Dynamic Recurrent Neural Network (DRNN) without sensor input. Bongard and Paul have evolved bipedal locomotion in a physically simulated agent through genetic encoding that comprises morphological as well as ANN parameters (Bongard and Paul, 2001). Researchers have evolved bipedal locomotion in a physically simulated robot that incorporates a model of neuromodulation (Ishiguro et al., 2003). Billard and Ijspeert have been successful in evolving Quadrupedal locomotion in a real robot (Billard and Ijspeert, 2000).

Unfortunately, there are remarkably few comparative studies that enable us to judge which of the diverse approaches taken by different researchers is the most expedient when applied to a specific problem. It is hoped that this paper provides a useful insight into the relative performance of some of the more common neural network varieties. This paper is intended as a complementary study to work carried out by the authors (McHale and Husbands, 2004) which sought to identify the relative performance of 14 different network varieties applied to bipedal locomotion.

Network Descriptions

Put in the simplest terms, Continuous Time Recurrent Neural Networks (CTRNNs) (Yamauchi and Beer, 1994) represent the “plain vanilla” form of DRNN’s, GasNets represent an approach to incorporate neuromodulation into a form of DRNN (Husbands et al., 1998), and Plastic Neural Networks (PNNs) seek to incorporate Hebbian dynamics (Floreato and Mondada, 1996). The particular variants used in this experiment are described in detail in this following section.

One thing that should be noted is that for each network, network morphology has been constrained to correspond more closely to that associated with coupled-oscillator circuitry. All networks comprise a total of 16 nodes or cells. An initial population is seeded with networks that have a single symmetry axis, such that we have two subnetworks of 8 nodes, each with identical parameters. In the case of GasNets (where nodes have a physical location in a 2 dimensional plane) the position of each node from one subnetwork is mirrored in an axis that divides the plane. This is shown more clearly in Figure 1. The mirrored nodes are interconnected via mutually inhibitory connections. Whilst initial populations comprise symmetrical networks, mutation and crossover results in the introduction of asymmetries over a period of time (symmetry is only enforced in the initial population).

Copyrighted Material




Figure 1: Schematic of the symmetrical distribution of GasNet nodes.

Center-Crossing CTRNNs

The characteristic equation of the conventional CTRNN (Beer, 1995) is shown below;

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i} (-y_i^t + \sum_{j=1}^N \omega_{ji} \sigma(y_j^t + \theta_j) + I_i) \quad (1)$$

$$i = 1, 2, \dots, N$$

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 .

τ_i is the time constant for the i 'th node calculated according to equation 18.

I_i a sensor input to the i 'th node where I is either 1 (in contact with the floor) or 0 (not in contact with the floor).

θ_j a bias term for the j 'th node where $\theta \in [-2, 2]$.

T is the time slice (in this case T is set to 1).

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.

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

The network is fully interconnected. Node connection weights and bias are under evolutionary control. This study uses a variant of the conventional CTRNN, referred to as the Center-Crossing CTRNN (Mathayomchan and Beer, 2002), where initial biases are calculated such that:

$$\theta_i = \frac{-\sum_{j=1}^N \omega_{ji}}{2} \quad (3)$$

Mathayomchan and Beer suggest that populations seeded with center-crossing networks may be more likely to yield a wider range of dynamics than a population of random networks.

The incorporation of a single symmetry axis (as used in this study) results in a final form, as described by;

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i} (-y_i^t + \sum_{j=1}^N \omega_{ji} \sigma(y_j^t + \theta_j - z_i^t) + I_i) \quad (4)$$

$$i = 1, 2, \dots, N$$

Where:

z_i^t is the activation of the corresponding i 'th node in the symmetrical subnetwork at time t .

N is the number of nodes in each subnetwork (in this case 8).

GasNets

GasNets are an example of a class of ANNs that seek to model aspects of neuromodulation. A key attribute of the GasNet model is that the transfer characteristics of network nodes are modified via the influence of diffused gases (modeled in a 2-dimensional plane). This network model is inspired by the action of Nitric Oxide in biological systems (Husbands et al., 2001). Earlier work has shown that GasNets are more *evolvable* than comparable networks that do not incorporate gas modulation, in simulation and when used in real robots (Smith et al., 2003).

In GasNets, node transfer functions can be modulated by local gas concentrations in the vicinity of the node. Nodes can also act as chemical emitters, under either gas or electrical stimulation. GasNet nodes exist in a geometric plane where internode distances determine gas concentrations and (in conjunction with additional genetic parameters) network connectivity. Under typical evolutionary parameters the GasNet connectivity rules result in a sparsely connected network.

$$y_i^{t+1} = \tanh[k_i^t (\sum_{j \in C_i} \omega_{ji} \sigma(y_j^t + I_i)) + b_i] \quad (5)$$

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 9.

C_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 $b_i \in [-2, 2]$.

The original GasNet diffusion model (upon which this implementation is based) is controlled by two genetically specified parameters, namely 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 r , with the concentration set to zero for all distances greater than r (Equation 6). 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 7 and 8 at a rate determined by s .

$$C(d,t) = \begin{cases} e^{-2d/r} \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (6)$$

$$T(t) = \begin{cases} H\left(\frac{t-t_e}{s}\right) & \text{emitting} \\ H\left[\frac{t-t_e}{s}\right] - H\left(\frac{t-t_s}{s}\right) & \text{not emitting} \end{cases} \quad (7)$$

$$H(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x < 1 \\ 1 & \text{else} \end{cases} \quad (8)$$

where $C(d,t)$ 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).

For mathematical convenience, in the basic GasNet there are two 'gases', one whose modulatory effect is to increase the transfer function gain parameter (k_i^t from equation 5) 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 (either when the 'electrical' activation of the node exceeds a threshold, or the concentration of a genetically determined gas in the vicinity of the node exceeds a threshold. Note these emission processes provide a coupling between the 'electrical' and 'chemical' mechanisms). The concentration-dependent modulation is described by Equation 9, with transfer parameters updated on every time step as the network runs.

$$k_i^t = k_i^0 + \alpha C_1^t - \beta C_2^t \quad (9)$$

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 $[0, 1]$. Thus the gas does not alter the electrical activity in the network directly but rather acts by continuously changing the mapping between input and output for individual nodes, either directly or by stimulating the production of further virtual gas. The concentration dependent modulation can, for instance, change a node's output from being positive to being zero or negative even though the input remains constant. Any node that is exposed to a non zero gas concentration will be modulated. This set of interacting processes provides the potential for highly plastic systems with rich dynamics.

CTRNN/PNN Hybrid

One of the underlying concepts associated with Plastic Neural Networks is that there is value in evolving artificial neural networks that are capable of exhibiting learning through ontogenetic change (Floreano and Mondada, 1996). Let us first start with a description of a basic PNN (Urzalai and Floreano, 2000). A key characteristic of PNN's is that connection weights vary over time based on Hebbian learning rules given by:

$$\omega_{ji}^t = \omega_{ji}^{t-1} + \eta \Delta \omega_{ji} \quad (10)$$

Where η is a learning rate ($0.0 < \eta < 1.0$) and ω_{ji} is the connection weight of the input to node i from node j . The adaptation rule $\Delta \omega_{ji}$ is genetically determined for each node. All inputs to a given node are subject to the same adaptation rule (referred to as node encoding by the original authors).

Where x is the activation of node j , which is an input to node i (which has an output activation of y), the adaptation rule is one of:

Plain Hebb Rule

$$\Delta \omega_{ji} = (1 - \omega_{ji}) x_j y_i \quad (11)$$

Post-Synaptic Rule

$$\Delta \omega_{ji} = \omega_{ji} (-1 + x_j) y_i + (1 - \omega_{ji}) x_j y_i \quad (12)$$

Pre-Synaptic rule Rule

$$\Delta \omega_{ji} = \omega_{ji} x_j (-1 + y_i) + (1 - \omega_{ji}) x_j y_i \quad (13)$$

Covariance Rule

$$\Delta \omega_{ji} = \begin{cases} (1 - \omega_{ji}) & \text{if } F(x_j, y_i) > 0 \\ (\omega_{ji}) F(x_j, y_i) & \text{otherwise} \end{cases} \quad (14)$$

Where:

$$F(x_j, y_i) = \tanh(4(1 - |x_i - y_j| - 2)) \quad (15)$$

All nodes in the PNN are fully interconnected. The rate of learning η can only assume one of four values (0.0, 0.3,

Copyrighted Material

0.6, 0.9). The characteristic equation for the PNN is shown below:

$$y_i^{t+1} = \sigma\left(\sum_{j=1}^N \omega_{ji}^t(y_j^t)\right) + I_i \quad i = 1, 2, \dots, N \quad (16)$$

Where:

ω_{ji}^t is the adaptive weight for the j 'th input to the i 'th node.

σ is the standard logistic activation function.

I_i a sensor input to the i 'th node where I is either 1 (in contact with the floor) or 0 (not in contact with the floor).

The CTRNN/PNN Hybrid used in this study is a variation on the conventional PNN, where activation signals are further modified by a node based time constant under evolutionary control (in a similar fashion to conventional CTRNNs). This model was first introduced by the authors of this paper in an attempt to create a PNN that exhibited richer frequency dynamics than those of the conventional PNN (McHale and Husbands, 2004). The range of y_i is $[0, 2]$ for input neurons and $[0, 1]$ for hidden and output neurons (Blynel and Floreano, 2002). It is modified again here such that the network comprises two symmetrical sub-networks, with mutual inhibition;

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i} (-y_i^t + \sum_{j=1}^N \omega_{ji}^t \sigma(y_j^t + \theta_j - z_i^t) + I_i) \quad (17)$$

$$i = 1, 2, \dots, N$$

Where:

ω_{ji}^t is the adaptive weight for the j 'th input to the i 'th node.

z_i^t is the activation of the corresponding i 'th node in the symmetrical subnetwork at time t .

N is the number of nodes in each subnetwork.

Genetic Algorithm

The Genetic Algorithm (GA) used in this experiment is the same as that used in the aforementioned comparative study. The population comprises a 2-dimensional grid of 100 individuals. A distributed steady-state GA is used, with a small 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 individual, then 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.

single-point crossover, followed by mutation. The replacement of 100 individuals corresponds to a single pseudo-generation. A total of 14 evolutionary runs were carried out for each network type. Each trial was allowed to continue for 50 pseudo generations.

The genetic encoding strategy follows a similar approach for all networks. Network parameters are stored on a node or cell basis. Each gene comprises a list of real valued and integer parameters (comprising 16 parameters per node for a conventional GasNet for example). Connection weights (where relevant) are also stored on a per node basis.

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. These mutation parameters were chosen in preliminary experiments to avoid premature convergence and maintain a reasonable degree of phenotypic diversity across the different network varieties during evolution.

CTRNNs undergo further mutation. Each randomly selected cell has all of its weights mutated (again by a factor of 4 percent with a 20 percent probability and 1 percent with an 80 percent probability).

Time constant initialization was devised to yield a wide range of values. An exponent f was randomly selected from the set:

$$f \in [-10, -8, -6, -4, -2, 0, 2, 4, 8, 10]$$

A second random variable $r \in [0.0, 1.0]$ was then used to scale the value such that the time τ constant is calculated from:

$$\tau = 1.0 + r(10^f) \quad (18)$$

The time constant mutation operator increments or decrements the exponent by 1 with a probability of 0.2, and generates a new value of $r \in [0.0, 1.0]$ with a probability of 0.8.

Experimental Setup

A screen shot of the quadruped used in this experiment is shown in Figure 2. Whereas the previous (bipedal locomotion) study used a physics package called AutoSim, this study uses an open source package called Open Dynamics Engine (ODE).

The quadruped torso is simulated with 6 physical degrees of freedom (unlike the previous biped study, where the biped was physically incapable of falling sideways). The quadruped comprises 9 rigid bodies, two rigid bodies for




Figure 2: Physically Simulated Quadruped

each leg, and a single rigid body for the quadruped torso. Lower limbs are connected to the the upper limbs via a limited hinge joint with a single rotational degree of freedom. Upper limbs are connected to the torso again with a limited hinge joint with one degree of freedom. The angular limits are shown in a scale diagram of the quadruped in Figure 3.

The assessed fitness of each individual is simply taken as the absolute distance traveled by the quadruped in a fixed time interval. The neural network is updated at half the frequency of the physics simulation, for a total of 5000 updates (approximately 20 seconds of real-time simulation).




Figure 3: Scale drawn diagram showing joint angular limits (in degrees)

Limb	Sensor Nodes
Right Rear	0,1
Left Front	5,6
Left Rear	8,9
Right Front	12,13

Sensor input to the neural network comprises simple contact sensors associated with each lower limb. When a lower limb is in contact with the ground, the sensor value is 1 at

all other times it is zero. Each contact sensor is connected to two network nodes as shown in Table 1.

Motor output nodes are shown in Table 2. The output signal of each motor node is mapped linearly into the hinge angular range. This becomes a target angular displacement. A velocity value for the joint is then calculated based on the difference between the current angular displacement and this target displacement. The physics engine then applies torque necessary to arrive at this joint velocity, constrained by a maximum torque value.

Table 2: Motor Nodes

Joint	Motor Nodes
Hip Right Rear	0
Knee Right Rear	2
Hip Left Front	4
Knee Left Front	6
Hip Left Rear	8
Knee Left Rear	10
Hip Right Front	12
Knee Right Front	14

Results

The results of each evolutionary run are shown in Table 3. The distance traveled is normalized by the body length of the quadruped so as to present the data in a more intuitive fashion. A distance traveled of 1.3 body lengths simply corresponds to the quadruped falling forwards. Between 2 and 3 body lengths, typically one or two steps have been taken. Distances greater than 4 body lengths usually correspond to a slow or unstable gait. Distances greater than this correspond typically to cyclical gaits. The global fitness peak is likely to be around 14 body lengths.

Although the results for GasNets and the hybrid CTRNN/PNN are very similar, there are differences in the stability of evolved gaits, with those of the GasNets exhibiting greater stability. Over all, the fittest individual was evolved using GasNets, however the CTRNN/PNN achieved a marginally higher average fitness measure. The results for the Center-Crossing CTRNN were generally poor, however one of the runs did discover the same gait as the fittest GasNet and CTRNN/PNN. As a consequence the fittest Center Crossing CTRNN individual attained a fitness very close to that of the other two networks considered.

Discussion

Although not obvious from Table 3, the quality of motion exhibited by the quadruped varied substantially with different networks. The fittest CTRNN driven quadruped exhibits motion that is similar to what we might expect from a Central Pattern Generator. The gait is symmetrical and the frequency of oscillation appears to be relatively stable.

Table 3: Distance Traveled by the Fittest Individual (normalized to quadruped body length). Letter superscripts correspond to distinct gaits described in the text.

Run Index	CTRNN	GasNet	CTRNN/PNN
1	1.4	7.2 ^b	4.2
2	1.5	4.6	5.7 ^e
3	1.3	2.0	1.3
4	1.3	3.0	1.6
5	11.7 ^a	2.0	1.3
6	1.3	3.4 ^d	7.2 ^f
7	1.4	2.5	6.2
8	1.4	5.2 ^c	1.5
9	1.4	2.7	1.3
10	1.3	2.0	4.9
11	1.3	1.7	6.2 ^e
12	1.3	2.6	13.6 ^a
13	1.4	13.7 ^a	1.1
14	1.4	3.2	1.3
Average.	2.1	4.0	4.1
Median.	1.4	2.8	2.9
Maximum.	11.7	13.7	13.6

The CTRNN produced a stable gait that continues for a prolonged time period after the end of a trial (if allowed to continue).

In contrast, the gaits generated by the CTRNN/PNN hybrid appear to be highly reactive, with little evidence of forced oscillations (excluding that of the fittest CTRNN/PNN individual). Although the order of stepping may assume a regular pattern, there is considerable variability in the speed of subsequent steps. In this respect motion closely resembles irregular passive dynamic walking. The GasNet demonstrates gaits which exhibit aspects of reactive behavior together with forced oscillations (producing the fasted quadruped within the evaluation time period).

The GasNet and CTRNN/PNN Hybrid exhibited the widest range of gaits. The CTRNN/PNN exhibited some gaits that were not discovered by GasNet, although the CTRNN/PNN gaits were relatively unstable. If we consider the form of the CTRNN/PNN network, it is clear that connection weights will gradually decline if there is a lack of coincident activity. In such a dynamic environment, rich external sensory input may play a more significant role, than it would in networks that exhibit strong intrinsic dynamic activity (such as self-oscillation). This may well go some way to explaining why the gaits exhibited by the CTRNN/PNN appear to be more reactive, but seem to lack strong oscillatory activity.

The evolved quadrupeds exhibit a variety of of the gaits, and body configurations. The quality of the motion varies from driven-oscillatory to ballistic-reactive. Some of the most distinctive patterns are described below.

correspond to that which appears next to the fitness value in Table 3).

- a The front legs hit the ground together, then the back legs, corresponding to the bound gait.
- b The quadruped jumps from its rear limbs, stopping itself with its fore-limbs, before returning to a squatting position. This cycle then repeats. This does not correspond to any of the commonly observed animal gaits.
- c In this case the left fore-limb remains in a forward position, whilst the right fore-limb remains in a rearward position. Rear limbs push off from the ground in a coordinated fashion. Motion resembles that of the three-legged bound gait.
- d In some runs, the quadruped assumes a crawling configuration. Early in the evolutionary run a suspended walking gait is evident. Two diametrically opposed limbs are always in contact with the ground.
- e This motion pattern most closely resembles ballistic walking. It is an highly irregular gait, with little evidence of regular oscillatory movement.
- f This is another bounding gait, however it makes use of “elbows” rather than “hands/feet” in its forelimbs.

Comparison with Prior Study

In the previous study, which considered bipedal locomotion, GasNets appeared to offer the best solution (the only network to achieve cyclical bipedal locomotion) followed by Center-Crossing CTRNNs. For detailed analysis of GasNet dynamics and performance, the reader is referred to (Smith et al., 2003) and (Philippedes et al., 2002). In this respect the results for the GasNet are broadly in line with those of the previous study.

In this study, only one of the CTRNN runs resulted in locomotion. This is in keeping with prior work (Reil and Husbands, 2002) where CTRNNs were evolved for bipedal locomotion control. Only 10% of runs generated oscillatory activity that resulted in bipedal locomotion. No stable cyclical gaits were generated for the biped in the previous comparative study using CTRNNs.

Compared to the previous study, the biggest difference is in the relative performance of the Hybrid CTRNN/PNN. Whilst at best mediocre in the prior study, the results when applied to quadrupedal locomotion are comparable with those of the GasNet. There are two possible explanations that spring to mind. Firstly, modifying the original Hybrid CTRNN/PNN so that it more closely resembles a coupled-oscillator, may result in dynamic activity that is more suitable for self-oscillation and locomotive control. Note that in

the previous study networks were single heterogeneous networks with no axis of symmetry. Secondly, quadrupedal locomotion may be more amenable to reactive solutions than the intrinsically less stable problem of bipedal locomotion.

CPGs are currently the dominant motor control paradigm. Work has shown that it is possible to model all the common quadrupedal gaits using a network of eight cells (Buono and Golubitsky, 2001). The results described here may lead us to question whether or not reactive responses are just as important in generating locomotive activity.

Conclusion

In conjunction with previous studies, GasNets appear to provide us with a reliable approach to evolving locomotion in physically simulated agents. However ANNs that exhibit substantially different dynamics (such as CTRNN/PNNs), may yield alternative solutions that are comparable in certain cases. Future work will investigate issues of gait stability and directed motion.

References

- Beer, R. (1995). On the Dynamics of Small Continuous-Time Recurrent Neural Networks. *Adaptive Behavior*, 3(4):469–509.
- Billard, A. and Ijspeert, A. K. (2000). Biologically Inspired Neural Controllers for Motor Control in a Quadruped Robot. In Amari, S., Giles, C., Gori, M., and Piuri, V., editors, *IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN 2000)*, volume 6, pages 6637–6641.
- Blynel, J. and Floreano, D. (2002). Levels of Dynamic and Adaptive Behavior in Evolutionary Neural Controllers. In Hallam, B., Floreano, D., Hallam, J., Hayes, G., and Meyer, J.-A., editors, *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior (SAB 02)*, pages 272–281. MIT Press.
- Bongard, J. and Paul, C. (2001). Making Evolution an Offer it Can't Refuse: Morphology and the Extradimensional Bypass. In Keleman, J. and Sosik, P., editors, *Proceedings of the Sixth European Conference on Artificial Life (ECAL 2000)*, pages 401–412.
- Buono, P.-L. and Golubitsky, M. (2001). Models of central pattern generators for quadruped locomotion. *Journal of Mathematical Biology*, 42:291–326.
- Floreano, D. and Mondada, F. (1996). Evolution of Plastic Neurocontrollers for Situated Agents. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S., editors, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB 96)*. MIT Press.
- Husbands, P., Philippides, A., Smith, T., and O'Shea, M. (2001). Volume Signalling in Real and Robot Nervous Systems. *Theory in Biosciences*, 120:253–269.
- Husbands, P., Smith, T., O'Shea, M., Jakobi, M., Anderson, J., and Philippides, A. (1998). Brains Gases and Robots. In et al, N. N., editor, *Proceedings of the 8th International Conference on Artificial Neural Networks (ICANN98)*, pages 51–64. Springer-Verlag.
- Ishiguro, A., Fujii, A., and Eggenberger Hotz, P. (2003). Neuromodulated Control of Bipedal Locomotion Using a Polymorphic CPG Circuit. *Adaptive Behavior*, 11:7–17.
- Mathayomchan, B. and Beer, R. (2002). Center-Crossing Recurrent Neural Networks for the Evolution of Rhythmic Behavior. *Neural Computation*, 14:2043–2051.
- McHale, G. and Husbands, P. (2004). GasNets and other Evolvable Neural Networks applied to Bipedal Locomotion. From Animals to Animats 8: Proceedings of the Eighth International Conference on Simulation of Adaptive Behavior (SAB 2004) - Accepted.
- Philippides, A., Husbands, P., Smith, T., and O'Shea, M. (2002). Fast and Loose : Biologically Inspired Couplings. In Standish, R., Bedau, M., and Abbass, H. A., editors, *Artificial Life VIII: Proceedings of the Eight International Conference on Artificial Life*, pages 292–301. MIT Press.
- 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., Philippides, A., and O'Shea, M. (2003). Temporally Adaptive Networks: 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.
- Urzelai, J. and Floreano, D. (2000). Evolutionary Robots with Fast Adaptive Behavior in New Environments. In Thompson, A., editor, *Third International Conference on Evolvable Systems: From Biology to Hardware (ICES2000)*. Springer Verlag.
- Yamauchi, B. and Beer, R. (1994). Integrating Reactive, Sequential and Learning Behavior Using Dynamical Neural Networks. In Cliff, D., Husbands, P., Meyer, J., and Wilson, S., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB 94)*, pages 205–224. MIT Press.