The Shifting Network: Volume Signalling in Real and Robot Nervous Systems

Phil Husbands and Andy Philippides and Tom Smith and Michael O'Shea

CCNR, University of Sussex, Brighton BN1 9QH, UK
philh@cogs.susx.ac.uk

Abstract. This paper presents recent work in computational modelling of diffusing gaseous neuromodulators in biological nervous systems. It goes on to describe work in adaptive autonomous systems directly inspired by this: an exploration of the use of virtual diffusing modulators in robot nervous systems built from non-standard artificial neural networks. These virtual chemicals act over space and time modulating a variety of node and connection properties in the networks. A wide variety of rich dynamics are possible in such systems; in the work described here, evolutionary robotics techniques have been used to harness the dynamics to produce autonomous behaviour in mobile robots. Detailed comparative analyses of evolutionary searches, and search spaces, for robot controllers with and without the virtual gases are introduced. The virtual diffusing modulators are found to provide significant advantages.

1   Introduction

This paper describes some of the main thrusts of an ongoing interdisciplinary study of diffusing neuromodulators in real and artificial systems. After explaining the motivations and biological background of the project, the key results from recent detailed computational models of nitric oxide (NO) diffusion from neural sources are discussed. This leads to a description of work on more abstract artificial neural systems heavily inspired by the biology. These so-called GasNets are used as artificial nervous systems for mobile autonomous robots. Detailed comparative studies of evolutionary robotics experiments involving GasNets and non-GasNets are introduced. These include investigations into the formal evolvability of such systems. The paper closes with a sketch of current and future directions of the project.

2   Biological Background and Motivation

As the Brain Sciences have advanced is has become more and more clear that nervous systems are electrochemical devices of enormous complexity and subtlety [8,15]. While the transmission of electrical signals across neuronal networks is regarded as a

fundamental aspect of the operation of nervous systems, neurochemistry adds many dimensions to the picture. Cells can respond to chemicals that they themselves synthesize (autocrine signaling), to chemicals that diffuse from very nearby sources (paracrine signaling) or to chemicals that diffuse over greater distances or are carried by blood and tissue fluids [15,16]. The responses that these chemicals elicit are legion and can vary according to the internal and environmental states of the cells involved. Important classes of chemicals involved in the functioning of nervous systems include: neurotransmitters, receptors, neuromodulators and second messengers. Traditionally, chemical signaling between nerve cells was thought to be mediated solely by messenger molecules or neurotransmitters which are released by neurons at synapses [16] and flow from the presynaptic to postsynaptic neuron. 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 various reactions. Together these features confine the spread of such neurotransmitters to be very close to the points of release and ensure that the transmitter action is transient. In other words, chemical synaptic transmission of the classical kind 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 a number of important discoveries have necessitated a fundamental revision of this model. It is now clear that many neurotransmitters, perhaps the majority, cannot be simply classified as excitatory or inhibitory [8]. 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 influencing 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. The action of neurotransmitters also depends on the receptors they bind to. Although most receptors are highly selective, responding to a single transmitter only, most transmitters can bind to a variety of receptors, with different consequences for different transmitter receptor pairings, even in the same cell [16]. There are a great variety of receptors on different types of cells suggesting the possibility of a combinatorially explosive range of pairings and effects. However, when coupled with this expanded picture of the nervous system, it is the recent discovery that the gas nitric oxide is a modulatory neurotransmitter that has opened entirely unexpected dimensions in our thinking about neuronal chemical signaling [5,6,10]. Because NO is a very small and nonpolar molecule it diffuses isotropically within the brain regardless of intervening cellular structures [25]. 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 diffusable modulators (carbon monoxide is another example). NO is generated in the brain by specialised neurons that contain the neuronal isoform of the calcium activated enzyme, nitric oxide synthase or nNOS [1]. 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 [1]. NO cannot be contained by biological membranes, hence its release must be coupled directly to its synthesis. Because the synthetic enzyme nNOS can be distributed throughout the neuron, 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 [9,18].

The emerging picture of nervous systems sketched above -- as being highly dynamical, with many codependent processes acting on each other over space and time -- is thoroughly at odds with simplistic connectionist models of neural information processing. Importantly, the discovery of diffusible modulators shows that neurons can interact and alter one another's properties even though they are not synaptically connected. Indeed, all this starts to suggest that rather than thinking in terms of fixed neural circuits, a picture involving shifting networks – continually functionally and structurally reconfiguring – may be more appropriate. Of course many Alife practitioners reject the simple information processing models, but even so, by far the most popular kind of system used as artificial nervous systems are networks of nodes connected by virtual 'wires' along which inhibitory or excitatory 'electrical' signals flow. Although few would claim these are adequate models of the brain, their origins are in principles abstracted from the neuroscience of several decades ago. Although there are many possible levels of abstraction, new styles of artificial nervous systems directly inspired by contemporary understandings of brains as electrochemical machines may be a very fruitful avenue in our quest to develop artificial systems capable of more interesting and useful adaptive behaviours than we can currently manage. At the same time, the study of such systems should bring us deeper understandings of the principles underlying the functioning of real brains.

Given the limitations of current technology, if implemented versions of such systems are to act in real time as sensorimotor control systems for autonomous agents, they must necessarily abstract away much of the detailed complexity of real nervous systems. However, we believe that enough will be left behind to make this a worthwhile endeavour. In tandem with this kind of biologically inspired investigation, there is another far more detailed, more direct, form of modelling that can also be very useful. It is, as yet, very difficult to gather detailed empirical findings on such phenomena as the diffusion dynamics of NO in different parts of the nervous system, because the necessary experimental apparatus has not yet been developed. However, it is possible to build detailed computational models that capture certain salient features of these phenomena in an accurate way. These models are computationally expensive and do not run in real time, but the data they produce can make important contributions to our understanding of the biological processes.

This paper discusses examples of both kinds of work; two aspects of our ongoing investigation of the role of diffusing neuromodulators.


3   Modelling NO Diffusion in Real Neuronal Networks

In the previous section the role of NO in neuronal volume signalling was sketched. NO spreads in three dimensions away from the site of synthesis regardless of intervening cellular or membrane structures [25]. Another very important feature of NO signalling follows from the fact that nitric oxide synthase is soluble and thus highly likely to be distributed throughout a neuron's cytoplasm. This means that the whole surface of the neuron is a potential release site for NO, in marked contrast to conventional transmitter release. These properties suggest that the 3D structure of the NO source, and of any NO sinks, will have a profound influence on the dynamics of NO spread. Hence an accurate structure-based model of neuronal NO diffusion is an indispensable tool in gaining deeper insights into the signalling capacity of the molecule.

Figure 1 shows the results generated by the first accurate model of NO diffusion from continuous biologically realistic structures [19]. The source is an irregular neuron-like structure where the main cell body is a hollow sphere (NO is synthesized in the cell walls but not in the interior of the sphere). A sink has been placed just to the right of the cell body. Diffusion was modelled using accurate difference equation methods on a fine grid [19]. The figure shows the evolution of NO concentration during and after a 100ms burst of synthesis. Two very interesting observation are that the concentration remains high near the centre of the cell body long after synthesis has finished and that there is a significant delay between the start of synthesis and a rise in concentration for points distant from the main cell body. These observations follow from a 'reservoir effect' where NO diffuses into the centre of the hollow structure and is then 'trapped' there by a pressure gradient resulting in a slow time-delayed release [19]. Such a phenomenon, with its possible functional implications, would not have been observed in a less accurate point-source type model [25].

NO is also synthesized in another kind of irregular structure – namely a mesh of fine neuronal fibres in the mammalian cerebral cortex [20]. This mesh, or plexus, arises from a small population of neurons. As one of the biological affects of NO is to dilate the walls of blood vessels, the plexus might mediate the link between increased neural activity and increased blood supply to the same volume of the cortex. However, the vast majority of fibres in the plexus have been shown to be too small to generate above (biological) threshold concentrations of NO. This situation is again ripe for investigation with computational models. Using the same techniques as for the study illustrated in Figure 1, Philippides et al. have modelled the diffusion of NO from plexus structures [20]. Figure 2 shows results from a model investigating the volume over which NO concentrations are above threshold for sources made from regular arrays of very fine tubular structures. We see that once the density of fibres rises above a certain limit, the concerted effect of several very fine sources is to raise concentrations to significant levels. Further computational studies have shown how a random mesh of fine (rather than course) fibres is an ideal structure to ensure a

uniform concentration over the plexus [20]. This is exactly the kind of structure found in the cortex, hence these models may point towards an important mechanism for allowing highly targeted NO 'clouds' in the brain.
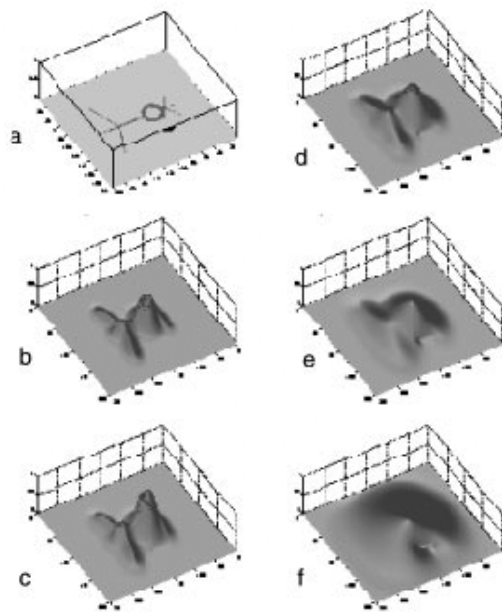
a) position of neuron and sink, b) NO concentration at t=50ms during synthesis, c) concentration at t=100ms: the end of synthesis, d) conc. at t=150ms, e) conc. at t=250ms, f) conc. at 750ms.

Fig. 1. Diffusion of NO from an irregular neuron being influenced by a nearby sink. NO concentration is shown at several time intervals following the initiation of a 100ms burst of synthesis. A 2D slice through the structure is illustrated here. See text for further details.
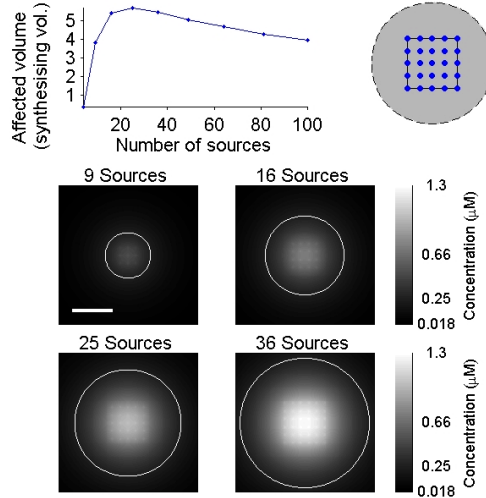
Fig. 2. Different numbers of very fine tubular NO sources arranged in regular arrays affect different volumes of tissue.

4   GasNets: From Neuroscience to Engineering

This section describes one style of artificial neural network from a class of networks whose operation is strongly inspired by those parts of contemporary neuroscience that emphasize the complex electrochemical nature of real nervous systems. So-called GasNets incorporate virtual diffusing gaseous neuromodulators and are used as artificial nervous systems for mobile autonomous robots. They are being investigated as potentially useful engineering tools and as a way of gaining helpful insights into biological systems. While a number of authors have incorporated global analogues of chemical signalling systems into agent control systems [2,7], as far as we are aware this work, which dates back to several years ago [11,13], is the first to concentrate on local processes, with virtual modulators diffusing over space and time.

   The basic GasNet networks used in many recent experiments [13] 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^t, of node i at time step t 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 virtual '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, as illustrated in Figure 3. This aspect of the networks is described in more detail later.
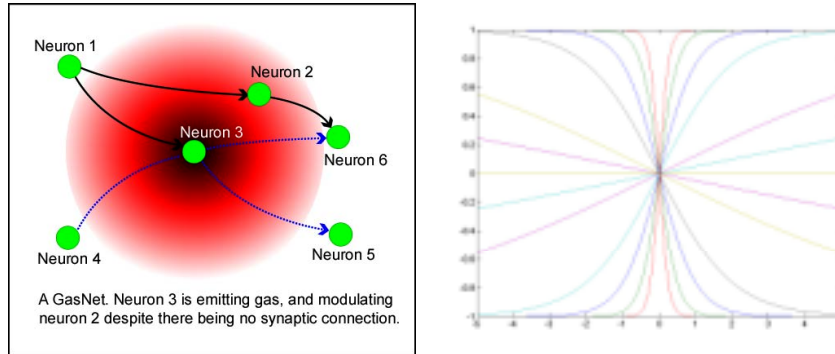
Fig.3. GasNet operation depends on the geometric layout of the nodes in a 2D plane. The righthand side of the diagram shows how the shape of the tanh transfer function depends on the gain parameter k_i^t , see text for further details.

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

Where C_i is the set of nodes connected to node i, I_i^t is the external (sensory) input to node i and b_i is a genetically set bias. Each node has a genetically set default transfer function gain k_i^0. The right hand side of Figure 3 shows the shape of the function tanh(kx) over the range [-5,5] for a discrete set of values of k between –4 and 4. It is this gain parameter that is modulated by the diffusing virtual gases in the networks. This means that while the gases are active the shapes of the node transfer functions are being altered from time step to time step. The mechanism for this is explained in the next section.

4.1 Diffusion and modulation

The virtual diffusion process is simple in order to be computationally fast so that GasNets can be used to control robots in real time. For mathematical convenience there are two gases, one whose modulatory effect is to increase the transfer function gain parameter and one whose modulatory effect is to decrease it. It is genetically

determined whether or not any given 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 one of the gases (genetically determined) 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. Allowing these two highly biologically inspired possibilities [6,10] is important – it provides a mechanism for rich interaction between two processes, the 'electrical' and the 'chemical'. 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_i is the genetically determined radius of influence of the ith node, so that concentration falls to zero for $d>r_i$. This is loosely analogous to the length constant of the natural diffusion of NO, related to its rate of decay through chemical interaction. $T_i(t)$ is a linear function that models the build up and decay of concentration after the node has started/stopped emitting. The slope of this function is individually genetically determined for each emitting node, C_0 is a global constant. For full details see [13].

$$C_i(d,t) = C_0 e^{-2d / r_i} \times T_i(t) \tag{2}$$

At each time step the gain parameter, k_i^t, for the node transfer function at each node (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^t in a concentration dependent way, while gas 2 decreases its value. Concentration contributions from nodes within range of any given site are simply added together. The modulatory effects of the two gases are then summed to calculate the value of k_i^t at each time step. Each node has its own default rest value for the gain parameter, the virtual gasses continually increase or decrease this value. Referring to the right-hand side of Figure 3, this modulation can potentially have drastic effects on a nodes's transfer function, dramatically increasing or decreasing, or even flipping the sign of, its slope. This means that the networks are usually in flux, with rich dynamical possibilities.

Since there were no pre-existing principles for the exact operation and design of such networks, it was decided to allow most of their detailed properties to be genetically specified, giving the possibility of highly non-uniform dynamically complex networks. Hence, in most experiments to date nearly everything is up for grabs: the number of nodes in a network; the way they are connected; the position of the nodes on the 2D plane; the individual properties of each node controlling when (if at all) they emit a gas; which gas is emitted and how strongly; how and if nodes are connected to sensors or motors, as well as various properties of the sensors and motors themselves [13]. About 20 variables per node are needed to describe all this. Our experience has been that a well setup evolutionary search algorithm is a good tool for exploring the space of such systems [12], looking for interesting and useful examples that deepen our understanding of autonomous adaptive systems or provide practical engineering advantages such as robustness and reliability [12,17]. The next section gives an example of using GasNets in such an evolutionary robotics setting.

5   Experimental Comparison

Various forms of GasNet have been used as robot controllers for a variety of tasks and robots [11,13]. A very large number of runs of one particular experimental setup have been carried out, giving us enough data to be able to make statistically significant claims. In this series of experiments GasNets were evolved to control a robot engaged in a visually guided behaviour involving shape discrimination. A simple robot with a fixed CCD camera as its main sensor moved in an arena as illustrated in Figure 4. Two light coloured shapes, a rectangle and a triangle, were placed against a darker background on one of the walls. The task was to reliably move to the triangle, while ignoring the rectangle, from a random initial orientation and position under highly variable lighting conditions. The relative positioning of the shapes, in terms of which was on the left and which on the right, was made randomly.

As well as network size and topology, and all the parameters controlling virtual gas diffusion and modulation, the robot visual morphology, i.e. the way in which the camera image was sampled, was also under unconstrained genetic control. This was achieved by allowing the evolutionary search algorithm to specify 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 the 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 the bottom left quadrant of Figure 4. Note that 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, the shape discrimination task becomes non-trivial.

All the evolutionary runs were carried out using a Jakobi minimal simulation of the robotic setup. The methodology behind these ultra-lean ultra-fast simulations was developed by Jakobi [14] to address one potential problem with evolutionary approaches to exploring classes of robotic control systems: the time taken to evaluate behaviours over many generations. Through a careful use of noise and important decisions about what not to model, a minimal simulation will run very fast but behaviours evolved in them will transfer to the real robots. For full details of the minimal simulation used for the triangle rectangle task see [14]. In the experiment described here, all successful evolved controllers crossed the reality gap: they generated the same behaviours on the real robot as in simulation. Success was defined as being able to move to the triangle and stay there 30 times in direct succession from random starting positions and orientations, under very noisy lighting conditions and irrespective of the relative positioning of the shapes on the same wall. The great advantage of using minimal simulations in this work is that we were able to perform many complete evolutionary runs and hence derive meaningful statistics.
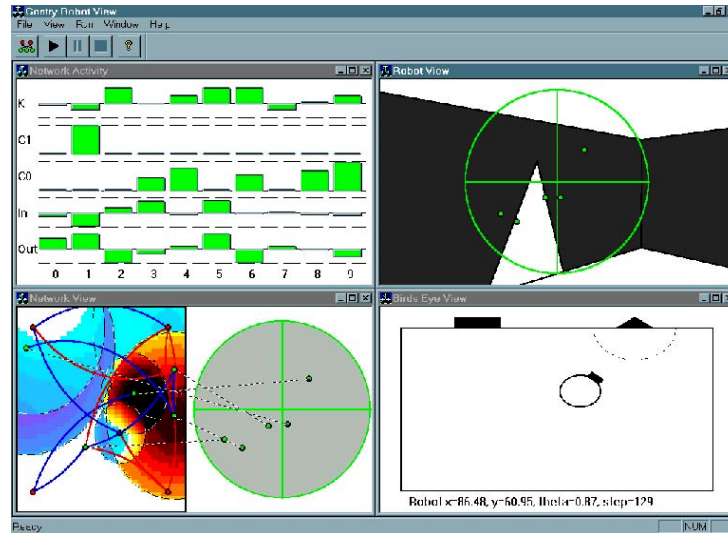
Fig. 4. The visualisation tool used with the minimal simulation of the shape discrimination task. The top right quadrant shows the view through the robot's camera, the bottom right gives a bird's eye view of the robot moving in the arena. The left-hand side of the screen illustrates the structure (including visual morphology) and functioning of the GasNet controlling the evolved robot. The shading in the network representation at extreme bottom left shows the gas concentrations in the network plane at the instant the snapshot was taken. The darker the shading the higher the concentration. See text for further details.

The initial set of evolutionary GasNet experiments with this task resulted in highly robust controllers emerging about 10 times faster than in earlier runs with conventional connectionist networks [13]. Subsequent comparative runs have concentrated on identifying whether or not the virtual gas modulation mechanism was at the root of this speed up. The key result is illustrated in Figure 5. In all experiments the genotypes were strings of integers encoding the various properties of the controllers and coupled visual morphologies, a geographically distributed genetic algorithm was used with a population of 100 [4], a number of different mutation operators were used in tandem, including node addition and deletion operators. The fitness function was based on a weighted average of final distances to the triangle over a set of evaluations from different initial conditions and different relative positioning of the shapes. Poor scores were weighted more heavily than good scores, encouraging robustness by requiring uniformly high scores across the whole evaluation set.
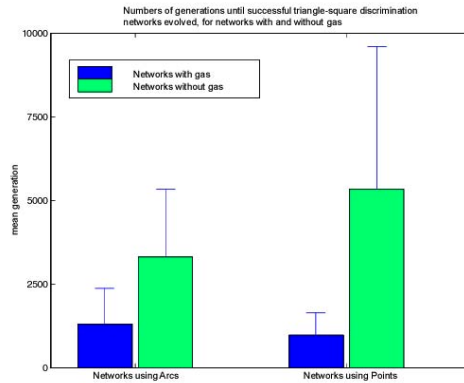
Fig.5. The average number of generations needed to find controllers giving perfectly successful behaviour on the triangle rectangle problem. The dark columns are for networks with the gas mechanism turned on. The light columns are for networks with the gas mechanism turned off. The figure illustrates two sets of 20 runs in each condition; the difference between the left and right-hand sets is the way in which the network connectivity was encoded. See text for further details.

It can clearly be seen in Figure 5 that controllers based on networks with the virtual gas diffusion and modulation mechanisms turned on evolve significantly faster than those that are identical in every respect (including genotypes and all the evolutionary machinery) except that the gas mechanisms are rendered inoperative. This result has been repeated under various different encoding schemes and for a wide range of mutation rates [23]. The clear implication is that GasNets are more evolvable – their search space is more amenable to the form of evolutionary search used – than the various other forms of network explored. Obviously this could be a potentially very useful property and it is looked at in more detail in the next section.

Nearly all the successful GasNet controllers that were examined in detail exhibited surprisingly simple structures (a typical example is shown in Figure 6) relying on a very small number of visual inputs, although their internal dynamics, supported by interwoven 'chemical' and 'electrical' processes, were often intricate [13]. A number of interesting sub-networks, such as oscillators making use of spatial aspects of the modulation and diffusion processes [13], were independently evolved in several runs, suggesting that they are easily found building blocks that the evolutionary process can make good use of.
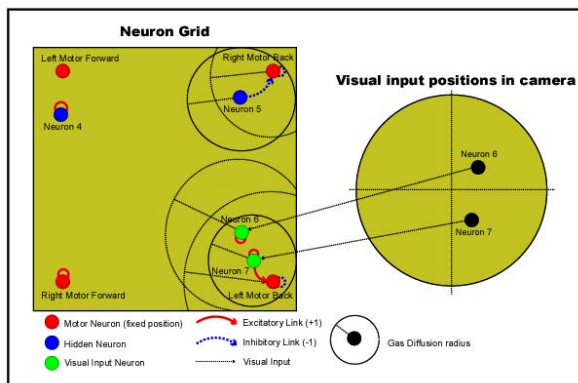
Fig. 6. A typical evolved GasNet controller for the triangle rectangle task illustrating the kind of structural simplicity often found in highly robust solutions. See text for further details.

6   Evolvability and Search Space Properties

The key result illustrated by Figure 5, that, for a particular evolutionary search algorithm, it is easier to find GasNet controllers for the triangle-rectangle task than non-GasNet controllers, tells us that there must be differences in the properties of the two search spaces. Understanding more about what this difference is may help us gain some valuable insights into the dynamics of artificial evolution and the nature of complex search spaces, as well as understanding more about the potential of GasNets. Smith et al. have published a number of papers working towards this goal [21,22,23]. The earliest studies in this series applied a whole range of standard search space 'complexity' and 'smoothness' metrics to the two spaces. These all failed to predict any difference between the spaces [22]. However, the research revealed a number of likely reasons for this: the spaces both appeared to be highly anisotropic, there is strong evidence for large neutral networks permeating the spaces, and a very large percentage of genotypes have negligible fitness. These and other properties combine to make the standard metrics useless for describing the pertinent differences between the spaces. Another (probably interrelated) reason is illustrated by the left-hand graphs in Figure 7 which show the median number of generations needed to reach a given fitness level for GasNets and non-GasNets. There is no difference in the two graphs for fitnesses of less than about 0.5. Fitnesses greater than this value are extremely unlikely to be found in random samples on which most of the basic metrics are based. Hence the focus of the work has shifted to the notion of evolvability and ways to measure it [23]. Evolvability is related to the ability of an individual or population to generate fit variants [24]. A useful measure of this is this transmission function which gives the distribution of all possible offspring for a given individual or

population [3]. In the work described here, a variety of mutation operators are used, but no crossover operator is employed. This has allowed transmission functions to be approximated through massive sampling of mutated populations saved from many evolutionary runs, so that many levels of initial fitness are represented. Smith has devised a number of evolvability metrics based on the transmission function [23]. The right-hand graphs in Figure 7 show plots of one of these measure (explained in the caption). The small, but significant, difference reveals that at higher fitnesses the GasNet space has fewer deleterious mutations than the non-GasNet space. This will lead to a larger variety of good solutions in the GasNet populations. While this is very likely not to be the whole story, and further investigations of search dynamics are underway, it must aid evolution. This discovery means the investigations are starting to bear fruit and we hope to soon have a full explanation of the differences in search performance on the two spaces.
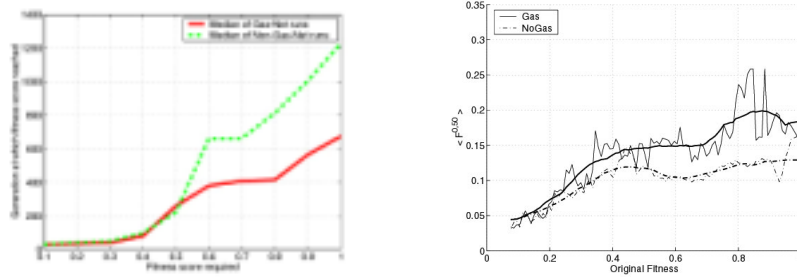
Fig. 7. The left-hand graphs show the median number of generations needed to reach a given fitness score for a large set of evolutionary runs. The right-hand graphs show a small, but significant, difference in the expected fitness of the bottom 50% of mutations applied to solutions of a given initial fitness. See text for further details.

7   Future Directions

There are many extensions to all aspects of the work described in this paper, some planned and some already underway. A number of these will be briefly discussed here.

As far as the computational modelling of the volume signalling roles of NO is concerned, obvious candidates for further work are: modelling diffusion from larger more complex structures and groups of structures, and introducing functional roles for the gas in detailed simulations of the behaviour of small neuronal networks. Both of these present non-trivial technical challenges and would require significant computing resources. However, given how much has been learnt from the studies carried out to date, as outlined in Section 3, it is important that such work is carried out.

The details of the 'electrical' and 'chemical' aspects of GasNets are, to some extent, rather arbitrary. There are numerous other forms that the modulations, node internal dynamics and virtual diffusion processes could take. Many of these are very worthwhile investigating in order to gain a deeper understanding of a whole class of systems. Such insights should usefully inform development in autonomous adaptive systems as well as systems level neuroscience. A number of interesting alternative modulation schemes being investigated include:

- Site specific modulations. The modulation depends on the presence of a 'receptor'. Virtual gases can trigger a range of modulations in a single network, depending on which receptors are present.

- Modulation of other adaptive processes (such as Hebbian synaptic changes). This could add a very powerful dimension to the evolution of adaptive mechanisms as advocated by Floreano et al. [17].

- Modulations at many different time scales, including permanent changes, are common in biology and are likely to play an important role in artificial systems.

Of course investigations into autonomous adaptive systems cannot focus solely on specific behaviour generating mechanisms. A bigger picture, involving overall architectures, body and sensor morphologies, developmental processes and a host of other issues surrounding embodied behaviour in the world, must be borne in mind.

The search space analyses outlined in the last section are part of an on-going investigation and in the future we wish to incorporate analyses of robot behaviours and their underlying controllers into the story.

8   Conclusions

The sciences of the real and the artificial have much to offer each other. This seems especially true in the study of adaptive behaviour and the mechanisms underlying it. This paper has outlined a multi-stranded interdisciplinary project that has striven to exploit the synergies at the interface between neuroscience and contemporary AI/Alife, and in so doing has advocated a shift towards rich electrochemical models and analogies. While the kind of creative exploratory work favoured by many Alife researchers are regarded as very important, and some of the work described here falls into that category, the authors stress the need for strong theoretical underpinnings and detailed analysis wherever possible. Without these, significant progress is unlikely.

References

1. Bredt DS and Snyder SH (1990) Isolation of nitric oxide synthetase, a calmodulin-requiring enzyme. *Proc Natl Acad Sci USA* **87**: 682-685.
2. 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.
3. Cavalli-Sforza, L. and Feldman, M. (1976). Evolution of continuous variation: Direct approaches through joint distribution of genotypes and phenotypes. *Proc. Nat. Academy of Sciences*, USA, **73**:1689-1692.
4. 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.
5. 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.
6. 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.
7. Grand, S. Creatures: An exercise in Creation, *IEEE Intelligent Systems magazine*, July/August 1997.
8. Hall ZW (1992) An Introduction to Molecular Neurobiology. Sinauer Associates Inc, Sunderland, Massachusetts.
9. 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.
10. Holscher, C. (1997) Nitric oxide, the enigmatic neuronal messenger: its role in synaptic plasticity. *Trends Neurosci.* **20**: 298-303.
11. Husbands, P. (1998) Evolving Robot Behaviours with Diffusing Gas Networks, In: P. Husbands and J.-A. Meyer (1998), 71-86.
12. P. Husbands and J.-A. Meyer (Eds) (1998*) EvoRobot98: Proceedings of 1st European Workshop on Evolutionary Robotics*, Springer-Verlag LNCS 1468.
13. P. Husbands and T. Smith and N. Jakobi and M. O'Shea. Better Living through Chemistry: Evolving GasNets for Robot Control, *Connection Science*, **10(3&4)**, 185-210, 1998.
14. Jakobi, N. (1998) Evolutionary Robotics and the Radical Envelope of Noise Hypothesis, *Adaptive Behavior*, **6(2)**: 325-368.
15. Kandel, E. (1976) *The cellular basis of behavior*. Freeman.
16. Katz B (1969) *The release of neural transmitter substances*. Liverpool University Press.
17. Nolfi, S. and Floreano, D. (2000). Evolutionary Robotics: The biology, intelligence and technology of self-organizing machines. MIT Press.
18. Park J-H, Straub V and O'Shea M (1998) Anterograde signaling by Nitric Oxide: characterization and in vitro reconstitution of an identified nitrergic synapse. *J Neurosci* **18.**
19. Philippedes, A. and P. Husbands and M. O'Shea. Four Dimensional Neuronal Signaling by Nitric Oxide: A Computational Analysis. *Journal of Neuroscience* **20(3)**: 1199--1207, 2000.

20. A. Philippedes and P. Husbands and T. Lovick and M. O'Shea (2001). Targeted gas clouds in the brain. (submitted)

21. T. Smith and P. Husbands and M. O'Shea (2001). Neutral Networks and Evovability with Complex genotype-Phenotype Mapping. *Proc. ECAL'01.* LNCS, Springer.

22. T. Smith and P. Husbands and M. O'Shea (2001). Not Measuring Evovability: Initial Investigations of an Evolutionary Robotics Search Space. In *Proc. CEC'01*, IEEE Press.

23. T. Smith and P. Husbands and M. O'Shea (2001). Evolvability, Neutrality and Search Difficulty. (submitted)

24. Wagner, G. and Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, **50(3)**:967-976.

25. 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.