

CSRP498

Better Living Through Chemistry: Evolving GasNets for Robot Control*

Phil Husbands^{1,2} and Tom Smith¹ and Nick Jakobi^{1,4} and Michael O’Shea^{1,3}

¹Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton, UK

²School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK

³School of Biological Sciences, University of Sussex, Brighton, UK

⁴AnimatLab, Ecole Normale Supérieure, Paris, France

{philh,toms}@cogs.susx.ac.uk, nick.jakobi@ens.fr, bafu7@central.sussex.ac.uk

Abstract

This paper introduces a new type of artificial neural network (GasNets) and shows that it is possible to use evolutionary computing techniques to find robot controllers based on them. The controllers are built from networks inspired by the modulatory effects of freely diffusing gases, especially nitric oxide, in real neuronal networks. Evolutionary robotics techniques were used to develop control networks and visual morphologies to enable a robot to achieve a target discrimination task under very noisy lighting conditions. A series of evolutionary runs with and without the gas modulation active demonstrated that networks incorporating modulation by diffusing gases evolved to produce successful controllers considerably faster than networks without this mechanism. GasNets also consistently achieved evolutionary success in far fewer evaluations than were needed when using more conventional connectionist style networks.

keywords: ANN, diffusible modulator, evolutionary robotics, GasNet.

1 Introduction

1.1 Robots

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 (Brooks, 1991a; Beer, 1990). 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. Brooks, 1994), 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. Purves, 1997). This in turn means that connectionist style networks, and even recurrent dynamical ones (Beer, 1990), 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

*A version of this paper will appear in *Connection Science* 10(4), 185–210, 1998.

inspiration in devising alternative styles of artificial neural network (Brooks, 1991b). The core of this paper is concerned with investigating abstractions of some of the extremely important chemical mechanisms of nervous systems and incorporating them into control networks for simple autonomous mobile robots.

1.2 Brains

Traditionally, chemical information flow in the brain has been thought to be mediated by messenger molecules or neurotransmitters which are released by neurons at points of close apposition known as synapses (Katz, 1969). Because most neurotransmitters are relatively large and polar molecules (amino acids, amines and peptides), they cannot diffuse through cell membranes and do not spread far from the release site. They are also rapidly inactivated by enzymatic hydrolysis and by active re-uptake. Together these features confine the spread of neurotransmitters very close to the points of release and ensure that the transmitter action is transient. In other words, chemical synaptic transmission of information operates essentially two-dimensionally (one in space and one in time). This conventional interpretation is coupled to the idea that neurotransmitters cause either an increase or a decrease in the electrical excitability of the target neuron. According to a traditional view of neurotransmission therefore, chemical information transfer is limited to the points of connection between neurons and neurotransmitters can simply be regarded as either excitatory or inhibitory.

In recent years two important discoveries have required a fundamental revision of this model. Firstly, many neurotransmitters, perhaps the majority, cannot be simply classified as excitatory or inhibitory (Hall, 1992). These messenger molecules are best regarded as ‘modulatory’ because among other things they regulate or modulate the actions of conventional transmitters. Modulatory neurotransmitters are also ‘indirect’ because they cause medium- and long-term changes in the properties of neurons by altering the rate of synthesis of so called ‘second messenger’ molecules. By altering the properties of proteins and even by changing the pattern of gene expression, these second messengers cause complex cascades of events resulting in fundamental changes in the properties of neurons. In this way modulatory transmitters greatly expand the diversity and the duration of actions mediated by the chemicals released by neurons. Secondly, the discovery that the gas nitric oxide (NO) is a modulatory neurotransmitter has opened entirely unexpected dimensions in our thinking about how chemical information is transmitted by neurons (Garthwaite et al., 1988; Gally, et al., 1990; Hölscher, 1997). Because NO is a very small and nonpolar molecule it diffuses isotropically within the brain regardless of intervening cellular structures (Wood and Garthwaite, 1994). NO therefore violates some of the key tenets of point-to-point chemical transmission and is the first known member of an entirely new class of transmitter, the gaseous diffusible modulators.

NO is generated in the brain by specialised neurons that contain the neuronal isoform of the calcium activated enzyme, nitric oxide synthase or nNOS (Bredt and Snyder, 1990). This enzyme catalyses the synthesis of NO from the amino acid L-arginine and molecular oxygen. NO synthesis is triggered when the calcium concentration in nNOS-containing neurons is elevated, either by electrical activity or by the action of other modulatory neurotransmitters. The existence of a freely diffusing modulatory transmitter suggests a radically different form of signalling in which the transmitter acts four-dimensionally in space and time, affecting volumes of the brain containing many neurons and synapses (Bredt and Snyder, 1992). The properties of NO that allow it to diffuse freely also prevent it from being stored prior to release in membrane bound vesicles at the synapse, as are the conventional neurotransmitters. This means that for NO to act as a signalling molecule, its release must be coupled directly to its synthesis. Because the synthetic enzyme nNOS can be distributed throughout the neuron and NO release does not require the synapse, NO can be generated and released by the whole neuron. NO is therefore best regarded as a ‘non-synaptic’ transmitter whose actions moreover cannot be confined to neighbouring neurons (Hartell, 1996; Park et al., 1998). NO cannot be classified conventionally as excitatory or inhibitory, it is a modulatory transmitter which activates the synthesis of cyclic-GMP,

an important second messenger which regulates a wide variety of cellular processes in target neurons, some of which underlie synaptic plasticity (Hölscher, 1997).

The discovery of diffusible gaseous modulators in the brain clearly challenges simplistic connectionist models of neural information processing (O’Shea et al., 1998). For example, it suggests a rich diversity of modulatory mechanisms with different temporal and spatial dynamics affect the properties of neurons. Importantly, the discovery of diffusible modulators shows that neurons can interact and alter one another’s properties even though they are not synaptically connected.

1.3 From Neuroscience to Control Systems

In this paper we have attempted to abstract some of these concepts and incorporate the elements of gaseous transmission into a fundamentally new class of artificial neural network. 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. This paper describes work where we have used this style of network to build control systems for autonomous mobile robots.

One of the new styles of AI to have emerged recently is Evolutionary Robotics (Cliff, Harvey and Husbands, 1993; Nolfi et al., 1994; Floreano and Mondada, 1994; Yamauchi and Beer, 1994; Husbands and Meyer, 1998). The evolutionary process, based on a genetic algorithm (Holland, 1975), 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 control networks involving many parameters and whose properties are difficult to predict in advance. The type of networks introduced in this paper are of that nature and have been investigated using evolutionary robotic techniques.

A word of warning: the focus of this paper is on ANNs using computationally efficient *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 (Philippides et al., 1998; Gally et al., 1990). Of course it should also be stressed that the physical language used in this paper (gas, diffusion etc.) is again only analogy. The networks are actually abstract discrete dynamical systems implemented as a C program. However, we feel that a profitable way to think of these systems are as computational simulations of physical devices. This particularly because they are used to control physical devices engaged in activities where space and time are highly pertinent. For instance, the coupling between the robot’s visual sensors and the control networks means that there is a direct relationship between the (modelled) spatial properties of the networks – crucial to the operation of the ‘diffusion’ processes – and the (actual) spatial properties of the sensor. It should be noted that it is quite possible that the class of networks introduced in this paper is mathematically equivalent, in a rather tortuous way, to varieties of other styles of ANNs. However, this issue will not be further considered here, although it is a topic worthy of future investigation. Rather than giving a formal treatment of the properties of this class of ANN, the focus of this paper is on exploring their evolvability when used for robotic control. Do they have advantages over other styles of ANN? What kinds of internal dynamics will emerge?

The next section of this paper introduces a class of ANNs inspired by gaseous modulators, so called GasNets. The next two sections after that describe the evolutionary robotics techniques and experimental setup used to evolve GasNet controllers for visually guided robots engaged in a target discrimination task under very noisy lighting conditions. In the results section 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. By performing many runs both with and without the gas modulation effects active, the importance of these effects in the performance of the networks is clearly demonstrated. The paper closes with conclusions and a brief discussion of future work.

2 GasNets

The networks used in the experiments described later are discrete time step dynamical systems built from units connected together by links that can be excitatory (with a weight of +1) or inhibitory (with a weight of -1). The output, O_i^n , of node i at time step n 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.

$$O_i^n = \tanh \left[k_i^n \left(\sum_{j \in C_i} w_{ji} O_j^{n-1} + I_i^n \right) + b_i \right] \quad (1)$$

Where C_i is the set of nodes with connections to node i , I_i^n is the external (sensory) input to node i at time n , and b_i is a genetically set bias. Each node has a genetically set default transfer function parameter k_i^0 . 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 be seen later, the value k_i^n for each node can be changed by diffusing gases as the network runs. Thus the actual shape of the node’s transfer function is altered via the gas modulation mechanism.




Figure 1: Family of curves defined by $y = \tanh(kx)$ transfer function for a range of values of k . Each curve shows the relationship between x (over the range $[-5,5]$) and y for a different value of k . The following set of k values are illustrated: $\{-4, -2, -1, -0.5, -0.25, -0.125, -0.0625, 0, 0.0625, 0.125, 0.25, 0.5, 1, 2, 4\}$.

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

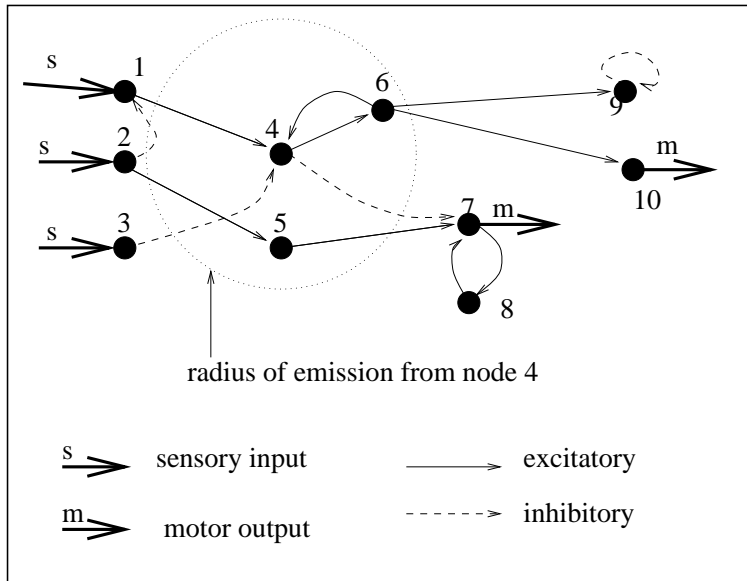


Figure 2: A GasNet: node 4 can emit gas and hence modulate nodes 5 and 6.

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). The electrical threshold used in the experiments described later was 0.5, the gas concentration threshold 0.1.

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 $d > r$. This is loosely analogous to the length constant of the natural diffusion of NO, related to its rate of decay through chemical interaction. $T(t)$ is a linear function that models the build up and decay of concentration after the node has started/stopped emitting (equation 3 and 4). The slope of this function is individually genetically determined for each emitting node, C_0 is a global constant.

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

$$T(t) = \begin{cases} H\left(\frac{t-t_e}{s}\right) & \text{emitting} \\ H\left(H\left(\frac{t_s-t_e}{s}\right) - H\left(\frac{t-t_s}{s}\right)\right) & \text{not emitting} \end{cases} \quad (3)$$

$$H(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x < 1 \\ 1 & \text{else} \end{cases} \quad (4)$$

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, and s (controlling the slope of the function) is genetically determined for each node.

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 $T(t)$ 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 transfer parameter value for the i th node at time step n , k_i^n (see equation 1), is changed (or *modulated*) by the presence of gases at the site of the node. Gas 1 increases the value of k_i^n in a concentration dependent way, while gas 2 decreases its value. This modulation is described by equations 5 to 7 and happens on every time step as the network runs. This provides a form of plasticity very different from that found in most traditional artificial neural networks.

$$k_i^n = \mathbf{P}[index_i^n], \quad \mathbf{P} = \{-4.0, -2.0, -1.0, -0.5, -0.25, 0.0, 0.25, 0.5, 1.0, 2.0, 4.0\} \quad (5)$$

Where,

$$index_i^n = f \left(index_i^0 + \frac{C_1^n}{C_0 \times K} (N - index_i^0) - \frac{C_2^n}{C_0 \times K} index_i^0 \right) \quad (6)$$

$$f(x) = \begin{cases} 0 & x \leq 0 \\ [x] & 0 < x < N \\ N & \text{else} \end{cases} \quad (7)$$

Where $\mathbf{P}[i]$ refers to the i th element of set \mathbf{P} , $index_i^n$ is node i ’s index into the set \mathbf{P} of possible discrete values k_i^n can assume, N is the number of elements in \mathbf{P} , $index_i^0$ is the genetically set default value for $index_i$, C_1^n is the concentration of gas 1 at node i on time step n , C_2^n is the concentration of gas 2 at node i on time step n , and C_0 and K are global constants (both set to 1 in this study). So, $index_i^n$ 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_i^n is changed over time by the presence of gases at the node’s site (the concentrations are governed by equation 2).

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 (Jakobi, 1998a, 1998b), 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.

In the experiments described later, all successful evolved controllers crossed the reality gap: they generated the same behaviours on the real robot. The great advantage of using minimal simulations in the work described here is that we were able to perform many complete evolutionary runs and hence to derive meaningful statistics.

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 (Husbands, 1998). 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 (section 3), 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.

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° 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. The camera image is 20 pixels in radius giving approximately 1250 pixels in total. However, as will be explained later, only a tiny proportion of these are used by the evolved controllers. The angle of acceptance of the camera is 39° . For full details see (Harvey et al., 1994; Husbands et al., 1997). 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 forward facing video camera mounted on top.

4.1 The Task

A task was chosen for which we already had results from various evolutionary experiments with different styles of networks (Jakobi et al, 1998). This would allow direct comparison of the performance of




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.*

the GasNets with more conventional connectionist nets. A further aim was to compare runs with and without the gas modulation effects active. The main dimension of comparison was number of evaluations needed until highly fit controllers emerged. Section 4.2.4 gives details of the exact investigations undertaken.

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); section 4.2.1 gives details of the network encoding and genetic algorithm. The robot visual morphology, i.e. the way in which the camera image was sampled, was also under genetic control. 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

4.2.1 Network encoding: genotype to phenotype mapping

Each network was encoded on a variable sized genotype coding for a variable number of nodes. A genotype consisted of an array of integer variables, each lying in the range $[0, 100]$. For continuous variables, the phenotype value is obtained by normalising the genotype value to lie in the range $[0.0, 1.0]$ and multiplying by the relevant variable range. For nominal values, such as whether the node has a visual input or not, the phenotype value = genotype value MOD N_{nom} , where N_{nom} is the number of possible nominal values, and MOD is the binary modular division operator (remainder of integer division). Each node in the network had either 19 or 21 variables associated with it, depending on which of two possible connectivity encoding schemes were used. These are described below. All variables were under evolutionary control. That is:

$\langle genotype \rangle :: (\langle gene \rangle)^*$

$\langle gene \rangle :: \langle x \rangle \langle y \rangle \langle R_p \rangle \langle \Theta_{1p} \rangle \langle \Theta_{2p} \rangle \langle R_n \rangle \langle \Theta_{1n} \rangle \langle \Theta_{2n} \rangle \langle vis_{in} \rangle \langle vis_r \rangle \langle vis_{\theta} \rangle$
 $\langle vis_{thr} \rangle \langle rec \rangle \langle TE \rangle \langle CE \rangle \langle s \rangle \langle R_e \rangle \langle index^0 \rangle \langle bias \rangle$

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(a). Any node that falls within a positive segment has an excitatory (+1 weighting) link made to it from the segment’s parent node. Any node that falls within a negative segment has an inhibitory (-1 weighting) link made to it from the segment’s parent node. If the segments intersect, nodes lying in the intersection will have both excitatory and inhibitory links made to them.

Figure 6(b) shows an alternative connectivity scheme that was also used. It is based on an encoding described in (Jakobi, 1998b). This time x, y coordinate points are used to define node connectivity. Each node has four outgoing connections, two excitatory and two inhibitory. Two variables (x, y), specify each link, defining the centre of a circle on the network plane. The nearest node to this centre within a threshold radius (10% of the plane) has a connection made to it. If no node lies within the

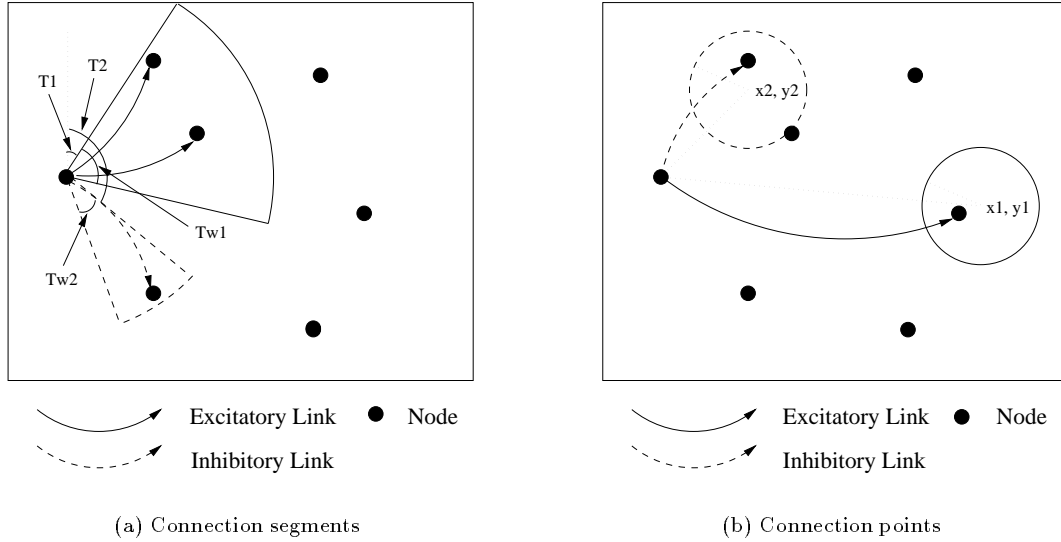


Figure 6: Connectivity of the network is defined by either positive and negative segments ($T=\theta$, $Tw=\theta_{width}$), or circles centred on x, y coordinates. Networks develop and function on a 2D plane.

threshold radius, no link is made. When this connectivity scheme is used, the six variables of the segments scheme are replaced with the eight needed to encode the four circle centres. Throughout the rest of this paper, these two schemes will be referred to as the *segments* and the *points* schemes respectively.

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 variables 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 (visual input is normalised to lie in the range $[0.0, 1.0]$, this is the range of the threshold). The value of rec determines whether the node has an excitatory recurrent connection, an inhibitory recurrent connection or no recurrent connection to itself. 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, its value ranges from 1 to 11. R_e is the maximum radius of gas emission, this ranges from 10%–60% of the plane dimension. $index^0$ is the default value for the index used in equation 5 to determine the transfer parameter value k_i^n for each node. Finally, $bias$ is the b_i term in the node transfer function (equation 1), restricted to the range $[-1.0, 1.0]$.

The first four nodes (the genotype must code for at least five nodes) are the motor nodes (left/right back and left/right forward), differing from other nodes in that the position on the plane is fixed (the four corners – see e.g. figure 10), and visual input is not permitted. Each of the four motor neurons is considered to be ‘on’ (+1.0) if $output > 0$, ‘off’ (0.0) otherwise. The actual wheel speeds are proportional to the output of the relevant forward node minus the output of the relevant backward node.

4.2.2 The Genetic Algorithm Scheme

The work described in this paper used a *distributed genetic algorithm* (Collins and Jefferson, 1991; Hillis, 1990) in which each solution is considered to occupy a unique position on a two-dimensional 10x10 toroidal solution grid. Initially, the grid is seeded with random genomes, each mapping on to a single neural network controller. Each solution receives a fitness score based on evaluation of its ‘triangle-finding’ ability (see section 4.2.3) before the main program loop is entered.

On each generation, the algorithm iterates *PopulationSize* times, choosing each time a random location on the solution grid. A *mating pool*, consisting of the current location plus the neighbouring 8 grid-points, is set up centred on the randomly chosen grid location. The mating pool solutions are ranked in order of fitness, and rank-based roulette selection is used to find the parent solution. Mutation (see below) of the parent solution produces an offspring solution which is evaluated and placed back in the mating pool at a point chosen by inverse rank-based roulette selection. Note no crossover is used; reproduction is asexual. Early exploratory experiments with various forms of crossover did not reveal any advantages in using any of these operators in addition to mutation, hence it was decided to employ a mutation only evolutionary algorithm for the study described here.

Rank-Based Roulette Selection

Linear rank-based selection is used, where the probability of selecting the i th solution from a mating pool of size N (where the solutions are ranked 0 to $N - 1$, with $N - 1$ being the fittest) is proportional to i :

$$P(i) = \frac{i}{\sum_{j=1}^{j=N} j} = \frac{2i}{N(N+1)} \quad (8)$$

Note this gives the lowest ranked solution zero probability of being selected as a parent, and the highest ranked solution zero probability of being replaced by the child.

Mutation

Four different mutation operators are applied, with different probabilities. The first two operate on the entire genotype, either deleting or adding a gene chosen or initialised at random. In the scheme used here, gene deletion is applied at a rate such that 0.8% of genotypes are altered, and gene addition at a rate of 1.0%. The third operator acts on every variable in the genotype, randomly altering it from its current value. The encoding scheme here uses integer variables in the range $[0, 100]$; the altering mutation operator rate is $0.8 \times 4\%$ per variable, and alters the value by an amount in the range ± 10 . This operator has no effect on nominally valued variables (see section 4.2.1). Finally, the random operator changes a variable to a new value chosen from a uniform random distribution over the whole range, its rate is $0.2 \times 4\%$ per variable. These values were based on those found to be successful in previous work (Harvey et al., 1994; Jakobi, 1998b).

4.2.3 Evaluation

N evaluations are carried out on an individual network, with scores f_i calculated on the fraction of the initial robot-triangle distance that the robot moves towards the triangle by the end of the evaluation; a maximum score of 1.0 is obtained by getting within 10.0cm of the triangle at any time during the evaluation. The controller only receives visual input; reliably getting to the triangle on a number of

trials with different starting conditions, can only be achieved by visual identification of the triangle. The evaluated scores are ranked, and the fitness F is the weighted sum of the N scores, with weight proportional to the *inverse* ranking i (ranking is from 1 to N , with N as the *lowest* score):

$$F = \frac{\sum_{i=1}^{i=N} if_i}{\sum_{i=1}^{i=N} i} = \frac{2}{N(N+1)} \sum_{i=1}^{i=N} if_i \quad (9)$$

Note the higher weighting on the poorer scores provides pressure to do well on *all* evaluations; a solution scoring 50% on every evaluation has fitness nearly 4 times that of one scoring 100% on half of the evaluations and zero on the other half.

4.2.4 Experimental Investigations

Early exploratory experiments with GasNets (Husbands, 1998; Jakobi, Husbands and Smith, 1998) on the gantry robot target discrimination task suggested that GasNets evolved considerably faster (fewer evaluations needed) than more conventional connectionist style networks used previously (Jakobi, 1998a). In order to probe this tentative result further, the following experiments were performed:

1. Ten runs to evolve GasNet-based controllers for the gantry triangle/rectangle task using the connection segments connectivity scheme (section 4.2.1).
2. Ten runs as in 1, but with the effects of the gases turned off.
3. Ten runs as in 1, but using the connection points connectivity encoding (section 4.2.1).
4. Ten runs as in 3, but with the effects of the gases turned off.

It was hoped that the results of these experiments would provide some insight into whether or not the gases played any significant role in the evolvability and functioning of the controllers, and whether or not the details of the connectivity encodings were important. Other experimental details were as described in the previous parts of this section.

5 Results

As mentioned earlier, the target discrimination task was chosen because a reasonable number of previous experiments had been carried out involving the evolution of more conventional connectionist style networks (Harvey et al., 1994; Jakobi, 1998b). An initial aim of the work described in this paper was to compare the performance of the GasNets with these other types of network. Therefore, before dwelling on the outcome of various GasNet runs, results of previous evolutionary experiments with this robot task are outlined.

5.1 Previous results

Jakobi had originally run the same experiment using two different styles of node and connection type networks (Jakobi 1998a, 1998b). He did a series of runs with simple *binary networks* consisting of nodes connected together by weighted links. Each unit used the transfer function given in equation

10. O_j is the output of the j th node and T_j is its threshold. The size and topology of the network was under evolutionary control, as were the connection weights, node thresholds and visual morphology. As with the GasNets, four motor units were used. Thresholds were real numbers in the range $[0.0, 1.0]$, the weights, w_{ij} , were real numbers in the range $[-2.0, 2.0]$.

$$O_j = \begin{cases} 0 & \sum_i O_i w_{ij} < T_j \\ 1 & \sum_i O_i w_{ij} \geq T_j. \end{cases} \quad (10)$$

He used a GA similar to that described in section 4.2.2 and a comparable encoding scheme to that described in section 4.2.1 (using the connections points method). He was able to consistently evolve robust successful controllers after about 6,000 generations. One of the best successful evolved binary net controllers is shown in figure 7. For a controller to count as sufficiently robust and successful, the robot it guides must move to the triangle on many (at least 30) *successive* trials on the real robot under full noisy lighting with random relative positioning of the two targets on the gantry wall as well as random initial positioning and orienting of the robot. This is the definition of robust, reliable, successful implicitly referred to in the rest of this paper.

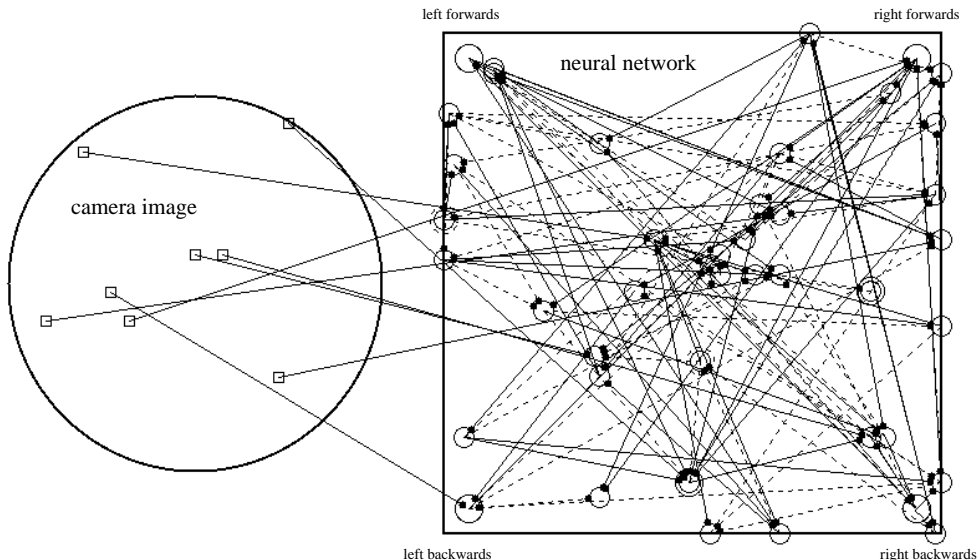


Figure 7: *Evolved network and visual morphology for binary networks. Solid lines are excitatory, dashed are inhibitory.*

Jakobi also successfully evolved controllers for this task using veto networks, a style of network that have been used at Sussex for various ER experiments over the past few years, including the original gantry work (Cliff, Harvey and Husbands, 1993; Harvey, Husbands and Cliff, 1994). The nodes in this style of network use separate channels for excitation and inhibition. Real values in the range $[0, 1]$ propagate along excitatory links. The inhibitory (or veto) channel mechanism works as follows. If the sum of excitatory inputs exceeds a threshold, $T_v = 0.75$, the value 1.0 is propagated along any inhibitory output links the unit may have, otherwise a value of 0.0 is propagated. Any unit that receives a non zero inhibitory input has its excitatory output reduced to zero (i.e. is vetoed). In the absence of inhibitory input, excitatory outputs are produced by summing all excitatory inputs, adding a quantity of noise, and passing the resulting sum through a simple linear threshold function, $F(x)$, given below. Noise was added to provide further potentially interesting and useful dynamics. The noise was uniformly distributed in the real range $[-0.1, 0.1]$. Again each network had four motor neurons (left/right forward/backward).

$$F(x) = \begin{cases} 0 & x \leq T_1 \\ \frac{x-T_1}{T_2-T_1} & T_1 < x < T_2 \\ 1 & x \geq T_2 \end{cases} \quad (11)$$

Where $T_1=0.0$ and $T_2=2.0$.

A simple direct encoding was used for the veto networks (Cliff, Harvey and Husbands, 1993). Again, the size and topology of the networks was under evolutionary control as was the visual morphology. Also in this case, about 6,000 generations were required to successfully evolve reliable robust veto net controllers (Jakobi, Husbands and Smith, 1998). An example is shown in figure 8.

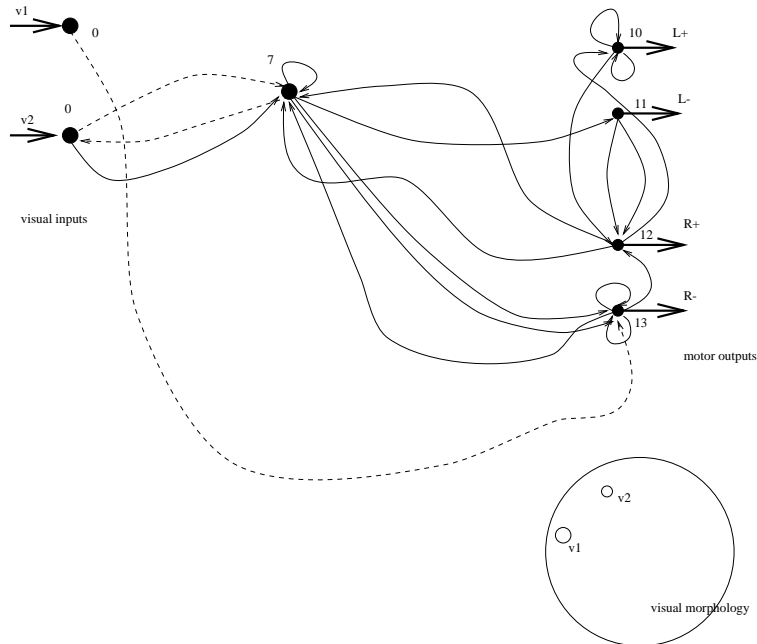


Figure 8: *Evolved network and visual morphology for veto networks. Active part of network is shown. Solid lines are excitatory, dashed are veto.*

5.2 GasNet Results

Many of the early exploratory GasNet runs produced successful controllers in less than 500 generations and they rarely needed more than 1,000 generations (Jakobi, Husbands and Smith, 1998). The successful networks were also structurally very simple in comparison to the evolved binary net and veto net based controllers. This suggested that the space of GasNet controllers was more amenable to evolutionary search than both the binary and veto net spaces. Table 1 summarises the results of 40 evolutionary runs designed to probe these initial tentative conclusions.

Table 1 shows the results for runs in each of the four conditions explained in section 4.2.4, and analysis of the difference between conditions with and without gas is shown in table 2. In both the connection segment and connection point schemes, significant difference ($P < 0.01$)¹ is seen between the with and without gas conditions. In both cases, the gas condition performed significantly better.

Clearly the addition of another distinct mechanism (gas modulation) to the fairly standard ‘processing units and wires’ networks has resulted in a class of networks underlying a space of robot behaviours

¹Significance is shown for the non-parametric Mann-Whitney U analysis. The parametric t-test gave similar results.

Run	Segs. w. gas	Segs. w/out gas	Points w. gas	Points w/out gas
1	300	1000	200	900
2	300	1000	400	1000
3	350	2000	450	1500
4	500	2400	500	2000
5	600	2800	700	2000
6	1500	3200	950	6000
7	1600	3300	1000	10000
8	2000	3950	1400	10000
9	2800	6400	1800	10000
10	3100	7000	2300	10000
Mean	1305	3305	970	5340
SD	1062	2029	677	4254
Median	1050	3000	825	4000
Best	300	1000	200	900
Worst	3100	7000	2300	10000

Table 1: Number of generations before consistent success is achieved. Data is shown for runs in each condition. NB runs not achieving consistent success by generation 10000 were terminated.

Scheme	Gas mean rank (N)	No gas mean rank (N)	U	2-tailed P
Segs.	7.1 (10)	13.9 (10)	16.0	0.0089
Points	6.75 (10)	14.25 (10)	12.5	0.0029

Table 2: Mann-Whitney U analysis of difference between with/without gas conditions.

that is easier to search, using an evolutionary approach, than those of more conventional networks. Successful GasNet-based controllers are relatively quickly found; it seems that the space is rich with useful network dynamics offering many relatively short paths to fit robot nervous systems. Although there is no restriction on the size of network used and the number of visual inputs employed, populations were initially seeded with fairly small networks (with 14 ± 1 units). We observed that the particular properties of the search space defined by the network encoding, the size of the network development/operation plane, the possible network dynamics, the robot geometry, the task, the properties of the visual sensor, and the *interaction* of all these things, has resulted in ‘easy’ routes to simple controllers employing very low bandwidth vision.

It can be seen that as long as the gas modulation mechanism is present there is little difference between the performances of the two network connectivity schemes. However, when the gases are turned off, the connection points scheme runs seem to quite often get stuck in low fitness areas of the search space (the 10,000 generation runs were terminated without finding a successful controller). This suggests that the connection segments scheme, which allows multiple connections to be made using a small number of variables, results in a slightly more amenable search space, possibly with less neutrality (Huynen et al., 1996). A more detailed investigation of this point would have to be made before anything very concrete could be concluded.

It should be noted that even without the gases active, the resultant heterogeneous networks are quite capable of generating successful behaviours. Indeed, on average the without-gas runs achieve success in significantly fewer generations than was necessary for the binary and veto nets discussed earlier.

Figures 9 and 10 show examples of typical evolved successful GasNet controllers. They are structurally very simple, indeed much simpler than previously evolved binary and veto networks, examples of which

are shown in figures 7 and 8. Although the nets were structurally sparse, the modulatory interaction of the spreading and decaying gases and the ‘electrical’ networks gave rise to surprisingly sophisticated internal dynamics.

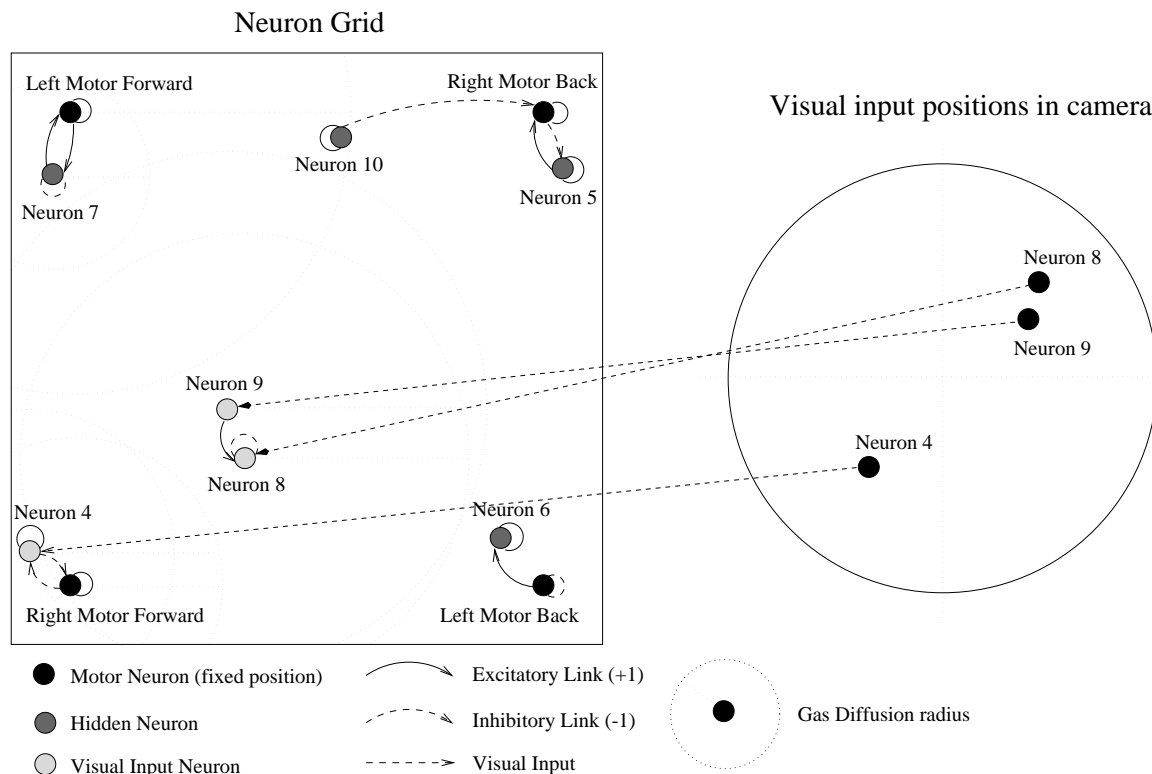


Figure 9: Closed-loop ‘tracking’ two-gas model triangle finding network, see text for details. NB gas radii are shown only where used.

Figure 11 shows pin traces of typical robot behaviours generated by the networks shown in figures 9 and 10. The head of the pin indicates the centre of the robot, and the tail shows the robot’s forward orientation. Movement towards the triangular target is always very direct. Remember the lighting conditions are extremely noisy and outrageously low bandwidth (2 or 3 pixel) vision is being employed.

5.3 Behavioural Strategies

A surprising observation is that *all* of the successful evolved GasNet controllers (more than 100 to date) employ one or both of *only two* (closely related) behavioural strategies. Indeed, all successful evolved binary net and veto net controllers also fell into these categories. The two strategies are illustrated in figure 12. The first strategy, illustrated in the left half 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 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

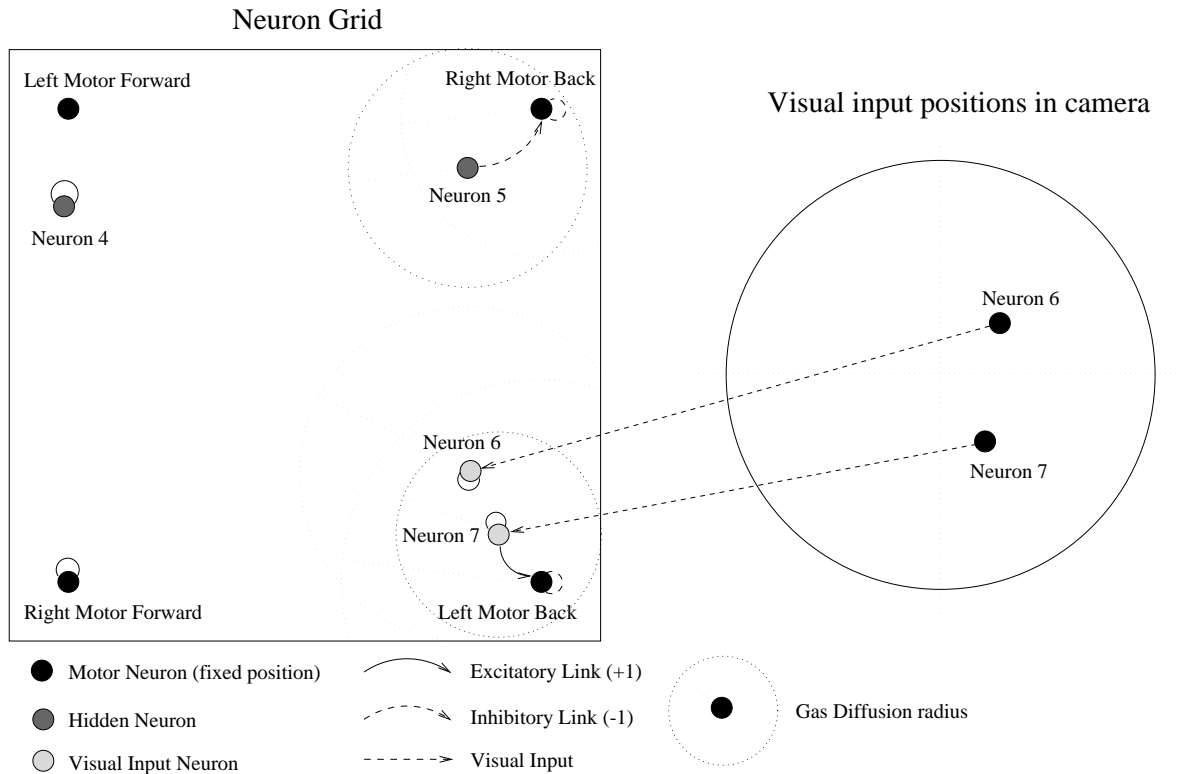


Figure 10: Open-loop ‘ballistic’ two-gas model triangle finding network, see text for details. NB gas radii are shown only where used.

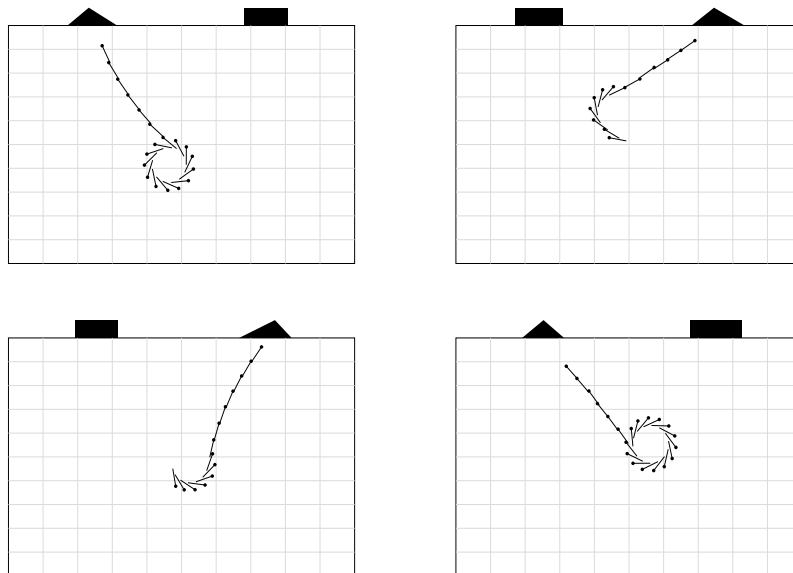


Figure 11: Four evaluations in simulation. The left-hand evaluations show the ‘ballistic’ open-loop triangle finder (from figure 10), while the right-hand evaluations show the closed-loop network (from figure 9). NB gantry robot and orientation are indicated by pin, and position is shown every 10 updates (roughly 1.0 seconds intervals).

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

5.4 Network Analysis

In order to gain deeper insight into the properties of GasNets and their suitability to act as robot nervous systems, it is necessary to analyse evolved successful controllers. There is not room in this paper to do a full network analysis (a paper devoted to that topic will appear in due course). Instead a summary of some of the key features arising from a detailed analysis will be given. Any analysis of a robot control network should take into account the robot-environment interaction dynamics, in this case partly dictated by the visual morphology, as well as the intrinsic network dynamics. Successful evolved GasNet controllers for the target discrimination task have all been structurally simple and have employed very low bandwidth sensing. These two factors have greatly helped in facilitating straightforward analysis. The mathematics of discrete dynamical systems (J. Sandefur, 1990) has proved a very useful framework for understanding the intrinsic properties of GasNets in detail, here only a brief qualitative analysis will be given.

The following paragraphs summarise an analysis of the controller shown in figure 10. This system has a number of interesting properties that nicely illustrate some general points about GasNets and evolved controllers at large.

The essential workings of the network are based on the two subnetworks in the right-hand corners of the network plane; both are required for accurate triangle finding behaviour, despite the lack of explicit interaction between them. None of the three neurons *not* involved in the subnetworks receive any external input; genetically set transfer function parameters and positive recurrency ensure that the outputs of neurons 0 and 1 (respectively the right and left forward motor neurons) stabilise at fairly high constant positive values. Neuron 4 has no functional role. Thus both forward motor neurons are continually on, and behaviour is governed by the two subnetworks acting on the back motor neurons.

The top right-hand corner subnetwork, built from units 2 and 5, produces a periodic output of neuron 2 (the right back motor neuron) in which the back motor neuron is generally off, but turns on for one time step in every eight; the motor neuron output trace is reminiscent of real neuron ‘spiking’ traces (Purves and Augustine, 1997), see figure 13. The bottom right-hand corner subnetwork, built from units 6 and 7, ‘switches’ into a particular extremely stable state, if and only if the robot scans visually across the triangle. Both networks rely heavily on gas diffusion effects (disabling diffusion results in the failure of both networks), and both are required for the overall triangle discrimination behaviour.

The unit 5 – unit 2 subnetwork behaviour is regulated by two interacting factors - the high negative recurrency on neuron 2, and the emission of gas by neurons 2 and 5. Neuron 2 emits gas1 when its electrical activation exceeds 0.5, neuron 5 is stimulated to emit gas2 by the presence of gas1. Neuron 5 receives no input whatsoever; genetically set transfer function parameters mean that its output




Figure 13: Output trace for ‘spiking’ neuron 2 over time. The vertical dashed lines mark 10 time step intervals, the area between the output curve and the time axis is shaded.

is always $\tanh(0.48) \approx 0.44$. Taking into accounts its genetically set parameters and the negative connection from unit 5, the neuron 2 transfer function can be written as:

$$y_2^n = \tanh(-k_2^n(y_2^{n-1} + 0.44) - 0.66)$$

The genetically set default value for k_2^n is 4, giving an unstable equilibrium solution at $y_2 \approx -0.48$. The instability comes from the high negative feedback; in the absence of any modulation the output of the unit swiftly turns into a saw-tooth oscillation between -1 and +1 with a period of 2 time steps. However, once its output reaches +1, neuron 2 emits gas1 with fast build-up/decay and low distance decay, affecting neuron 5 strongly and immediately for one time-step only. Neuron 5 in turn emits gas2 for one time-step, stimulated by the concentration of gas1. The subsequent concentration of gas2 at neuron 2 reduces k_2^n from 4 to 0.25. With this new value of k_2^n the subnetwork enters a different dynamical regime where a stable attractor exists for the output of unit 2 ($y_2 \approx -0.55$). However, the high gas concentration only lasts one time-step and k_2^n increases back to 4. In the single time-step with reduced k_2^n , output is dampened to $y_2 \approx -0.48$, then the next five time-steps are a ‘refractory’ period where the motor neuron output starts from near the unstable equilibrium slowly returning to ± 1 oscillation, and the cycle begins again. Thus the right back motor neuron has positive output once every eight time-steps, causing inhibition of the right motor by periodically turning it off (right forward is always on). This causes an average slowing of the right wheel and hence a slow curve to the right (if the left motor is producing positive output). This kind of oscillator circuit, where modulation causes a change in intrinsic dynamics, was observed to evolve on several evolutionary runs, and played a variety of functional roles. The oscillator’s period depends on the spatial relationship between the units and the genetically set parameters governing gas emission and transfer function properties.

The only neurons receiving external environmental input (6 and 7) are involved in the triangle discrimination network, which regulates the left back motor neuron through electrical and gas diffusion effects. Now, both neurons 6 and 7 have the same high threshold on the input below which visual input is not accepted; only intensities above 0.84 will have any affect. It can be shown that when neither visual input exceeds this threshold both units 6 and 7 stabilise with high negative output. This results in the left back motor neuron (with its negative recurrency) stabilising at a high positive output. This leads to the left motor being off (left forward is also high). The net effect is for the robot to circle anticlockwise (remember the role of the top righthand network is to periodically reduce a constant positive right motor signal).

Both units 6 and 7 emit gas2 when their electrical activity exceeds 0.5. Taking into account all the genetically set gas diffusion and transfer function parameters, it can be shown that if unit 6 receives bright input before, or only slightly after, unit 7, unit 7 is inhibited and does not emit gas or change the stable state of the left back motor output. However, if unit 7 receives high visual input first the left back motor unit is immediately inhibited and unit 7 emits gas2. This results in an inhibition of unit 6. Its transfer function is changed such that, no matter how high its input, unit 6’s activation cannot go

high enough to cause it to emit gas. Recurrency allows the subnetwork to stabilise with a high positive output for unit 7 (ensuring continued gas emission). The result is a highly stable state (that cannot be perturbed) in which the left back motor output is inhibited. Hence the overall behaviour is a slow turn to the right (left forward motor on all the time, right forward on most of the time; periodically turned off). From figure 10 we see that the visual inputs are geometrically arranged in the camera field of view such that input to unit 6 is stacked vertically above input to unit 7 with a reasonable separation between the two pixels.

Putting these pieces together we see that when both the visual inputs are less than very high (i.e. not looking at one of the targets) the robot circles anticlockwise. When unit 6 gets high visual input before, or at about the same time as, unit 7, the rotational movement continues unchanged. Note this will occur if the robot scans across the rectangle. However, if unit 7 gets high visual input before unit 6, the left motor is activated and the robot moves clockwise in a wide arc circle. Note this will occur if the robot scans across the triangle. Once this behaviour is initiated it cannot be turned off – the robot is in ballistic open-loop mode. From any starting position and orientation, the robot rotates anticlockwise until it scans across the triangle. Momentum effects always means that it slightly overshoots the triangle before the open-loop behaviour kicks in. The slight curve to the right means that the robot always ends up very close to the triangle although it is no longer using vision.

It should be noted that most successful GasNet based controllers employed closed-loop visually based control, always relying on an unbroken visuomotor feedback loop (e.g. the controller shown in Figure 9). However, the open-loop network described above illustrates some interesting and general internal mechanisms, and is a reminder of the fact that evolutionary processes often find ways of satisfying the fitness criteria that go against our intuitions as to how the problem should be ‘solved’.

5.5 Success with other Forms of GasNet

The basic network model presented in this paper is a refinement of a slightly earlier system described in (Husbands, 1998). The main difference between the earlier networks and those described in the current paper is that the units were governed by the transfer function shown in equation 12

$$O_i(S_i) = \begin{cases} 0 & S_i < 0 \text{ and } (a < 0 \text{ or } b < 0) \\ (S_i^a + S_i^b)/2 & \text{else} \end{cases} \quad (12)$$

Where O_i is the output of the i th node, S_i is the normalised input to the i th node, and a and b are parameters that are altered by 4 gases (2 for a and 2 for b) in much the same way as k is altered by 2 gases in the systems described in this paper (see section 2.2). All other details were more or less the same as described here.

Using this form of network robot controllers were successfully evolved in a few hundred generations to achieve the target discrimination task under very noisy lighting conditions. This is mentioned because it strongly suggests that the increased speed of evolution is not tied to a particular type of modulation of a particular type of network; there is nothing magic about the *tanh* transfer function used in this paper.

The earlier form of network was also used to successfully evolve various memory based behaviours on another style of mobile robot (Husbands, 1998). Again there was a considerable reduction in the number of evaluations needed in comparison with runs where more conventional networks were used. This, coupled with the results presented here for an essentially reactive visually guided task, suggests that the diffusing gas modulation mechanism may be fairly widely applicable.

6 Related Work

The authors are not aware of any directly comparable work to that described in this paper. However, a number of other pieces of research will be mentioned as having some kind of indirect relationship to the GasNets project. Floreano and Mondada have used a genetic algorithm to develop plastic controllers for a khepera robot (Floreano and Mondada, 1996) engaged in obstacle avoidance behaviours. The genetic algorithm was used to develop neural structures that were continually modified during the robot's lifetime according to mechanisms that were specified on the genotype. This was achieved by allowing the type of Hebbian-style rule operating at each connection to be selected by the genetic algorithm. They evolved learning structures with emergent fast adaptation properties. Interestingly, some of their best robot control networks generated stable behaviour through continuously changing synapses which were dynamically stable. There are perhaps strong similarities with the dynamic modulated states of the GasNets described earlier.

Ziemke has used higher order recurrent network architectures for controlling a khepera robot (Ziemke, 1996). He used two networks: one coupling sensors and motors and another which dynamically adapts the sensorimotor network. Essentially, this allows the sensorimotor control network to dynamically adapt its own weights to its current context. This allows a richer class of adaptations than is possible with more standard connectionist architectures.

The work of Beer and his students (e.g. Yamauchi and Beer, 1994) should be mentioned as providing a number of examples of robot controllers whose adaptivity is a result of the intrinsic dynamics of evolved recurrent dynamic ANNs rather than of an imposed adaptive mechanism in the network.

Of course there are a number of explicit computational modelling projects in which the mechanisms and functional role of various kinds of modulation are studied, based on specific biological data (Fellous and Linstner, 1998). More specifically, there is some work on modelling the diffusion and signalling properties of NO in real neuronal networks (Gally et al., 1990; Phillipides et al., 1998). The level of modelling used in such research is considerably more detailed than the abstractions employed in the GasNets work. Such modelling is computationally much more expensive than the algorithms underlying the GasNets. The loose abstractions used in the GasNets were chosen to ensure that they would be computationally efficient and capable of acting as real time controllers for a mobile robot. However, different, less arbitrary, abstractions may be used in the future, guided by the results of the detailed computational modelling.

7 Discussion and Conclusions

This paper has introduced a new class of ANNs incorporating principles abstracted from contemporary neuroscience. A simple form of modulation by processes analogous to diffusing gas, emitted by some nodes in the networks, have been added to heterogeneous arbitrarily recurrent networks (GasNets) used as artificial nervous systems for autonomous mobile robots engaged in visually guided behaviours. Evolutionary robotics techniques were used to evolve control networks and visual morphologies to enable a robot to achieve a target discrimination task under very noisy lighting conditions. A series of evolutionary runs with and without the gas modulation active demonstrated that networks incorporating modulation by diffusing gases evolved to produce successful controllers considerably faster than networks without this mechanism. GasNets also achieved evolutionary success much faster (often by an order of magnitude) than more conventional styles of networks previously used. The successful GasNet based controllers were structurally very simple (far simpler than successful controllers based on conventional networks) but exhibited intricate internal dynamics making full use of the modulatory effects of the diffusing gases.

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.

Acknowledgements

The initial phase of this work was done while PH was a visiting Professor at UniVap, Sao Jose dos Campos, SP, Brazil. We gratefully acknowledge the sponsorship of the British Council (SPA/126/881) and of FAPESP (proj. 96/7200-8) for that visit. Many thanks to Pedro Paulo Balbi de Oliveira, Andy Philippides, John Anderson, Emmet Spier, Tom Ziemke, Carmen Ruybal, three anonymous reviewers and various members of the CCNR, COGS and BIOLS for useful discussions of this work.

References

- Beer R.D. (1990) *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology*, Academic Press.
- Bredt DS and Snyder SH (1990) Isolation of nitric oxide synthetase, a calmodulin-requiring enzyme. *Proc Natl Acad Sci USA* **87**: 682-685.
- Brooks, R.A. (1991a) Intelligence without Representation. *Artificial Intelligence* **47**: 139-159.
- Brooks, R.A. (1991b) Challenges for Complete Creature Architectures. In: J.-A. Meyer and S.W. Wilson (Eds.), *From Animals to Animats: Proceedings of The First International Conference on Simulation of Adaptive Behavior*, 434-433, MIT Press/Bradford Books, Cambridge, MA.
- Brooks, R.A. (1994) Coherent Behavior from Many Adaptive Processes. In: D. Cliff and P. Husbands and J.-A. Meyer and S.W. Wilson (Eds.), *From Animals to Animats 3: Proceedings of The Third International Conference on Simulation of Adaptive Behavior*, 22-29, MIT Press/Bradford Books, Cambridge, MA.
- Cliff, D. and Harvey, I. and Husbands, P. (1993) Explorations in Evolutionary Robotics, *Adaptive Behavior* **2**(1): 73-110.
- Collins, R. and Jefferson, D. (1991) Selection in massively parallel genetic algorithms. In: R. K. Belew and L. B. Booker (Eds), *Proceedings of the Fourth Intl. Conf. on Genetic Algorithms, ICGA-91*, 249-256, Morgan Kaufmann.
- Fellous, J-M. and Linster, C. (1998) Computational Models of Neuromodulation, *Neural Computation* **10**: 771-805.
- Floreano, D. and Mondada, F. (1994) Automatic Creation of an Autonomous Agent: Genetic Evolution of a Neural-Network Driven Robot. In: D. Cliff and P. Husbands and J.-A. Meyer and S. Wilson (Eds.),

From Animals to Animats 3, Proc. of 3rd Intl. Conf. on Simulation of Adaptive Behavior, SAB'94, 421–430, MIT Press/Bradford Books.

Floreano, D. and Mondada, F. (1996) Evolution of Plastic Neurocontrollers for Situated Agents. In: P. Maes et al. (Eds), *From Animals to Animats 4, Proc. of 4th Intl. Conf. on Simulation of Adaptive Behavior, SAB'96*, 401–410, MIT Press/Bradford Books.

Gally JA, Montague PR, Reeke Jnr GN and Edelman GM (1990) The NO hypothesis: possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. *Proc Natl Acad Sci USA* **87**: 3547-3551.

Garthwaite J, Charles SL and Chess-Williams R (1988) Endothelium-derived relaxing factor release on activation of NMDA receptors suggests role as intracellular messenger in the brain. *Nature* **336**: 385-388.

Hall ZW (1992) *An Introduction to Molecular Neurobiology*. Sinauer Associates Inc, Sunderland, Massachusetts.

Hartell NA (1996) Strong activation of parallel fibres produces localized calcium transients and a form of LTD that spreads to distant synapses. *Neurons* **16**: 601-610.

Harvey, I. and Husbands, P. and Cliff, D. (1994) Seeing The Light: Artificial Evolution, Real Vision, In: D. Cliff and P. Husbands and J.-A. Meyer and S. Wilson (Eds.), *From Animals to Animats 3, Proc. of 3rd Intl. Conf. on Simulation of Adaptive Behavior, SAB'94*, 392–401, MIT Press/Bradford Books.

Hillis, W.D. (1990) Co-Evolving Parasites Improve Simulated Evolution as an Optimization Procedure *Physica D* **42**: 228–234.

John Holland (1975) *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, USA.

Hölscher C (1997) Nitric oxide, the enigmatic neuronal messenger: its role in synaptic plasticity. *Trends Neurosci* **20**: 298-303.

Husbands, P. and Harvey, I. and Cliff, D. and Miller, G. (1997) Artificial Evolution: A new path for AI? *Brain and Cognition* **34**: 130–159.

Husbands, P. (1998) Evolving Robot Behaviours with Diffusing Gas Networks, In: P. Husbands and J.-A. Meyer (1998), 71–86.

P. Husbands and J.-A. Meyer (Eds) (1998) *Evolutionary Robotics: Proceedings of 1st European Workshop, EvoRobot98*, Springer-Verlag LNCS vol.1468.

Huynen, M.A., Stadler, P.F. and Fontana, W. (1996). Smoothness Within Ruggedness: The Role of Neutrality in Adaptation. *Proc. Natl. Acad. Sci. (USA)* **93**:397-401.

Jakobi, N. (1998a) Evolutionary Robotics and the Radical Envelope of Noise Hypothesis, *Adaptive Behavior* **6**(2): 325–368.

Jakobi, N. (1998b) *Minimal Simulations for Evolutionary Robotics*, DPhil thesis, University of Sussex.

Jakobi, N. and Husbands, P. and Smith, T. (1998) Robot Space Exploration by Trial and Error. In: J. Koza et al. (Eds.), *Proceedings of GP98*, Morgan Kaufmann.

- Katz B (1969) *The release of neural transmitter substances*. Liverpool University Press.
- Nolfi S and Floreano D and Miglino O and Mondada F (1994). How to Evolve Autonomous Robots: Different Approaches in Evolutionary Robotics. In: R. Brooks and P. Maes (Eds.), *Artificial Life IV*, 190-197, MIT Press/Bradford Books.
- O'Shea M, Colbert R, Williams L and Dunn S (1998) Nitric oxide compartments in the mushroom bodies of the locust brain. *NeuroReport* **9**: 333-336.
- Park J-H, Straub V and O'Shea M (1998) Anterograde signaling by Nitric Oxide: characterization and in vitro reconstitution of an identified nitreergic synapse. *J Neurosci* **18** (in press)
- Philippides A and Husbands P and O'Shea M (1998) Neural Signalling – it's a Gas! In: L. Niklasson, M. Boden and T. Ziemke (Eds.) *ICANN98 – Proc. of the 8th Int. Conf. on Artificial Neural Networks*, 51–63, London: Springer-Verlag.
- D. Purves and G. Augustine (1997) *Neuroscience*, Sinauer.
- J. Sandefur (1990) *Discrete Dynamical Systems: Theory and Applications*, OUP.
- Wood J and Garthwaite J (1994) Model of the diffusional spread of nitric oxide - implications for neural nitric oxide signaling and its pharmacological properties. *Neuropharmacology* **33**: 1235-1244.
- Yamauchi, B. and Beer, R. (1994) Integrating Reactive, Sequential, and Learning Behavior Using Dynamical Neural Networks. In: D. Cliff and P. Husbands and J.-A. Meyer and S. Wilson (Eds.), *From Animals to Animats 3, Proc. of 3rd Intl. Conf. on Simulation of Adaptive Behavior, SAB'94*, 382–391, MIT Press/Bradford Books.
- Ziemke, T. (1996) Towards adaptive behaviour system integration using connectionist infinite state automata. In: P. Maes et al. (Eds), *From Animals to Animats 4, Proc. of 4th Intl. Conf. on Simulation of Adaptive Behavior, SAB'96*, 145–154, MIT Press/Bradford Books.