GasNets and other Evolvable Neural Networks applied to Bipedal Locomotion

Gary McHale* **Phil Husbands****
*Department of Informatics,
**Centre for Computational Neuroscience and Robotics
Centre for Research in Cognitive Science
University of Sussex, Falmer, Brighton BN1 9QH, UK
cglife@hotmail.com

Abstract

Evolutionary robotics relies upon techniques involving the evolution of artificial neural networks to synthesize sensorimotor control systems for actual or physically simulated robots. This paper is a comparative study of three principal types of artificial neural networks; the Continuous Time Recurrent Neural Network (CTRNN), the Plastic Neural Network (PNN) and the GasNet. An attempt is made to evolve networks capable of achieving locomotion with a physically simulated biped. Of the 14 distinct networks tested, GasNets were the only network to achieve cyclical locomotion, although CTRNNs were able to attain a higher level of average fitness.

1. Introduction

Evolutionary Robotics seeks to create real and simulated robotics that have economic value, and to gain insight into biological systems through suitable abstractions of these systems.

In Evolutionary Robotics, biologically inspired Artificial Neural Networks (ANN) are evolved using Genetic or Evolutionary Algorithms (GAs or EAs) to generate sensorimotor control systems. A significant challenge faced in this area is that the evolved neural network circuitry is rarely tractable to functional decomposition or deconstruction other than in the most trivial of cases (Beer et al., 1999). Whilst the characteristic equations associated with a specific network are a compact description of the network, we are as yet unable to predict from these equations, the dynamic characteristics of the network when it is embodied in an environmental agent.

In this paper empirical investigation is used to compare the performance of different network formulations in a common task; that of evolving bipedal locomotion in a physically simulated biped. It is hoped that through considering a number of networks applied to different but related tasks, we may discover certain network formulations that are of general or specific use. It is intended as a complementary study to work that sought to evolve locomotion in a physically simulated quadruped (McHale and Husbands, 2004). Whilst an in-depth analysis of the reasons why one network is superior to another is beyond the scope of this paper, a summary of relevant work is given in the latter stages of this report. A total of 14 different ANNs are assessed in this study.

2. Related Work

Continuous Time Recurrent Neural Network (CTRNNs), Plastic Neural Networks (PNN) and GasNets are all forms of Dynamic Recurrent Neural Networks (DRNNs). CTRNNs are the one of the simplest forms of DRNN. PNNs (Floreano and Mondada, 1996) seek to incorporate aspects of run-time learning through Hebbian adaptive network weights. GasNets are inspired by the action of Nitric Oxide as a neuromodulator (Husbands, 1998). GasNets and CTRNNs particularly have the potential to generate intricate dynamics on several different time scales, albeit by rather different means. This very likely, at least partially, accounts for their observed advantages over other less dynamic and unmodulated ANNs (Husbands et al., 1998, Smith et al., 2002).

This paper is the first study that includes GasNets, CTRNNs, PNNs and other forms of DRNN, in a performance comparison involving the generation of a suitably dynamic behavior - bipedal locomotion. This report extends work using a conventional DRNN applied to bipedal locomotion (Reil and Husbands, 2002). Other researchers have also addressed the issue of bipedal locomotion using evolutionary algorithms (Hase and Yamazaki, 1999) that incorporate neuromodulation (Ishiguro et al., 2003), variable morphology (Bongard and Paul, 2001) and genetic programming (Ok et al., 2001).

3. Experimental Setup

Parameters were chosen as outlined by the original authors wherever possible (given the differences in genetic encoding). These original papers should be consulted for further experimental details if more information is required. Variations from their implementations are stated where relevant.

For the purposes of a fair comparison no assumption is made about the coupling of the underlying network nodes. The chosen task is to achieve bipedal locomotion *ab initio*, without predetermining that nodes are configured as coupled oscillators (Matsuoka, 1985).

3.1 Genetic Algorithm

The same distributed steady-state genetic algorithm is used on all networks. The population grid has dimensions of 10 by 10, to yield a total of 100 individuals. Competition is tournament based, with a tournament group comprising three individuals. A *principal* is selected at random from the grid. Two other population members are selected by a random walk originating at the principal's grid cell. The length of the random walk is an integer value in the range 1-4.

If the principal is the fittest then the weakest member is replaced by a mutated version of the principal. Otherwise the weakest member of the tournament set is replaced by a recombination of the two fittest individuals (using single point cross-over). This recombined genome is then mutated.

In these experiments the generation index is incremented after the evaluation of 100 individuals (comprising a pseudo-generation). Each neural network type was evaluated for 200 generations. This was carried out ten times with different random seeds for each network type.

Each fitness assessment starts with the biped in a stable standing position. The trial lasts a simulated 20 seconds. The fitness of the individual is taken to be the minimum of the distance traveled by either of the biped's feet, or hips. Each time-step for the neural network update lasts 0.025 seconds (or 1/40th of a second). Thus a 20 second trial corresponds to 800 neural network time-steps.

The trial was terminated if the biped fell below 50 percent of its original height and a fitness value assigned based on the distance traveled up to that point in time.

3.2 Physical Model

The computer code for the biped physical model was generated with the aid of a product called Autosim. Joints are simulated as torsional springs. Strictly speaking the motor output is actually a control signal. This signal is mapped to an angular displacement that corresponds to the rest position of a torsional spring. A change in the angular displacement of this rest position, will result in a torque applied to the lower limb attached to the joint (as the spring seeks to restore the joint to its new rest position).

The kinematic root has five degrees of freedom (two rotational and three translational). The biped is physically incapable of rotating in its roll axis. This prevents the biped from falling over on its side, although it is still free to fall forwards and backwards. Whilst this is not entirely physically realistic, it is sufficient for the purposes of this comparative study. This still remains a non-trivial problem; feet are modeled as point contact points, resulting in a dynamically unstable model after an initial displacement of the biped.

3.3 Sensorimotor System

Motor output signals are taken from the 1st 5th, 9th, 13th nodes in the network . These nodes were used to control the rest position of the torsional springs in the biped's left hip, right hip, left knee and right knee respectively. The output of these nodes was mapped to a range of -94 to 101 degrees angular hip displacement, and 0 to -89 degree angular knee displacement.

There are two sets of sensor configurations used in this experiment (nominally referred to as R and I configurations). The majority of simulations were carried out with a regular connection pattern (R configuration). Sensor input from the right foot contact sensor occurs at the 2nd, 6th, 10th and 14th nodes. Sensor input from the left foot contact sensor occurs at the 3rd, 7th, 11th and 15th nodes. Sensor input consisted of a binary 1.0 or 0.0 value depending upon whether or not the feet of the biped were in contact with the ground.

Additional simulations were carried out on GasNet networks with an irregular sensor connection pattern (I configuration). In this case sensor signals were input to the 8th, 9th and 10th nodes for the right foot, and the 11th, 12th, 13th and 14th nodes for the left foot. One of the reasons for this second sensor configuration was to see if there were any significant changes in the performance of circuits where sensor nodes were connected directly to motor nodes (i.e. in the case of the 9th and 13th nodes in the I configuration).

3.4 Genetic Encoding and Mutation

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. Connection weights (where relevant) are also stored on a per node basis. The number of genetic parameters per node for each network type is shown in Table 1. Note that CTRNN's require us to store node weights on a per connection basis, in a fully connected network this means that we have N weights per node as evolvable parameters, where N is the total number of nodes in the network, plus the time constant and bias parameters, to yield a total of 18 parameters in this case.

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.

In addition to this, those networks where connection weights are under evolutionary control (such as conventional 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 of values. An exponent f was randomly selected from the set:

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

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

$$\tau = 1.0 + r(10^f) \tag{1}$$

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

4. Network Details and Characteristic Equations

This section describes the network details of each of the varieties of networks tested. The focus is on the characteristic equations that govern the dynamics of each network variety.

A list of the network types is shown in Table 1. Each network type is assigned an index for reference purposes numbered between 1 and 14. The R letter denotes the regular sensor configuration, the I denotes the irregular sensor configuration.

It should be noted that Conventional GasNets comprise nodes that are spatially distributed. A parametric coding strategy is used where connections are determined for each node based on genetic parameters that define geometric arcs originating at each node. A node that falls within an excitatory or inhibitory arc are deemed to be electrically connected to the node at origin of the arc (Husbands, 1998). The consequence of this, is that conventional GasNets are only sparsely connected. This approach is used in the the following network types; 6R, 8R, 9R, 13I and 14I. All other network types are fully interconnected.

Table 1: List of network types, and number of genetic parameters per node.

Type	Params.
1R Conventional CTRNN	18
2R Center-Crossing CTRNN	18
3R Basic PNN	3
4R Gas-Modulated PNN	17
5R CTRNN/PNN Hybrid	20
6R Conventional GasNet	15
7R Fully Recurrent GasNet	15
8R CTRNN/GasNet Hybrid no Gas	5
9R CTRNN/GasNet Hybrid with Gas	18
10I Conventional GasNet 16 Cell	15
11I Conventional GasNet 32 Cell	15
12I Fully Recurrent GasNet	15
13I CTRNN/GasNet Hybrid no Gas	5
14I CTRNN/GasNet Hybrid with Gas	18

4.1 Conventional CTRNN - type 1R

This is a conventional CTRNN (Beer, 1995), fully recurrent, where node connection weights and biases are under evolutionary control.

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i}(-y_i^t + \sum_{j=1}^{N} \omega_{ji} \sigma(y_j^t + \theta_j) + I_i) \tag{2}$$

$$i = 1, 2, ..., N$$

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 equation 1.

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 $\theta \epsilon [-2, 2]$.

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 $\omega \epsilon [-4.0, 4.0]$.

σ is the logistic activation function.

$$\sigma(z) = \frac{1}{(1+e^{-z})} \tag{3}$$

4.2 Center-Crossing CTRNN - type 2R

The characteristic equation of the Center-Crossing CTRNN (Mathayomchan and Beer, 2002) is the same as that of the CTRNN (type 1). However initial biases are calculated such that:

$$\theta_i = \frac{-\sum_{j=1}^{N} \omega_{ji}}{2} \tag{4}$$

These authors 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.

4.3 Basic PNN - type 3R

The key characteristic of PNN's (Urzelai and Floreano, 2000) is that connection weights vary over time based on Hebbian learning rules given by:

$$\omega_{ji}^{t} = \omega_{ji}^{t-1} + \eta \Delta\omega_{ji} \tag{5}$$

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

$$\Delta\omega_{ji} = (1-\omega_{ji})x_j y_i \tag{6}$$

Post-Synamptic Rule

$$\Delta\omega_{ji} = \omega_{ji}(-1+x_j)y_i + (1-\omega_{ij})x_j y_i \tag{7}$$

Pre-Synaptic rule Rule

$$\Delta\omega_{ji} = \omega_{ji}(-1+y_i) + (1-\omega_{ji})x_j y_i \tag{8}$$

Covariance Rule

$$\Delta\omega_{ji} = \begin{cases} (1-\omega_{ji}) & \text{if } F(x_j, y_i) > 0 \\ (\omega_{ji})F(x_j, y_i) & \text{otherwise} \end{cases} \tag{9}$$

Where:

$$F(x_j, y_i) = \tanh(4(1-|x_i - y_j| - 2) \tag{10}$$

All nodes in the PNN are fully interconnected (with self-connections also supported). The rate of learning η can only assume one of four values (0.0, 0.3, 0.6, 0.9). The characteristic equation for the PNN is shown below:

$$y_i^{t+1} = \sigma(\sum_{j=1}^{N} \omega_{ji}^{t}(y_j^{t})) + I_i \quad i = 1, 2, ..., N \tag{11}$$

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 term "basic" is used to differentiate it from that used by (Tuci and Quinn, 2003) where an additional bias input was used for each node. In common with the implementation described in (Blynel and Floreano, 2002) the range of y_i is $[0, 2]$ for input neurons and $[0, 1]$ for hidden and output neurons.

4.4 Gas-Modulated PNN - type 4R

This is essentially the same as the basic PNN (type 3) with the exception that nodes whose weights are genetically determined to be modified by the Plain Hebb Rule, or Post-Synaptic Rule have their weights modified by diffused gases. For these two varieties of nodes, the weight modification rule becomes:

Gas Modified Plain Hebb Rule

$$\Delta\omega_{ji} = (\frac{c_{1i}^{t}}{c_{1i}^{t} + c_{2i}^{t}})(1-\omega_{ji})x_j y_i \tag{12}$$

Gas Modified Post-Synaptic Rule

$$\Delta\omega_{ji} = (\frac{c_{2i}^{t}}{c_{1i}^{t} + c_{2i}^{t}})(\omega_{ji}(-1+x_j)y_i + (1-\omega_{ij})x_j y_i) \tag{13}$$

Where:

c_{1i}^{t} is the concentration of gas 1 at the i'th node.

c_{2i}^{t} is the concentration of gas 2 at the i'th node.

The significance of these equations is that the rate of change in the weights of the inputs to these nodes will vary continuously with changes in the relative concentration of these two gases. When both gases have zero concentrations there is no change in weight.

4.5 CTRNN/PNN Hybrid - type 5R

This is a modification of the conventional PNN, with activation signals modified by a node based time constant under evolutionary control (in a similar fashion to conventional CTRNNs).

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i}(-y_i^t + \sum_{j=1}^{N} \omega_{ji}^t \sigma(y_j^t + \theta_j) + I_i) \quad (14)$$

$$i = 1, 2, ..., N$$

Where:

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

4.6 Standard GasNet - type 6R

In GasNets (Husbands et al., 1998), 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.

$$y_i^{t+1} = \tanh[k_i^t(\sum_{j \epsilon C_i} \omega_{ji}\sigma(y_j^t + I_i)) + b_i] \quad (15)$$

Where:

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

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.

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 16). 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 17 and 18 at a rate determined by s.

$$C(d,t) = \begin{cases} e^{-2d/r} \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (16)$$

$$T(t) = \begin{cases} H\left(\frac{t-t_e}{s}\right) & \text{emitting} \\ H\left[H\left(\frac{t_s-t_e}{s}\right) - H\left(\frac{t-t_s}{s}\right)\right] & \text{not emitting} \end{cases} \quad (17)$$

$$H(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x < 1 \\ 1 & \text{else} \end{cases} \quad (18)$$

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 19) 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 19, with transfer parameters updated on every time step as the network runs.

$$k_i^t = k_i^0 + \alpha C_1^t - \beta C_2^t \quad (19)$$

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.

4.7 Fully Recurrent GasNet - type 7R

The fully recurrent GasNet uses the GasNet model of gas diffusion but adopts a fully recurrent connectivity model.

Whereas a conventional GasNet is sparsely connected, this network is fully connected.

$$y_i^{t+1} = \tanh[k_i^t(\sum_{j=1}^{N} \omega_{ji}\sigma(y_j^t + I_i)) + b_i] \qquad (20)$$

$$i = 1, 2, ..., N$$

4.8 CTRNN/GasNet Hybrid no Gas - type 8R

This is a variation of the conventional CTRNN. In this case inter-node connectivity is determined by the approach used in GasNets. It uses a sigmoid transfer function, with bias and node time constants under evolutionary control.

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i}(-y_i^t + \tanh[K_i(\sum_{j \epsilon C_i} \omega_{ji}\sigma(y_j^t + I_i)) + b_i]) \qquad (21)$$

Where:

K_i is a transfer function constant.

T is the time slice constant.

τ_i is the time constant for the i'th node.

4.9 CTRNN/GasNet Hybrid with Gas - type 9R

This is another variation of a conventional CTRNN, but in this case it is more extensively modified along the lines of GasNets. The network nodes transfer function is gas modulated, and network connectivity is based on the GasNet model. What remains of the the original CTRNN is the node time constant.

$$y_i^{t+1} = y_i^t + \frac{T}{\tau_i}(-y_i^t + \tanh[k_i^t(\sum_{j \epsilon C_i} \omega_{ji}\sigma(y_j^t + I_i)) + b_i]) \qquad (22)$$

Where:

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

4.10 Conventional GasNet 16 Cell - type 10I

This network is based on the Conventional GasNet (type 6R) but uses the irregular sensor configuration. Another difference is that bias values are set such that $b_i \epsilon [-4, 4]$.

4.11 Conventional GasNet 32 Cell - type 11I

Other than the number of cells in this network, all other parameters are the same as those of the 16 cell gas net (type 10I). Note that of all the networks tested, this is the only one that is comprised of 32 cells or nodes.

Given the nature of the the GasNet connectivity algorithm 32 cell GasNets are likely to be more highly interconnected than 16 cell GasNets. The size of the plane remains constant, so a higher number of cells means a higher cell density. A larger number of cells will tend to fall within a given connection arc, thus resulting in a higher number of inter-cell connections per cell.

4.12 Fully Recurrent GasNet - type 12I

This network is based on the Fully Recurrent GasNet (type 7R) but uses the irregular sensor configuration. Another difference is that bias values are set such that $b_i \epsilon [-4, 4]$.

4.13 CTRNN/GasNet Hybrid no Gas - type 13I

This network is based on the CTRNN/GasNet Hybrid without Gas (type 9R) but uses the irregular sensor configuration. Another difference is that bias values are set such that $b_i \epsilon [-4, 4]$.

4.14 CTRNN/GasNet Hybrid with Gas - type 14I

This network based on the CTRNN/GasNet Hybrid with Gas (type 9R) but uses the irregular sensor configuration. Another difference is that bias values are set such that $b_i \epsilon [-4, 4]$.

5. Results

Table 2 shows the peak fitness of the best individual in the ten evolutionary runs for each network type, together with the standard deviation of the fitness of the fittest individual across all ten runs. Table 3 shows the average and median peak fitness value across all runs for each network type.

To put these fitness values in perspective 2.5 meters is attainable through a fast walk in about 10 seconds (attained by the Conventional GasNet type 10I), and a slow walk in about 20 seconds (attained by the Conventional GasNet type 6R). Bipeds that have traveled around 1 meter have typically taken around two steps. The worst performing Basic PNN typically only extended one leg, before slightly drawing the lagging leg forward.

Figure 1 shows three snapshots of the biped walking. This rendering is based on the results obtained from the 16 cell GasNet (type 10I).

5.1 Principal Results Summary

We can summarize the most important results as follows:

1. Conventional GasNets were the only networks to achieve cyclical bipedal locomotion.
2. Center-Crossing CTRNNs achieved the second highest peak fitness after the conventional GasNets.
3. Center-Crossing CTRNNs and CTRNNs attained the highest average fitness values.
4. Center-Crossing CTRNNs marginally outperformed conventional CTRNNs in peak and average fitness.
5. Basic PNNs without dynamic attributes performed the worst.
6. Both a CTRNN/PNN Hybrid and a Gas Modulated PNN improved on the performance of the Basic PNN.
7. The CTRNN/PNN Hybrid achieved comparable results to the Conventional CTRNN.
8. Fully Connected GasNets perform badly compared to Conventional GasNets. GasNet performance seems to decline with increased inter-connectivity (see type 11I vs 10I and 6R).
9. Conventional GasNets exhibit a higher variation in phenotype fitness than other network types. Conventional CTRNNs exhibit relatively low phenotypic variation.
10. Gas Modulated networks generally outperform their un-modulated counter-parts (see type R8 vs R9 and I13 vs I14).
11. GasNet performance is largely unaltered by minor sensorimotor configuration changes (see R vs I sensor configurations).

These results are generally consistent with those reported by GasNet researchers and the comparative studies cited earlier.

6. Discussion

An obvious question to ask when considering the relative performance of these networks is whether or not the relative performance differences are primarily a consequence of the number of degrees of freedom in parameter space. Table 1 lists the number of parameters per node. The basic PNN (3R) has the lowest number of parameters and displays the worst performance. However, the CTRNN/PNN Hybrid (5R) has the largest number of evolvable parameters, but it does not have the best performance. Whilst degrees of freedom in parameter space may be a contributing factor in the results described here, the internal dynamics of each network also play a significant role. The 16 cell GasNet (10I) has substantially superior peak fitness performance to the 32 cell GasNet (11I) and fully connected GasNet (12I),

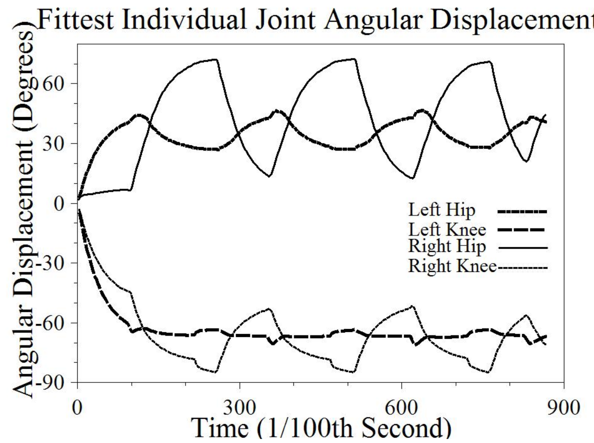

Figure 1: GasNet Biped Kinematic Data (Upper; Left Hip and Right Hip, Lower Right Knee and Left Knee).

Table 2: Peak fitness across all evolutionary runs, and Standard Deviation of the peak fitness (meters).

Network	Peak	SD.
6R Conventional GasNet	2.63	0.70
10I Conventional GasNet 16 Cell	2.62	0.66
2R Center-Crossing CTRNN	1.63	0.24
11I Conventional GasNet 32 Cell	1.31	0.29
5R CTRNN/PNN Hybrid	1.27	0.33
14I CTRNN/GasNet Hybrid with Gas	1.25	0.21
1R Conventional CTRNN	1.24	0.15
9R CTRNN/GasNet Hybrid with Gas	1.22	0.25
8R CTRNN/GasNet Hybrid no Gas	1.18	0.23
4R Gas-Modulated PNN	1.06	0.30
7R Fully Recurrent GasNet	0.98	0.09
13I CTRNN/GasNet Hybrid no Gas	0.72	0.13
12I Fully Recurrent GasNet	0.53	0.21
3R Basic PNN	0.24	0.01

Table 3: Average and median of distances traveled by the fittest individual at the end of each run (meters).

Network	Avg.	Med.
2R Center-Crossing CTRNN	1.13	1.19
1R Conventional CTRNN	1.11	1.18
6R Conventional GasNet	0.97	0.82
9R CTRNN/GasNet Hybrid with Gas	0.95	0.96
11I Conventional GasNet 32 Cell	0.94	1.08
14I CTRNN/GasNet Hybrid with Gas	0.92	0.96
10I Conventional GasNet 16 Cell	0.92	0.93
5R CTRNN/PNN Hybrid	0.90	1.00
8R CTRNN/GasNet Hybrid no Gas	0.90	0.95
13I CTRNN/GasNet Hybrid no Gas	0.61	0.66
4R Gas-Modulated PNN	0.57	0.72
12I Fully Recurrent GasNet	0.47	0.36
7R Fully Recurrent GasNet	0.35	0.38
3R Basic PNN	0.23	0.23

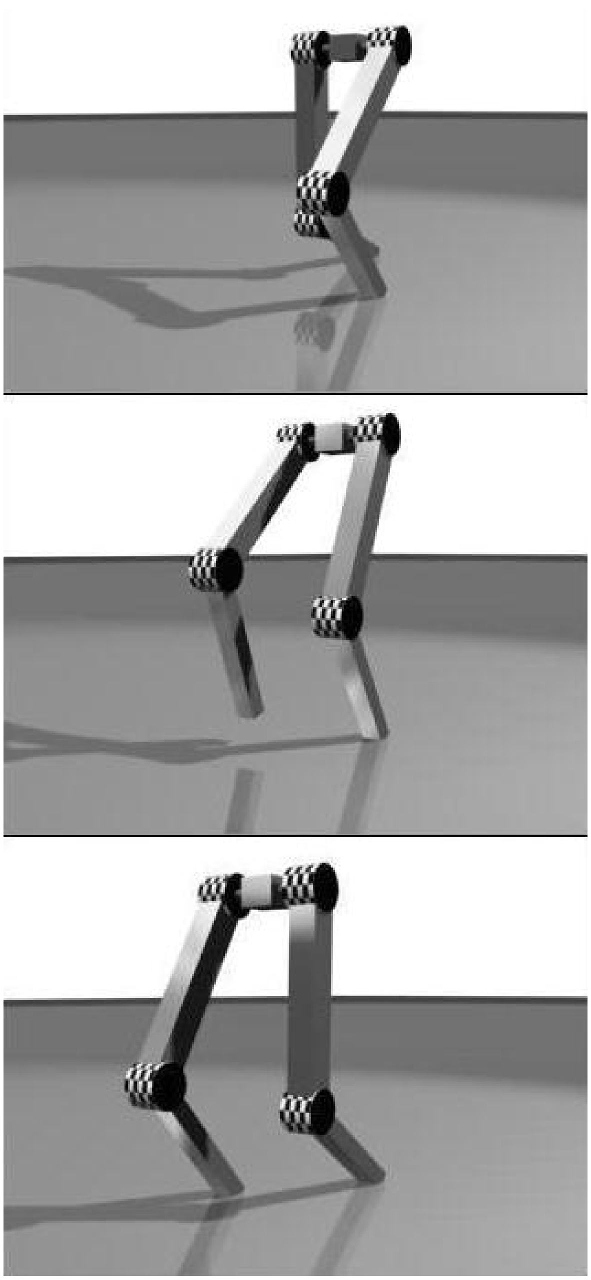

Figure 2: Renderings of the fittest GasNet biped at 5.0, 5.875 and 6.25 seconds.

despite having the same number of evolvable parameters. The number of parameters for a given network is a consequence of the network formulation, and whilst the low number of parameters used in the basic PNN may lead us to believe that it is under-specified in comparison with other networks, the key determinant is the saliency of a given parameter to the exploration of phenotype space. This is not something that can be easily determined without referring to the underlying dynamics of the alternative network formulations. As a consequence, the discussion in this section focusses largely on the underlying dynamics, rather than specifically address the issue of differences in paramater space.

6.1 Temporal Adaptivity and Evolvability

It is claimed that GasNets have a high evolvability due to their high capacity for *temporal adaptation* (Smith et al., 2002). A significant factor determining the activity of the network is the coupling of electrical and gas dynamics (Philippedes et al., 2002). In this work the authors showed that reducing the likelihood that a node pair are both chemically and electrically coupled, can improve the evolvability of the GasNet network even further. In particular they state;

> .. systems involving distinct yet coupled processes are highly evolvable when there is a bias towards loose coupling between the processes; this allows the possibility of 'tuning' one against the other without destructive interference.

The relatively poor performance of fully connected GasNets, where a high degree of coupling is forced into the network, tends to support this claim. Philippides et al. also refer to the multiple redundancies inherent in loosely coupled GasNets which potentially lead to increased numbers of routes through the evolutionary search space. These factors may help to account for the relatively high variance in GasNet fitness in comparison with CTRNNs, as well as the significantly higher peak fitness. The inherent dynamics are being shaped and explored in a very different way.

In the case of the one of the fittest GasNet bipeds (type I10) the reactive response of the right hip to the transition from non-contact to contact of the left foot appears to be primarily gas-mediated. In contrast the joint angular displacement dynamics are governed primarily by electrical activation signals. Whilst this may be an over-simplification (GasNets are integrated systems with co-dependencies between both gas and electrical signalling), it is easy to imagine circumstances in which there are independent phenotypic processes that have intrinsic time dynamics associated with them. The ability to explore these phenotypic temporal dynamics in parallel may be a significant factor in the evolvability of GasNets.

Concepts such as of temporal adaptivity and system coupling may be useful concepts in helping us to understanding the dynamics of such systems. However we still need to try and identify the specific characteristics of these networks that support temporal adaptivity. Two of the most obvious areas to consider are the frequency and phase characteristics of the GasNet.

6.2 Phase Space Exploration

Phase relationships between signals in an articulated agent are extremely important to achieve coordinated activity. One aspect of GasNets is that nodes are physically distributed in a virtual 2-dimensional space. A variation in the distance between nodes (through mutation for example) results in a phase lag or lead in the modulation of gas-coupled nodes. In this sense phase relationships are under direct evolutionary control. None of the other networks investigated in this study embody these characteristics. Simply put, in GasNets, node position mutations are operating directly in phenotype phase space (where nodes are chemically coupled).

In direct encoding, if a single parameter maps onto a specific phenotypic attribute that is largely independent of other attributes, this is likely to aid the efficient exploration of phenotypic space via mutation operators, since it may result in a smoother fitness landscape. In the case of articulated bipedal locomotion, we might imagine that the phase relationship between hip joints is a significant phenotypic attribute that affects overall fitness. A network model that can explore this phase relationship through mutation operators may exhibit greater temporal adaptation than one that cannot.

6.3 Frequency Space Exploration

Examination of the motor output signals of GasNets compared to other networks indicates that GasNets generate a motor control signals with a wide range of frequency components (particularly high frequency). It is easy to imagine that a network that exhibits signification oscillatory behavior over a wide bandwidths may have some advantages over networks in seeking to discover a solution suitable for articulated motion.

6.4 Temporal Dynamics and PNNs

It should be stated that there is no reason to expect that the Basic PNN should have performed particularly well in this problem domain. The environment remains constant over evolutionary time. As such there is no additional benefit to be gained by in-trial learning. Limited sensor stimulation may have resulted in a rapid decay in node activity in this implementation. However, the modification of the Basic PNN to incorporate richer time dynamics (e.g. type 4R and 5R networks) can improve the performance of PNN networks.

6.5 Complementary Study

A complementary study carried out by the authors of this paper (McHale and Husbands, 2004) investigated three of the network varieties described here (Conventional GasNet, Center-Crossing CTRNN and CTRNN/PNN Hybrid) in the evolution of locomotion in a physically simulated quadruped. Network morphologies were constrained to resemble coupled-oscillators. The initial population was comprised solely of symmetrical networks, with corresponding nodes connected by mutually inhibitory connections. The Open Dynamics Engine (ODE) physics simulation package was used, without any restrictions on the degrees of freedom of the quadrupeds kinematic root.

Whilst the GasNet network attained the highest fitness, the CTRNN/PNN Hybrid (introduced in this study) achieved comparable results, discovering several gaits between them. Whilst the Center-Crossing CTRNN also attained a high fitness level, it discovered only one gait, and average fitness was approximately half that of the other two networks. Qualitatively the gaits discovered by the CTRNN/PNN Hybrid appeared to more reactive than those discovered by the GasNet, exhibiting some gaits that resembled ballistic walking. The evolved GasNets discovered more stable gaits (that continued if the trial period was extended), and generally appeared to rely more upon forced oscillations than the CTRNN/PNN Hybrid.

7. Acknowledgements

The original physical biped model was developed in collaboration with C.Buckley. Thanks should be given to the two anonymous referees for their constructive criticism and recommendations.

8. Conclusion

The results described here, taken in conjunction with those of complementary studies, suggest that GasNets offer a reliable solution in solving sensorimotor control problems for physically simulated agents. This paper also introduces new variants of conventional networks that may merit further study, such as the gas-modulated Hebbian network (4R) and CTRNN/PNN Hybrid (5R). The results of the complementary study described earlier remind us that we should be careful not to generalize from single experiments. The consideration of alternative dynamic models should be an ongoing part of the characterization of alternative network formulations. Future work will seek to carry out additional comparative studies for physically simulated hexapods and bipeds with unconstrained kinematic roots in a continuation of the work described here.

References

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

Beer, R., Chiel, H., and J.C.Gallagher (1999). Evolution of CPGs for walking II. General principles and individual variability. *Journal of Computational Neuroscience*, 7(2):119–147.

Blynel, J. and Floreano, D. (2002). Levels of Dynamic and Adaptive Behaviorin Evolutionary Neural Controllers. In Hallam, B., Floreano, D., Hallam, J., Hayes, G., and Meyer, J.-A., (Eds.), *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., (Eds.), *Proceedings of the Sixth European Conference on Artificial Life (ECAL 2000)*, pages 401–412.

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., (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB 96)*. MIT Press.

Hase, K. and Yamazaki, N. (1999). Computational evolution of human bipedal walking by a neuro-musculo-skeletal model. *Artificial Life and Robotics*, 3:133–138.

Husbands, P. (1998). Evolving Robot Behaviours with Diffusing Gas Networks. In Husbands, P. and Meyer, J.-A., (Eds.), *EvoRobot98: Proceedings of 1st European Workshop on Evolutionary Robotics*, pages 71–86. Springer-Verlag LNCS.

Husbands, P., Smith, T., O'Shea, M., Jakobi, M., Anderson, J., and Phillipedes, A. (1998). Brains Gases and Robots. In Niklasson, N., (Ed.), *Proceedings of the 8th International Conference on Artificial Neural Networks (ICANN98)*, pages 51–64. Springer-Verlag.

Ishiguro, A., Fujii, A., and Hotz, P. E. (2003). Neuromodulated Control of Bipedal LocomotionUsing a PolymorphicCPG Circuit. *Adaptive Behavior*, 11:7–17.

Mathayomchan, B. and Beer, R. (2002). Center-Crossing Recurrent Neural Networks for the Evolutionof Rhythmic Behavior. *Neural Computation*, 14:2043–2051.

Matsuoka, K. (1985). Sustained-Oscillations Generated by Mutually Inhibiting Neurons with Adaptation. *Biologocal Cybernetics*, 52:367–376.

McHale, G. and Husbands, P. (2004). Quadrupedal Locomotion : GasNets, CTRNNs and HybridCTRNN/PNNsCompared. Artificial Life 9: Proceedings of the Ninth International Conference on Artificial Life- Submitted.

Ok, S., Miyashita, K., and Hase, K. (2001). Evolving bipedal locomotion with genetic programming - preliminary report. In *Proceedings of the 2001 Congress on Evolutionary Computation (CEC2001)*, pages 1025–1032. IEEE Press.

Philippedes, A., Husbands, P., Smith, T., and O'Shea, M. (2002). Fast and Loose : Biologically Inpsired Couplings. In Standish, R., Bedau, M., and Abbass, H. A., (Eds.), *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. (2002). M. neuronal plasticity and temporal adaptivity: Gasnet robot control networks. *Adaptive Behavior*, 10(3/4):161–184.

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., (Eds.), *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 et al., A. T., (Ed.), *Third International Conference on EvolvableSystems: From Biology to Hardware (ICES2000)*. Springer Verlag.