Husbands, P., Smith, T.M.C., O'Shea, M., Jakobi, N., Anderson, J. and Philippides, A. **Brains, Gases and Robots**. In Niklasson, N. *et al.* (editors). *Proceedings of the 8th International Conference on Artificial Neural Networks: ICANN98*. pp51-64. Springer. 1998.

Brains, Gases and Robots

Phil Husbands[1], Tom Smith[1], Michael O'Shea[1]
Nick Jakobi[1,2], John Anderson[1] and Andy Philippides[1]
[1]Centre for Computational Neuroscience and Robotics
University of Sussex
Brighton, UK
[2]AnimatLab, ENS
Paris, France

15th July 1998

1 Introduction

Over the past decade there has been renewed interest within AI in building simple autonomous 'creatures' as a way of investigating mechanisms underlying the generation of adaptive behaviour [4, 1]. The vast majority of researchers in this field use some form of artificial neural network (ANN) as the basis of the 'nervous system' of their agents. These networks can be envisaged as simple nodes connected together by directional wires along which signals flow. As has been pointed out by various people (e.g. [3]), advances in neuroscience have made it clear that the propagation of action potentials, and the changing of synaptic connection strengths, is only a very small part of the story of the brain (e.g [17]). This in turn means that connectionist style networks, and even recurrent dynamical ones, are generally very different kinds of systems from those that generate sophisticated adaptive behaviours in animals. Although our picture of biological neuronal networks changes every few years, contemporary neuroscience can provide a rich source of inspiration in devising alternative styles of artificial network [2].

In the last few years it has become clear that freely diffusing Nitric Oxide (NO) acts as a neurotransmitter and is involved in a range of modulatory processes. NO can act in space and time over volumes containing many synapses and nerve cells [8]. This is very different from the action of classical neurotransmitters which signal at precise synaptic sites. We have developed a number of ANNs based on abstractions of these phenomena and have used them to build control systems for autonomous mobile robots. Nodes in a spatially distributed network can emit 'gases' which diffuse through the network. The 'gases' can modulate intrinsic properties of nodes and connections in a concentration dependent fashion [11]. This paper concentrates on some of this work.

One of the new styles of AI is Evolutionary Robotics [5, 15]. The evolutionary process, based on a genetic algorithm [10], involves evaluating, over many generations, whole populations of control systems specified by artificial genotypes. These are interbred using a Darwinian scheme in which the fittest individuals are most likely to produce offspring. Fitness is measured in terms of how good a robot's behaviour is according to some evaluation criterion. This selectionist approach is particularly suited to the *exploration* of classes of networks involving many parameters and whose properties are difficult to predict in advance. The class of networks introduced in this paper are of that nature and have been investigated using evolutionary robotic techniques.

The focus of this paper is on ANNs using loose abstractions of biological phenomena; there is no *modelling* involved. However, for brevity and convenience, biological terminology is used frequently – it should be taken as analogy only. Having said that, the kind of work described in this paper can potentially have a useful relationship with more explicit modelling studies [16, 7]. Predicting how the spatial distribution of NO changes over time within defined neural structures is essential if we are to understand how nervous system function depends on gaseous signalling molecules. By modelling NO diffusion we can validate experimental results obtained from an intact biological preparation and thus help to establish that NO is both necessary and sufficient for the observed behaviour. Modelling also provides the opportunity to test hypotheses which cannot be validated in situ. Although the NO metabolic pathway can be interfered with pharmacologically, it is much less feasible to alter the properties of NO gas directly, for example, its diffusion constant or half life, which are likely to have significant effects on its functional role in the vervous system.

The remainder of this paper introduces a class of ANNs inspired by gaseous modulators, so called GasNets, experiments are described in which robot controllers built from these kinds of networks were evolved. Significant advantages over more standard ANNs are demonstrated, including a large reduction in the number of evaluations needed to develop successful controllers for visually guided behaviours.

2 GasNets

The networks used in the experiments described later consist of units connected together by excitatory links, with a weight of +1, and inhibitory links, with a weight of -1. The output , O_i, of a node i is a function of the sum of its inputs, as described by equation 1. In addition to this underlying network in which positive and negative 'signals' flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit 'gases' which diffuse and are capable of modulating the behaviour of other units by changing their transfer functions in ways described in detail later. This form of modulation allows a kind of plasticity in the network in which the intrinsic properties of units are changing as the network operates. The networks function in a 2D plane; their geometric layout is a crucial element in the way in which the 'gases' diffuse and affect the properties of network nodes. This aspect of the networks is described in more detail later.

$$O_i = tanh[k_i(t)(\sum_{j \in C_i} w_{ji} O_{j[t-1]} + I_{i[t]}) + b_i] \quad (1)$$

Where C_i is the set of nodes with connections to node i, $I_{i[t]}$ is the external (sensory) input to node i and b_i is a genetically set bias. Each node has a genetically set default $k_i(t)$. Figure 1 shows the family of curves generated for $tanh(kx)$ when k varies over a discrete set of values in the range [-4,4]. As can be seen, a wide range of output responses to a given input are possible, depending on the values of the parameter k. As will seen later, the default value of k for each node can be changed by diffusing gases as the network runs. This changes the shape of the node's transfer function providing the mechanism of modulation referred to in the previous paragraph.

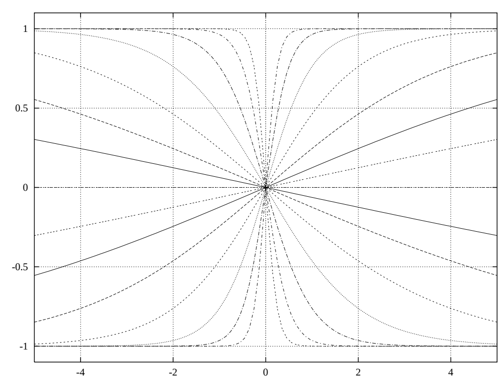

Figure 1: *Family of curves defined by $y = tanh(kx)$ transfer function for a range of values of k.*

Figure 2 shows a possible GasNet configuration. Node 4 can emit a gas and hence modulate nodes 5 and 6.

2.1 Gas Diffusion in the Networks

It is genetically determined whether or not a node will emit one of 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).

A very abstract model of gas diffusion is used. For an emitting node, the concentration of gas at distance d from the node is given by equation 2. Here r is the genetically determined radius of influence of the node, so that concentration falls to zero for $d > r$. $TC(t)$ is a linear function that models build up and decay of concentration after the node has started/stopped emitting. The slope of this

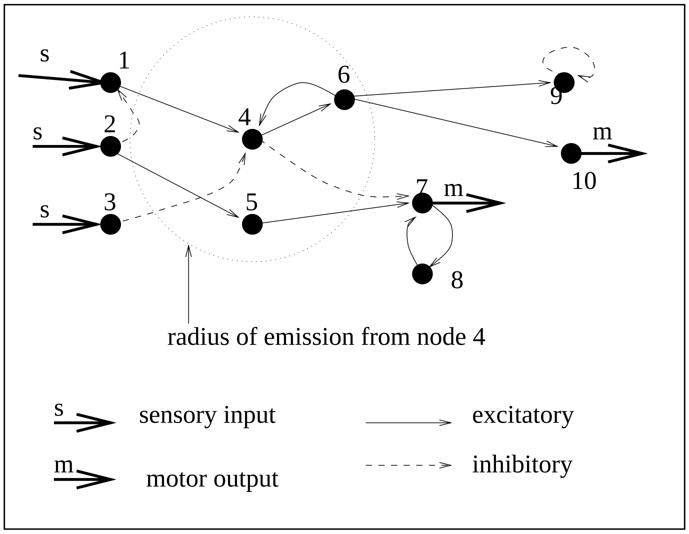

Figure 2: *A GasNet, node 4 can emit and hence modulate nodes 5 and 6.*

function is individually genetically determined for each emitting node, C_0 is a global constant.

$$C(d,t) = \begin{cases} C_0 \times e^{\frac{-2d}{r}} \times TC(t), & d < r \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

$$TC(t) = \begin{cases} H(\frac{(t-t_e)}{s}), & \text{emitting} \\ H(H(\frac{(t_s-t_e)}{s}) - H(\frac{(t-t_s)}{s})), & \text{not emitting} \end{cases} \tag{3}$$

Where, t_e is the time at which emission was last turned on, t_s is the time at which emission was last turned off, s (controlling the slope of the function) is genetically determined for each node and:

$$H(x) = \begin{cases} x, & x < 1 \\ 0, & x \leq 0 \\ 1, & \text{otherwise.} \end{cases} \tag{4}$$

In other words, the 'gas' concentration varies spatially as a Gaussian centred on the emitting node. The height of the Gaussian at any point within the circle of influence of the node is linearly increased or decreased depending on whether the node is emitting or not. Note $TC(t)$ saturates at a maximum of 1 and a minimum of 0. The total concentration at any point in the network is found by summing the concentrations from all emitting nodes.

2.2 Modulation by the Gases

The value of $k(t)$ in the network node transfer function (see equation 1) is changed (or *modulated*) by the presence of gases at the site of a node. Gas 1 increases the value of k in a concentration dependent way, while gas 2 decreases its values. This modulation is described by equations 6–8 and happens continually as the network runs. This provides a form of plasticity very different from that found in most traditional artificial neural networks. At every time step the value of k for node$i$, k^i, is updated according to equation 5.

$$k^i = PP[k^i_{index}] \tag{5}$$

Where,

$$k^i_{index} = S(N, k^i_{ni}) \tag{6}$$

Here k^i_{index} is node$i$'s index into the set PP of the possible discrete values k can assume. N is the number of elements in PP. In the experiments described later, $PP = \{-4.0, -2.0, -1.0, -0.5, -0.25, 0.0, 0.25, 0.5, 1.0, 2.0, 4.0\}$. At each time step k^i_{ni} is updated according to equation 7. The linear (thresholded) function S is described by equation 8.

$$k^i_{ni} = k^i_{def.index} + \frac{C_1}{C_0 \times K} \times (N - k^i_{def.index}) - \frac{C_2}{C_0 \times K} \times (k^i_{def.index}) \tag{7}$$

Where $k^i_{def.index}$ is the genetically set default value for k^i_{index}, C_1 is the concentration of gas 1 at the site of node$_i$, C_2 is the concentration of gas 2 at the site of node$_i$ and C_0 and K are global constants. So, k^i_{index} increases in direct proportion to the concentration of gas 1, and decreases linearly with respect to the concentration of gas 2. In this way the value of k for node$_i$ is changed by the presence of gases 1 and 2 at the node's site.

$$S(N, x) = \begin{cases} x, & 0 \leq x \leq N \\ 0, & x < 0 \\ N, & x > N \end{cases} \tag{8}$$

3 Minimal Simulations

Before describing in detail evolutionary robotics experiments using GasNets, some of the experimental methodology will be introduced. One potential problem with evolutionary approaches to exploring classes of robotic control systems is the time taken to evaluate behaviours over many generations. Recently, Jakobi has proposed new ways of thinking about and building fast-running easy-to-design minimal simulations for the evaluation of robot controllers. This methodology is described in detail elsewhere [12, 13], but since the experiments reported in this paper make extensive use of it, we offer a brief sketch here:

1. A small *base set* of robot-environment interactions that are sufficient to underly the behaviour we want to evolve must be identified and modelled. Because only this base set is modelled, some features of the simulation will have a basis in reality (the *base set aspects*), and some features will derive from the simulation's implementation (the *implementation aspects*).

2. Every implementation aspect of the simulation must be randomly varied from trial to trial so that controllers are unable to rely on them to perform the behaviour. In particular, *enough* variation must be included so that the only practicable evolutionary strategy is to actively ignore each implementation aspect entirely.

3. Every base set aspect of the simulation must be randomly varied from trial to trial. The extent and character of this random variation must be sufficient to ensure that reliably fit controllers are able to cope with the inevitable differences between the robot-environment interaction model and reality, but not so large that they fail to evolve at all.

The power behind these ideas derives from the fact that we only have to model a sufficient number of real-world features, and these do not even have to be modelled particularly accurately. This means that such simulations can be easily constructed and made to run extremely fast. As long as the right amount of variation is included according to the methodology outlined above, controllers that evolve to be reliably fit will almost certainly transfer into reality.

4 Experiments

A fairly large number of experiments have now been completed in which GasNet based robot controllers were developed for various tasks and robots [11]. Here we describe just one set of experiments on evolving GasNets to control a robot engaged in a visually guided behaviour.

This series of investigations made use of the Sussex Gantry Robot. In each case controllers were evolved using a minimal simulation. As explained earlier, such radical simulations run much faster than real time and have played a crucial role in allowing us to repeat the evolutionary experiments a sufficient number of times to gather meaningful statistics. Controllers evolved in minimal simulation work perfectly on the real robot. For details see [12, 13].

The gantry-robot is shown in figure 3. The robot body is cylindrical, some 150mm in diameter. It is suspended from the gantry-frame with stepper motors that allow translational movement in the X and Y directions, relative to a co-ordinate frame fixed to the gantry. Such movements, together with appropriate rotation of the sensory apparatus, correspond to those which would be produced by left and right wheels. The visual sensory apparatus consists of a CCD camera pointing down at a mirror inclined at 45^o to the vertical (see figure 4). The mirror can be rotated about a vertical axis so that its orientation always corresponds to the direction the 'robot' is facing. For full details see [9]. The gantry is a very useful apparatus for controlled experiments in the evolution of visually guided behaviours, but is probably best thought of as if it were a two wheeled mobile robot with a fixed video camera mounted on top.

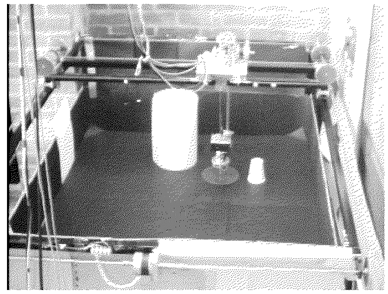

Figure 3: *The Gantry viewed from above. The horizontal girder moves along the side rails, and the robot is suspended from a platform which moves along this girder.*

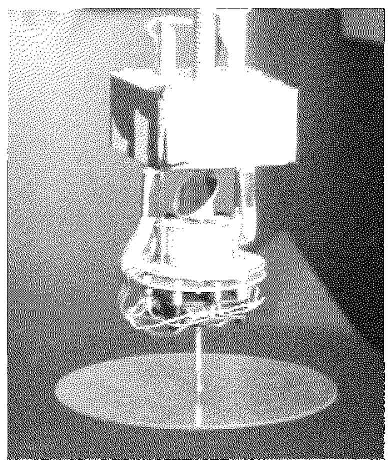

Figure 4: *The gantry-robot. The camera inside the top box points down at the inclined mirror, which can be turned by the stepper-motor beneath. The lower plastic disk is suspended from a joystick, to detect collisions with obstacles.*

4.1 The Task

A task was chosen for which we already had results from various evolutionary experiments with different styles of networks. This would allow direct comparison of the performance of the GasNets with more conventional connectionist nets.

Control networks were evolved for a target discrimination task. Two white paper targets were fixed to one of the gantry walls; a rectangle and an isosceles triangle with the same base width and height as the rectangle. Starting from a random position and orientation, the robot was required to move to the triangle while ignoring the rectangle. This was to be achieved under extremely variable and noisy lighting conditions in which the illumination intensity at any point in the gantry arena can vary by up to 100%. This was achieved by fixing a rig of spotlights above the gantry — the lights were randomly turned on and off at widely varying frequencies.

The network size and topology, as well as various other properties detailed below, were under unconstrained evolutionary control in every experiment (i.e. arbitrarily recurrent networks were possible). So was the robot visual morphology, i.e. the way in which the camera image was sampled. This was achieved by genetically specifying the number and position of *single* pixels from the camera image to use as visual inputs. The grey scale intensity value of these pixels (normalised into range [0.0,1.0]) were fed into the network, one for each genetically specified visual input node in the net. This is illustrated in figure 5. Note this means that the evolved control systems were operating with extremely minimal vision systems, just a few single pixel values. Given the very noisy lighting conditions and the minimal visual input, this was a non-trivial task.

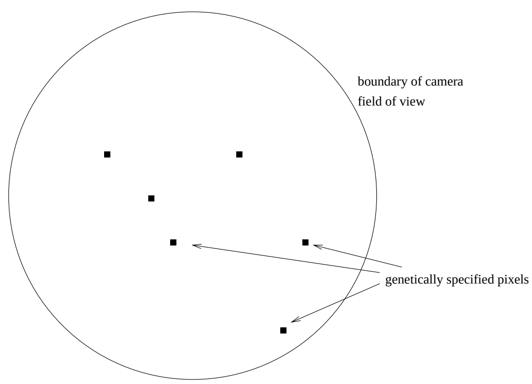

Figure 5: *Evolved visual morphology. Visual input is taken only from the genetically specified single pixels. The rest of the camera image is thrown away.*

4.2 Experimental Setup

GasNets were encoded on a genotype consisting of an array of parameter values. Each node in the network had 19 parameters associated with it, these are all under evolutionary control. That is:

$< geneotype > \;\; :: \;\; (< gene >)^*$

$< gene > \;\; :: \;\; < x >< y >< R_p >< \Theta_{1p} >< \Theta_{2p} >< R_n >< \Theta_{1n} >< \Theta_{2n} >$
$< vis_{in} >< vis_r >< vis_\theta >< vis_{thr} >< rec >< TE >< CE >< s >< R_e >$
$< k_{def.ind} >< bias >$

This encoding was used to generate networks conceptualized to exist on a 2D Euclidean plane. x and y give the position of a network node on the plane. The next six numbers define two segments of circles, centred on the node. These segments are used to determine the connectivity of the network. R_p gives the radius of the 'positive' segment, Θ_{1p} its angular extent and Θ_{2p} its orientation.

R_n, Θ_{1n} and Θ_{2n} define a 'negative' segment. The radii range from zero to half the plane dimension, the angles range from zero to 2π. The segments are illustrated in figure 6. Any node that falls within a positive segment has an excitatory (+1) link made to it from the segment's parent node. Any node that falls within a negative segment has an inhibitory (-1) link made to it from the segment's parent node.

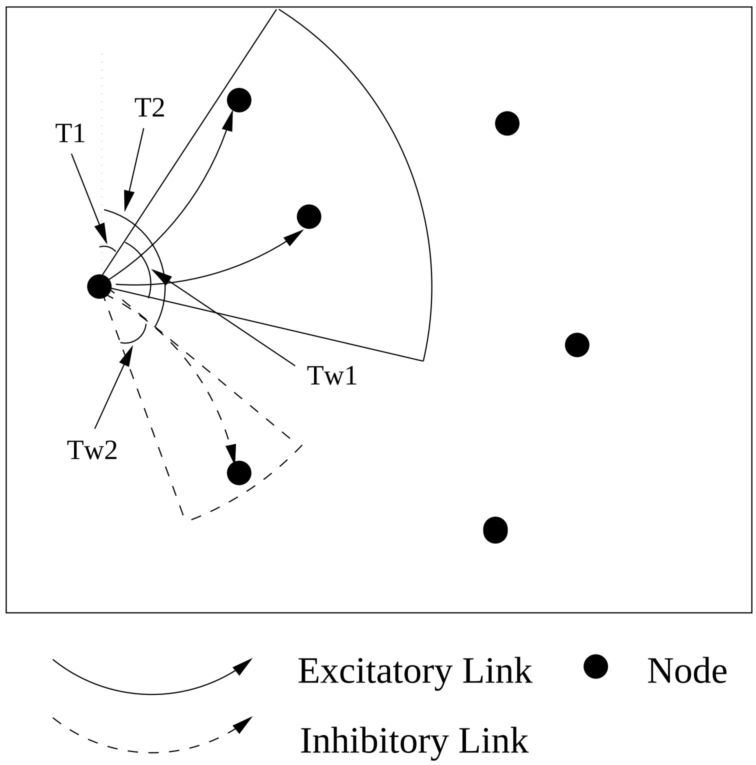

Figure 6: *Positive and negative segments define the connectivity of the network. The network develops and functions on a 2D plane.*

The rest of a gene is interpreted as follows. vis_{in} is a binary switch that determines whether or not a node has visual input. If it does, the following three parameters encode the polar coordinates of a pixel in the camera image the node will take input from, and a threshold below which input from that pixel is ignored. The value of rec determines whether the node has no recurrent connection to itself, an excitatory recurrent connection or an inhibitory recurrent connection, respectively. TE provides the circumstances under which the node will emit a gas. These are: not at all, if its 'electrical' activity exceeds a threshold, or if the concentration of the referenced gas (1 or 2) at the node site exceeds a threshold. CE gives the gas the node can emit. s is used to control the rate of gas build up/decay as described earlier by equation 3. R_e

is the maximum radius of gas emission, this ranges from 2 to half the plane dimension. $k_{def.ind}$ is the default value for the index used in equation 5 to determine the default values of k for the node. Finally, *bias* is the b_i term in the node transfer function (equation 1). The motor output nodes were put in fixed positions in the corners of the network grid.

The space of GasNet controllers was searched with an asynchronous distributed style GA in which a population of size 100 evolved on a torroidal grid [6]. Local neighbourhood selection rules were employed. No crossover operator was employed, the entire search process relied on various forms of mutation. There was a 3% chance of any parameter being mutated by an amount in the range ±10% of its full range. It was ensured that 20% of all mutations effected a small number of parameters that had a large effect on the network ($< rec >< vis_{in} >< TE >< CE >$). There was also a 0.3% chance that a selected genotype would have a gene (the full 19 parameters for a node) deleted or a random gene added.

On each fitness trial a robot controller was evaluated 8 times, starting from random positions and orientations. The evaluation score was the final distance away from the triangle after a fixed time period. The average of the eight scores was taken as the fitness of the controller.

5 Results

There is not enough room to report the results of the experiments in much detail. However, key observations can be made.

Jakobi had originally run the same experiment using binary networks in which the connectivity, weights on the connections and node thresholds were genetically encoded, along with the visual morphology [12]. He used a similar GA and comparable encoding scheme and was able to consistently evolve robust successful controllers after about 6,000 generations. Many of the GasNet runs produced successful controllers in less than 500 generations and they very rarely needed more than 1,000 generations.

Figures 7and 8 show examples of typical evolved successful controllers. They are structurally very simple, indeed much simpler than previously evolved binary networks [14]. Although the GasNets were very minimal, the modulatory interaction of the spreading and decaying gases and the sparse 'electrical' networks gave rise to surprisingly sophisticated internal dynamics.

A surprising observation is that *all* of the successful[1] evolved GasNet controllers (more than 50 to date) employ one or both of *only two* (closely related) behavioural strategies. The two strategies are illustrated in figure 9. The first strategy, illustrated to the left of the figure, involves moving until one of two strategically positioned visual inputs gives a high signal while the other gives a lower signal. The geometric layout of the sensors is such that this will only be

[1] To count as successful a controller must move to the triangle on many (at least 30) *successive* trials on the real robot under full noisy lighting and with random relative positioning of the two targets on the gantry wall.

reliably achieved when the robot is facing towards the triangle. The other strategy involves two vertically aligned visual sensors and is illustrated to the right of the figure. As the robot swings round towards a target, the bottom sensor will go high significantly earlier than the top sensor in the case of a triangle, but not for a rectangle. There is not enough space to describe the workings of the controllers in detail. However, it should be noted that many of them used a small number of additional visual inputs and various subtle internal dynamics to generate highly robust behaviours capable of coping with the extreme lighting conditions. In each case the visual morphology played a vital role. In each successful controller there was a perfect balance between the sensor geometry and robot motion resulting in active visual strategies. A traditional cognitive science perspective would think of the sensori capabilities as being passive and the sensor morphology as almost incidental; it is the internal processing where the real work is done. This is very clearly not the case in *any* of our evolved robots. The number and position of the visual inputs was under evolutionary control; it has clearly been demonstrated that very simple extremely low bandwidth sensors, when appropriately coupled to a dynamic controller, are sufficient for this kind of task.

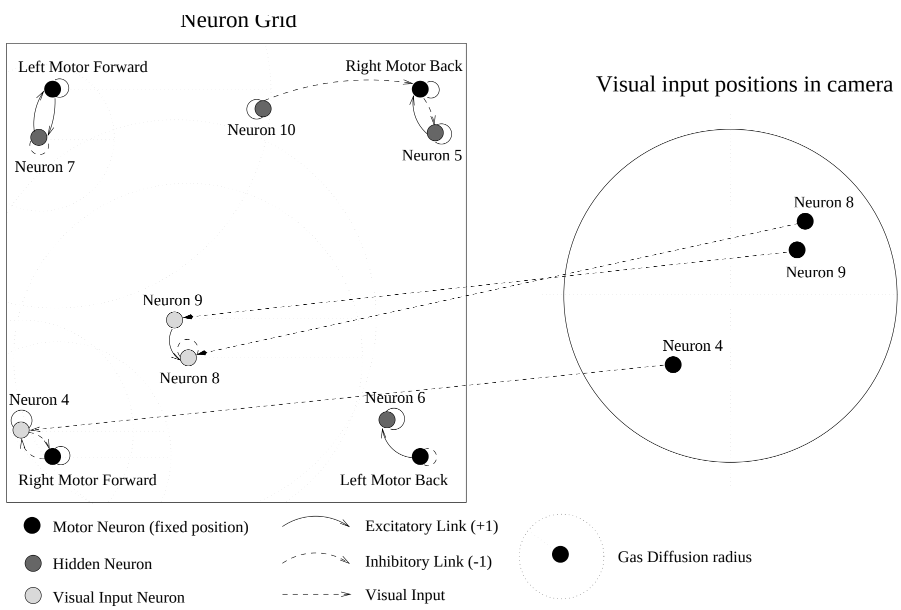

Figure 7: Closed-loop ‘tracking’ two-gas model triangle finding network, see text for details. NB gas radii are shown only where used.

Preliminary conclusions were that far fewer evaluations were needed to develop successful controllers using GasNets rather than more conventional binary networks (often an order of magnitude less), and that GasNets can provide successful robust controllers that are extremely simple structurally. But did this

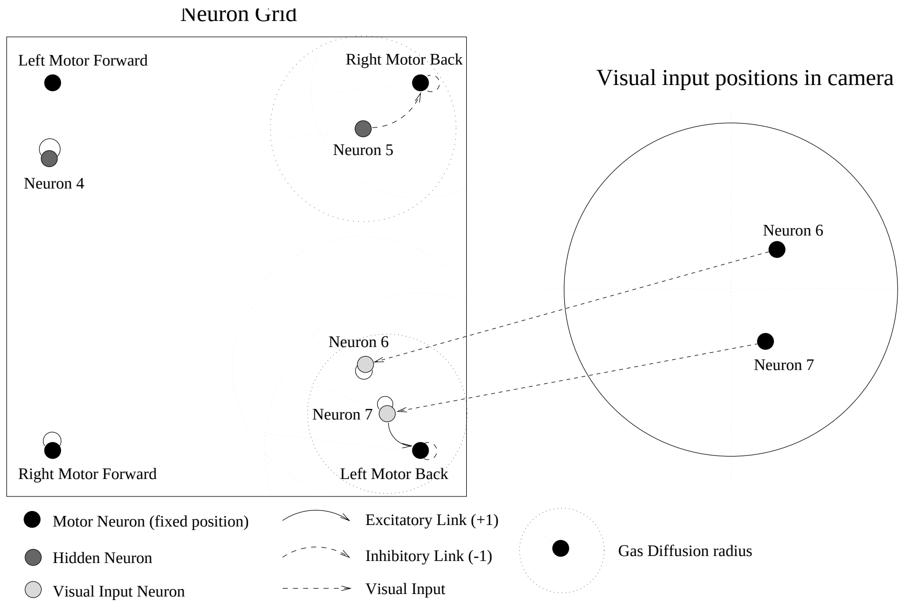

Figure 8: Open-loop 'ballistic' two-gas model triangle finding network, see text for details. NB gas radii are shown only where used.

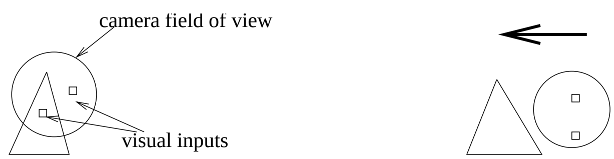

Figure 9: *The only two classes of successful behavioural strategy that we have observed to date.*

have anything to do with modulation by the gases? Two further sets of results very strongly suggest yes. The same experiment was run repeatedly with a different style of network that involved modulation by 4 gases [11]. Very similar results were observed, with rapid evolution of successful robots. For the style of networks described in this paper, the experiment was repeated 10 times as before and 10 times with the gas modulation turned off. Results are summarised in table 1.

Variable	Num. Cases	Mean	SD	SE of Mean
without gas	10	3305	2029.292	641.719
with gas	10	1305	1061.563	335.696

Table 1: *Number of Generations Before Consistent Success*

This gives a mean difference of 2000 generations. Levene's test for equality of variance gives F=1.846, P=0.191. Results of a t-test for equality of means is summarised in table 2. Clearly we can conclude that there is a significant (indeed, very significant) difference in evolving with and without the gases.

Variances	t-value	df	2-Tail Sig	SE of Diff	95% CI for Diff
Equal	2.76	18	0.013	724.220	(478.47, 3521.53)
Unequal	2.76	13.58	0.016	724.22	(442.213, 3557.787)

Table 2: *Summary of t-test.*

From these two results (the successful evolution of another type of GasNet, and the t-test summarised in the tables) we can conclude that the increased speed of evolution is not tied to a particular type of modulation of a particular type of network, and that the addition of a diffusing gas modulation mechanism produces a class of networks that are highly evolvable. This suggests that the space of possible behaviours open to being generated by GasNets is somehow 'thicker' than for more conventional networks. It is easier to find successful controllers in this space; it is rich with useful network dynamics and mechanisms.

6 Discussion

These preliminary investigations suggest that ANNs incorporating mechanisms analogous to those provided by diffusing gaseous neurotransmitters have interesting properties worth investigating further. The selectionist methodology of evolutionary robotics has proved to be a useful tool in exploring this class of networks. There are many possible future directions for this investigation. High on our agenda are studies involving larger, possibly more structured, networks; the investigation of a wider range of modulations – particularly longer lasting ones; the investigation of the concurrent evolution of structures acting as diffusion barriers or sinks within the networks; the use of gases to locally modulate hebbian style adaptive processes. Some of these studies will be more explicitly aimed at trying to better understand biological phenomena as well as developing artificial nervous systems.

References

[1] R. D. Beer. *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology.* Academic Press, 1990.

[2] R. A. Brooks. Challenges for complete creature architectures. In J.-A. Meyer and S.W. Wilson, editors, *From Animals to Animats: Proceedings of The First International Conference on Simulation of Adaptive Behavior*, pages 434–433. MIT Press/Bradford Books, Cambridge, MA, 1991.

[3] R. A. Brooks. Coherent behavior from many adaptive processes. In D. Cliff, P. Husbands, J.-A. Meyer, and S.W. Wilson, editors, *From Animals to Animats*

3: Proceedings of The Third International Conference on Simulation of Adaptive Behavior, pages 22–29. MIT Press/Bradford Books, Cambridge, MA, 1994.

[4] R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[5] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):73–110, 1993.

[6] R. Collins and D. Jefferson. Selection in massively parallel genetic algorithms. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth Intl. Conf. on Genetic Algorithms, ICGA-91*, pages 249–256. Morgan Kaufmann, 1991.

[7] J. Gally, P. Montague, G. Reeke, and G. Edelman. The no hypothesis: Possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. *Proc. Nat. Acad. Sci., USA*, 87:3547–3551, 1990.

[8] J. Garthwaite. Glutamate, nitric oxide and cell-cell signalling in the nervous system. *Trends in Neuroscience*, 14:60–67, 1991.

[9] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *From Animals to Animats 3, Proc. of 3rd Intl. Conf. on Simulation of Adaptive Behavior, SAB'94*, pages 392–401. MIT Press/Bradford Books, 1994.

[10] John Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, USA, 1975.

[11] P. Husbands. Evolving robot behaviours with diffusing gas networks. In P. Husbands and J.-A. Meyer, editors, *Evolutionary Robotics*, pages 71–86. LNCS vol.1468, Springer-Verlag, 1998.

[12] N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6(2):325–368, 1998.

[13] N. Jakobi. *Minimal Simulations for Evolutionary Robotics*. PhD thesis, University of Sussex, 1998.

[14] N. Jakobi, P. Husbands, and T. Smith. Robot space exploration by trial and error. In J. Koza, editor, *Proceedings of GP98*, pages 807–815. Morgan Kaufmann, 1998.

[15] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada. How to evolve autonomous robots: Different approaches in evolutionary robotics. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 190–197. MIT Press/Bradford Books, 1994.

[16] A. Philippides, P. Husbands, and M. O'Shea. Neural signalling – its a gas! In *Proc. ICANN98*. Springer-Verlag, 1998.

[17] D. Purves and George Augustine. *Neuroscience*. Sinauer, 1997.