

**CSRP498**  
**Better Living Through Chemistry: Evolving  
GasNets for Robot Control\***

**Phil Husbands<sup>1,2</sup> and Tom Smith<sup>1</sup> and Nick Jakobi<sup>1,4</sup> and Michael O'Shea<sup>1,3</sup>**

**1**

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

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








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 tg00.4(reliably)-14999.4(fit)-16006theely







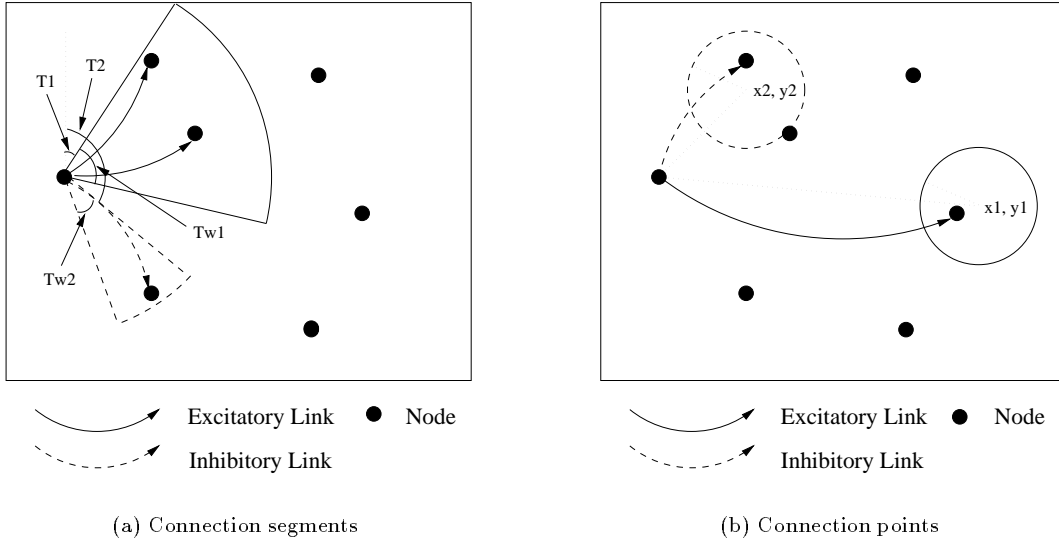


Figure 6: Connectivity of the network is defined by either positive and negative segments ( $T=\theta$ ,  $Tw=\theta_{width}$ ), or circles centred on  $x, y$  coordinates. Networks develop and function on a 2D plane.

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.  $E$  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.  $tproartorsvariablesrTfTreferrdtheTddthilSatdTdsTdfi fthadethreshrenysofreferrduserTfT$

## 4.2.2 The Genetic Algorithm Scheme

The work

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*

10.  $O_j$

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

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.



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  $14 \pm 1$  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







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 functions (syst19.48.88010Tdatcon)1003.34000.67(e000.373(extre02-11000.2-14




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  $\pm 1$  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 eniies.

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.

## 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*

*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

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: