

Flexible Couplings: Diffusing Neuromodulators and Adaptive Robotics

Andy Philippides², Phil Husbands², Tom Smith¹, and Michael O’Shea¹

Centre for Computational Neuroscience and Robotics (CCNR)

¹Department of Biology, ²Department of Informatics

University of Sussex, Brighton, UK

(andrewop|philh|toms|M.O-Shea)sussex.ac.uk

Abstract

Recent years have seen the discovery of freely diffusing gaseous neurotransmitters, such as nitric oxide (NO), in biological nervous systems. A type of artificial neural network (ANN) inspired by such gaseous signalling, the GasNet, has previously been shown to be more evolvable than traditional ANNs when used as artificial nervous system in an evolutionary robotics setting. Evolvability is used in the sense of consistent speed to very good solutions, here appropriate sensorimotor behaviour generating systems. We present two new versions of the GasNet which take further inspiration from the properties of neuronal gaseous signalling. The plexus model is inspired by the extraordinary NO producing cortical plexus structure of neural fibres and the properties of the diffusing NO signal it generates. The receptor model is inspired by the mediating action of neurotransmitter receptors. Both models are shown to significantly further improve evolvability. We describe a series of analyses suggesting that the reasons for the increase in evolvability is related to the flexible loose coupling of distinct signalling mechanisms, one ‘chemical’ and one ‘electrical’.

Keywords: GasNets, neuromodulation, evolvability, coupling, evolutionary robotics

Corresponding author: Phil Husbands

1 Introduction

This paper describes recent results from a two pronged research effort concerned with understanding more about the role of freely diffusing neurotransmitters in biological nervous systems and in exploring properties of artificial nervous systems based on principles abstracted from the biological research. The synthetic approach is embraced in both contexts. Artificial systems embodying the particular phenomena and principles in question are synthesized then analyzed post-hoc [4, 9, 7].

Both strands of the work focus on volume signalling, whereby neurotransmitters freely diffuse into a relatively large volume around a nerve cell, potentially affecting many other neurons [22, 54]. This exotic form of neural signalling does not sit easily with classical pictures of brain mechanisms and is forcing a radical re-think of existing theory [15, 24, 43, 26]. The first part of the work employs detailed computational and mathematical modelling to study how the dynamics of diffusion interact with neural morphology and influence potential functional roles of this form of signalling [43, 44]. The second strand, which is represented more heavily in this paper, looks at the efficacy of novel forms of artificial neural networks incorporating abstract analogs of volume signalling influenced by the detailed modelling work [28, 27]. Instances of these networks are synthesized, via an evolutionary search process, to generate behaviour in autonomous mobile robots. Shedding light on the reasons for a significant variation in evolvability of different forms of these networks is the central focus of this paper. Evolvability is used in the sense of consistent speed to very good solutions, here appropriate sensorimotor behaviour generating systems.

The class of ANNs we developed to explore analogues of volume signalling, we have termed GasNets. These are essentially a standard ANN augmented by a chemical signalling system comprising a diffusing *virtual* gas which can modulate the response of other neurons. The original GasNet has previously been shown to be more evolvable than traditional ANNs in the context outlined above [28]. In this paper we present two new versions of the GasNet which take further inspiration from the properties of neuronal gaseous signalling. The plexus model is inspired by the extraordinary NO producing cortical plexus structure of neural fibres and the properties of the diffusing NO signal it generates. The receptor model is inspired by the mediating action of neurotransmitter receptors. Both new GasNet variants are shown to exhibit significantly improved evolvability relative to the original. One of the most significant features of these two new types of ANN is the fact that the nature of the coupling between the ‘electrical’ and ‘chemical’ processes is more controllable and flexible than in the original GasNet. We later provide evidence, through a series of analyses described in Section 8, indicating that this has an

important influence on evolvability.

Since we develop artificial nervous systems for autonomous mobile robots engaged in sensorimotor behaviours, this work sits squarely in the area often referred to as new AI [11, 40] or embodied cognition [3, 13].

The next two sections provide biological background and motivation. These are followed by a brief description of work concerned with the detailed modelling of diffusing neurotransmitters in real nervous systems, which acts as explicit inspiration for the GasNets variants described in Section 5. Results from comparative evolutionary robotics experiments on the evolvability of various forms of GasNet are detailed in Section 6. These are followed by a series of analyses aimed at trying to explain the differences in evolvability highlighted in the robotics experiments. The paper finishes with conclusions.

2 From Neuroscience to Robotics

Modern Neuroscience has developed enormously since its beginnings in 19th century Physiology [21]. Its quest to provided us with an understanding of how nervous systems work, although very far from complete, has made great strides in the past few decades. Crucial to these later successes was the early twentieth century pioneering work of the likes of Adrian and Sherrington in elucidating key electrical properties and functions of nerve cells [10, 2, 47], work for which these two illustrious figures shared the 1932 Nobel prize for physiology. With the great benefit of hindsight, it is illuminating to step back to this time and consider the following quote from Adrian’s acceptance speech:

“The nerve fibre is clearly a signalling mechanism of limited scope. It can only transmit a succession of brief explosive waves, and the message can only be varied by changes in the frequency and in the total number of these waves. But this limitation is really a small matter, for in the body the nervous units do not act in isolation as they do in our experiments. A sensory stimulus will usually affect a number of receptor organs, and its result will depend on the composite message in many nerve fibres.” *Lord Adrian, Nobel Acceptance Speech, 1932 [1]*.

In this description we can clearly recognise the connectionist point-to-point electrical transmission picture that dominated thinking about the nervous system for several decades, and which was the original inspiration for artificial neural networks (ANNs) [36, 46] and is still a major focus of computational neuroscience today [16]. However, the past decade or so has

brought many fresh insights into the range of interacting signalling mechanisms underlying neural activity. These processes act over a wide range of both temporal and spatial scales and, as mentioned earlier, include such unexpected phenomena as chemical signalling, via diffusing neurotransmitters, between neurons that are not electrically connected [22]. Research is starting to reveal a far richer and more complex picture than was previously imagined [52, 5]. So although Adrian’s notion of understanding the ‘composite message’ is still highly relevant, the details of what might constitute such a message have significantly altered.

However, it is as yet very difficult to gather accurate empirical data for processes such as non-classical neurotransmitter diffusion [44]. Therefore it is natural to turn to computational modelling to shed light on volume signalling. As will be seen in Section 4, detailed models involving diffusing chemicals are computationally very expensive and at the moment have to be restricted to small numbers of neurons rather than whole behaviour generating neuronal circuits. Hence in order to further investigate functional roles involved in the generation of behaviour, we advocate the study of more abstract artificial robot nervous systems incorporating mechanisms based on volume signalling. These systems are computationally tractable and can generate sensorimotor behaviours in real time in the real world.

In order to provide background and context for the experimental work described later, the next section gives more details on volume signalling.

3 Volume Signalling: Beyond Connectionism

In the traditional connectionist model of neuron-to-neuron communication referred to above, neurons generate brief electrical signals (action potentials) which propagate along wire-like axons terminating at highly localized junctions (synapses) on other neurons where the release of a chemical signalling molecule, or neurotransmitter, is triggered. The neurotransmitter is confined to the region of the synapse and here the receiving neuron is equipped with receptors which directly translate the chemical signal into a brief electrical signal, either excitatory or inhibitory [45, 10]. Hence in standard ANNs based on this incomplete model, the notion of chemical signalling can be safely factored out, leaving only the idea of electrical signals flowing between nodes in a network.

But some receptors do not directly activate electrical events in the receiving neuron at all. They interact with messenger molecules that are best regarded as modulatory because among other things they regulate, or modulate, the actions of conventional transmitters. Modulatory neurotransmitters are ‘indirect’ because they can cause a range of 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 [45]. In this way modulatory transmitters greatly expand the diversity and the duration of actions mediated by the chemicals released by neurons. Thus while simple direct transmission between neurons certainly does exist, it operates in parallel with and sometimes conjointly with indirect chemical signalling systems that operate on an extended temporal scale [12].

The precise 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 [45, 30]. There are a great variety of receptors on different types of cells suggesting the possibility of a combinatorially explosive range of pairings and effects.

In addition to this, the spatial scale over which neurons can communicate is extended by the recent discovery of non-synaptic chemical signalling [50]. The most important feature of this derives from the ability of some neurotransmitters, in particular small gaseous molecules, to diffuse away from their site of release and to occupy a volume of the nervous system perhaps containing many other neurons and synapses [18]. To date three gaseous neurotransmitter molecules, *NO*, *CO* and *H₂S* have been identified, all of them, curiously, are highly poisonous. The most studied among them by far is *NO*, which is known to diffuse at above threshold concentrations many tens of microns away from a site of release [43]. Diffusion takes time; so this is not a rapid signalling system. Furthermore the main receptor for *NO* is of the indirect type and can have long-term modulatory effects on neurons [5, 50]. So not only can *NO* operate over a large region, it can also mediate extended temporal changes in the chemical and electrical properties of neurons within that volume.

Transmission by gases is not and perhaps cannot be confined to the highly localized region of the synapse, as in classical point-to-point signalling, loosening the tight coupling between electrical and chemical signals. Thus the concept of *volume signalling* can now be added to the growing list of phenomena in the nervous system that might be a source of inspiration for new and perhaps improved styles of ANNs. This is probably especially true for ANNs intended for use as artificial nervous systems, an area where taking inspiration from biology is often particularly fruitful.

Interneural chemical 'communication' systems of the volume signalling variety operate on

different temporal and spatial scales to electrical signalling. Thus an important principle we have used in the GasNets is the existence of two separate, yet interacting, signalling systems: one ‘electrical’ and one ‘chemical’.

Before moving on to the GasNet work, the next section covers elements of the detailed biological modelling work mentioned above. This provides further context and background to the GasNet studies, particularly the plexus GasNet.

4 Modelling NO diffusion in real brains

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 [54]. Another very important feature of NO signalling follows from the fact that nitric oxide synthase (responsible for the production of NO) 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 a continuous biologically realistic structures [43]. 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). Diffusion was modelled using accurate difference equation methods on a fine grid [44]. Equation 1 gives the modified diffusion equation approximated:

$$\frac{\partial C(\mathbf{x}, t)}{\partial t} - D\nabla^2 C(\mathbf{x}, t) = P(\mathbf{x}, t) - S(\mathbf{x}, t)C(\mathbf{x}, t) - \lambda C(\mathbf{x}, t) \quad (1)$$

where the terms on the left hand side describes general diffusion and those on the right hand side take into account NO production and depletion processes. Here $C(\mathbf{x}, t)$ is the concentration at point \mathbf{x} , D is the diffusion coefficient, $P(\mathbf{x}, t)$ is the concentration of NO produced per second at point \mathbf{x} . NO does not have a specific inactivating mechanism, and is lost through reaction with oxygen species and metals, as well as heme containing proteins [33, 51]. This means that the movement of other molecules and receptors and their interactions need not be modelled and instead a more general loss function can be used. Hence $S(\mathbf{x}, t)$ is a depletion function to model NO sinks (such as blood vessels with their very high concentrations of NO-binding hemes), and

λ is a general inactivation rate for all points outside sinks reflecting background oxidization and binding events. In this study $P = Q\rho$ for points inside a source during synthesis and zero elsewhere. Q is the amount of NO produced per second by a single NO producing ‘unit’ and ρ is the density of such units. For full details of modelling methods, all parameter values used, and their biological justification, see [43, 44].

Figure 1 illustrates the evolution of NO concentration during and after a 100ms burst of synthesis. Two very interesting observations 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 [43]. Such a phenomenon, with its possible functional implications, would not have been observed in a less accurate point-source type model [54], or indeed by purely empirical methods as they are as yet too crude.

Another very interesting NO signalling phenomenon amenable to computational modelling involves the study of the combined action of many small NO sources which individually cannot generate concentrations above the threshold level needed to have any effect. In particular, NO producing plexus structures, or meshes, of axonal fibres with diameters of a few microns or less are found in various parts of the brain. A knowledge of the characteristics of the NO signal created by such a structure could be crucial in helping to understand NO’s neuromodulatory role. For instance, while NO producing neurons account for only about 1% of cell bodies in the cerebral cortex, their processes spread so extensively, that almost every neuron in the cortex is exposed to these small fibres, the vast majority of which have diameters of a micron or less [5, 20]. Such structures are common across a range of species, another example is found in the locust optic lobe where there are highly ordered sets of NO producing $2\mu\text{m}$ diameter fibres [19]. Extending the techniques used for the study illustrated in Figure 1, Philippides et al. have modelled the diffusion of NO from multiple small sources including plexus type structures [44].

An illustration of how co-operation can occur in a model of fibres in the locust optic lobe is provided by figure 2 which shows the spatial extent of the NO signal generated by a single fibre of $2\mu\text{m}$ diameter and by ordered arrays of nine, sixteen, twenty five and thirty six identical fibre sources separated by $10\mu\text{m}$. An ordered array of N^2 fibres is an evenly spaced out 2D arrangement of $N \times N$ fibres. As the fibres are parallel the solution is symmetric along the z -axis (out of the page) and so we give results only for cross-sectional slices in a plane perpendicular to

the direction of the fibres. The single fibre does not achieve an above threshold signal principally because the great speed of NO diffusion means that NO will spread rapidly over a large volume. So while NO does not reach threshold anywhere, the volume occupied by NO at a significant fraction of threshold is large relative to the source size. Thus NO derived from small and well-separated individual sources can summate to produce an effective NO cloud. But what are the characteristics of such a signal and what do they imply for the way NO functions? The first thing to notice from Figure 2 is that while the ‘reservoir effect’ is still present, with NO accumulating in the centre of mass of the sources, the concentration profile appears to be flatter than for a single source. The effect of the spacing between the sources is further illustrated in Figure 3, where the area over threshold generated by 100 fibres is shown as a function of the total cross-sectional area of source. What is immediately obvious from this Figure is that if the goal is to reach the largest number of potential targets with a given volume of source fibres, then a dispersed source is better than a single source. Indeed, with multiple fine sources, it is possible to affect a volume over twice the size as for a solid source simply by dispersing them correctly.

Figure 3B reveals that the impact of the interaction effect will vary depending on the spacing used. In particular, the delay before the start of interaction can vary from nothing to more than a second, with the delay growing as the spacing is increased. This means that a system with optimal spacing, in terms of extent of the affected region, will experience a considerable delay before the region begins to be affected, with the total area affected suddenly rising sharply at the end of the delay. This raises the intriguing functional possibility of a system which is completely unaffected by NO for a given length of time (a period which would be tunable by changing the spacing), but once past this point, large regions are rapidly ‘turned-on’ by the cloud of NO.

Further computational studies have shown how a random mesh of fine fibres, as well as enabling similar signalling dynamics to the uniform meshes shown above, is an ideal structure to ensure a uniform concentration tightly targeted over the volume occupied by such a plexus [42, 44]. This is exactly the kind of structure found in the cortex, hence these models suggest a functional explanation for the extraordinary morphology of the plexus and point towards an important mechanism for allowing highly targeted NO ‘clouds’ in the brain. In highlighting the functional importance of brain morphology, these phenomena take us increasing further away from connectionist ideas and suggest that Pfeiffer’s notion of ecological balance, which requires a harmonious relationship between an agent’s morphology, materials and control [41], can perhaps be taken inside the head.

The next section describes a set of ANNs whose properties draw on those of real NO volume signalling systems, including some discovered using the computational modelling approach outlined in this section.

5 GasNets

The *GasNet* incorporates a mechanism based on the neuron-modulating properties of a diffusing signalling gas into a more standard sigmoid-unit neural network [28]. In previous work the networks have been used in a variety of evolutionary robotics tasks, including the one used in this paper, comparing the speeds of evolution for networks with and without the gas signalling mechanism active, showing that GasNets are consistently faster to evolve than more standard networks [28]. More recent work has demonstrated advantages of GasNets over other dynamical networks in bipedal and quadrupedal locomotion [38, 37], a quite different task from the one used in the study highlighted in this paper. A number of related studies have investigated the nature of the GasNet fitness landscapes [48] in order to elucidate the reasons for the faster evolutionary search. Other authors have used abstract notions of chemical modulation in neural networks used for controlling agents [32, 25] but on a more global level which does not involve the detailed spatiotemporal aspect we incorporate into our systems. In this section we introduce the basic GasNet model, and two further biologically motivated variants (plexus and receptor).

5.1 The basic GasNet

The ‘electrical’ network underlying the GasNet model is a discrete time step, recurrent neural network with a variable number of nodes. These nodes are connected by either excitatory (with a weight of +1) or inhibitory (with a weight of -1) links with the output O_i^n , of node i at time step n determined by a continuous mapping from the sum of its inputs, as described by the following equation:

$$O_i^n = \tanh \left[k_i^n \left(\sum_{j \in C_i} w_{ji} O_j^{n-1} + I_i^n \right) + b_i \right] \quad (2)$$

where C_i is the set of nodes with connections to node i and $w_{ji} = \pm 1$ is a connection weight. 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 , which can be altered at each time-step according to the concentration of the diffusing ‘gas’ at node i to give k_i^n (as described later in the section on modulation).

5.2 Gas diffusion in the networks

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. The networks occupy a 2D space; the diffusion processes mean that the relative positioning of nodes is crucial to the functioning of the network. The original GasNet diffusion model is controlled by two genetically specified parameters, namely the radius of influence r and the rate of build up and decay s . Spatially, the gas concentration varies as an inverse exponential of the distance from the emitting node with a spread governed by r , with the concentration set to zero for all distances greater than r (Equation 3 and Figure 4). The maximum concentration at the emitting node is 1.0 and the concentration builds up and decays from this value linearly as dictated by the time course function, $T(t)$, defined by Equations 4 and 5.

$$C(d, t) = \begin{cases} e^{-2d/r} \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (3)$$

$$T(t) = \begin{cases} H\left(\frac{t-t_e}{s}\right) & \text{emitting} \\ H\left[H\left(\frac{t_s-t_e}{s}\right) - H\left(\frac{t-t_s}{s}\right)\right] & \text{not emitting} \end{cases} \quad (4)$$

$$H(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x < 1 \\ 1 & \text{else} \end{cases} \quad (5)$$

where $C(d,t)$ is the concentration at a distance d from the emitting node at time t . t_e is the time at which emission was last turned on, t_s is the time at which emission was last turned off, and s (controlling the slope of the function T) is genetically determined for each node. The total concentration at a node is then determined by summing the contributions from all other emitting nodes (nodes are not affected by their own concentration, to avoid runaway positive feedback).

5.3 Modulation by the Gases

For mathematical convenience, in the basic GasNet there are two ‘gases’, one whose modulatory effect is to increase the transfer function gain parameter (k_i^n from equation 2) and one whose effect is to decrease it. It is genetically determined whether or not any given node will emit one of

these two gases (gas 1 and gas 2), and under what circumstances emission will occur (either when the ‘electrical’ activation of the node exceeds a threshold, or the concentration of a genetically determined gas in the vicinity of the node exceeds a threshold. Note these emission processes provide a coupling between the ‘electrical’ and ‘chemical’ mechanisms). The concentration-dependent modulation is described by Equation 6, with transfer parameters updated on every time step as the network runs.

$$k_i^n = k_i^0 + \alpha C_1^n - \beta C_2^n \quad (6)$$

where k_i^0 is the genetically set default value for k_i , C_1^n and C_2^n are the concentrations of gas 1 and gas 2 respectively at node i on time step n , and α and β are constants. Both gas concentrations lie in the range $[0, 1]$. Thus the gas does not alter the electrical activity in the network directly but rather acts by continuously changing the mapping between input and output for individual nodes, either directly or by stimulating the production of further virtual gas. The concentration dependent modulation can, for instance, change a node’s output from being positive to being zero or negative even though the input remains constant. Any node that is exposed to a non zero gas concentration will be modulated. This set of interacting processes provides the potential for highly plastic systems with rich dynamics. Typically many aspects of a functioning GasNet, such as gas concentrations at any point or the gain parameter of a given node, are in continuous flux as it generates behaviour in a mobile robot engaged in a sensorimotor task. The general form of the diffusion is based on the properties of a single source neuron as modelled in Section 4. The modulation chosen is motivated by what is known of NO modulatory effects at synapses [5].

5.4 Extensions to the basic GasNet I: The plexus model

In this section we introduce a GasNet variant, the plexus model, directly inspired by the type of signalling seen in the mammalian cerebral cortex as described earlier in Section 4. Recall that the NO signal is generated by the combined action of many fine NO producing fibres, producing a targeted ‘cloud’ which is distant from the neurons from which the fibre plexus emanates. Figure 2 shows how NO derived from small and well-separated individual sources can summate to produce an effective NO cloud.

This method of signalling has several interesting implications for the spatio-temporal nature of the ensuing volume signal, making it very different to a signal generated by a single neuron of the same size (the inspiration for the basic GasNet). Section 4 showed that the summation of NO

from several separated fibres gives a smoothed flat concentration distribution, with the above threshold concentration part extending over a wider area. Figure 3 illustrates the phenomenon whereby the whole volume affected by the cloud suddenly goes above threshold and is ‘turned on’.

What, though, if anything, do these features bring to evolved artificial nervous systems? In an attempt to answer this question we developed the plexus GasNet, whose diffusion properties are modified so as to produce an abstraction of the type of signal produced by plexus type structures. Firstly, we changed the spatial distribution of gas concentration. In the original GasNet this was modelled as an exponentially decaying function (Equation 3) which is loosely based on the type of spatial distribution of NO one would see outside a single neuron (Figure 4B). For the plexus model this has been modified to a uniform distribution over the volume of affect (Equation 7, $T(t)$ refers to the function in equation 4), with a peak concentration half that of the original (illustrated in Figure 4A):

$$C(d, t) = \begin{cases} 0.5 \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (7)$$

The second change is to allow the centre of this gas diffusion cloud to be distant from the controlling node (which, by analogy, is the source of the ‘plexus’). Note that this model requires two extra parameters for the gas diffusion centre (x, y) coordinates. Thus the plexus model produces constant concentration within the area of effect, with this area centred anywhere in the space (Figure 5). All other details of the models are identical to the original GasNet model, as described earlier.

5.5 Extensions to the basic GasNet II: The receptor model

An aspect of biological neuronal networks that has no analog in the vast majority of ANNs is the role of receptor molecules. All neural signalling is mediated by a diverse group of proteins which act as receptors to which neurotransmitters bind. The act of binding triggers chemical processes which result in functional changes to the neuron involved [12, 45]. In classical synaptic neurotransmission two basic classes of receptors have been identified: ionotropic and metabotropic. Ionotropic receptors are linked directly to ion channels in the postsynaptic membrane. These channels are opened or closed in response to transmitter binding, thus changing the postsynaptic membrane potential and hence mediating the postsynaptic electrical response. This type of receptor is generally involved in rapid timescale effects acting over milliseconds. Metabotropic receptors are not directly linked to ion channels but affect them by the activation of intermedi-

ate G-proteins [45]. G-proteins can interact directly with ion channels or with effector enzymes that give rise to intracellular second messengers that lead to complex biochemical signalling cascades, most of which are as yet poorly understood. Hence they can give rise to a wide range of modulatory effects that act over timescales ranging from seconds to hours or even months and years. The picture is significantly complicated by the fact that a single transmitter can activate both classes of receptors at a single site. As has already been stated, non-classical transmitters, such as NO, are not confined to act at localized synaptic sites, but diffuse freely. Accordingly, NO receptors are not membrane associated and can have a wide spatial distribution, and can be found anywhere in the nerve cell. NO triggers a variety of modulations through second messenger pathways that have the potential to interact in even more complex ways because of the spatially extended aspect of its action.

Although neuroscience is a long way from a full understanding of receptor mechanisms, especially those involved in indirect modulation by second messenger intracellular pathways, there are a number of powerful systems level ideas we can abstract and incorporate into our ANNs. This we have done with the second new GasNet variant: the receptor model, again taking inspiration directly from contemporary neuroscience .

Details are similar to the basic GasNet except there is now *only one* virtual gas and each node in the network can have one of three discrete quantities (zero, medium, maximum) of N possible receptors. Each diffusing neurotransmitter receptor pairing gives rise to a separate modulation to the properties of the node. The strength of a modulation at node i at time n , ΔM_j^n , is proportional to the product of the gas concentration at the node, C_i^n and the relevant receptor quantity, R_j as described by equation 8. Each modulation makes some change to one or more function parameters of the node. All the variables controlling the process are again set for each node by an evolutionary search algorithm.

$$\Delta M_j^n = \rho_i C_i^n R_j \tag{8}$$

In the original GasNet any node that was in the path of a diffusing transmitter would be modulated in a fixed way. The receptor model allows site specific modulations, including no modulation (zero quantity of receptors) and multiple modulations at a single site. This provides a powerful context ‘switching’ mechanism that pulls the ‘chemical’ and ‘electrical’ processes further apart, allowing (but not forcing) looser coupling, while further increasing the potential for complex network dynamics. A number of different receptor linked modulations have been experimented with, including:

- Action of receptor1: increase gain of node transfer function as in original GasNet
- Action of receptor2: decrease gain of node transfer function as in original GasNet
- Action of receptor3: increase proportion of retained node activation from last time step (i.e. treat node as a leaky integrator in which the time constant can be modulated upwards)
- Action of receptor4: if above a threshold switch transfer function of node for sustained period (i.e. fundamentally change transfer function, for instance from a sigmoid to a constant output or a periodic function)

Note the first two modulation are immediate and short-lived while the last two operate over a longer time-scale. Each possible subset of these receptors proved to be at least as evolvable (in terms of speed to a very good solution) as the original GasNet, while some were significantly better. A variant that proved particularly successful used receptor1 only. This is the model that will be referred to as the Receptor GasNet in the following sections on the comparative studies of the evolvability of different types of GasNet.

6 Comparative Experiments

Although most of the GasNet variants described in this paper have been successfully used in a number of robotic tasks, their evolvability and other properties were thoroughly *compared* on a robotic visual discrimination task. Starting from an arbitrary position and orientation in a black-walled arena, a robot equipped with a forward facing camera must navigate under extremely variable lighting conditions to one shape (a white triangle) while ignoring the second shape (a white square). The relative position of the shapes varied from trial to trial, as did their size within a variation of 10% in both dimensions. This task has been used before in detailed comparisons between the basic GasNet and other styles of ANN [28], hence it is appropriate to use it in this comparison between GasNet variants. Because of the noise and variation, and limited sensory capabilities, this task is quite challenging in this setup. Both the robot control network, one or other form of GasNet, and the robot sensor input morphology, i.e. the number and position of the input pixels on the visual array, were under evolutionary control. This is illustrated in figure 6. Fitness over a single trial was taken as the fraction of the starting distance moved towards the triangle by the end of the trial period, and the evaluated fitness was returned as the weighted sum of 16 trials of the controller from different initial conditions:

$$F = \frac{2}{N(N+1)} \sum_{i=1}^{i=N} i \left(1 - \frac{D_i^F}{D_i^S}\right) \quad (9)$$

where D_i^F is the distance to the triangle at the end of the i th trial, and D_i^S the distance to the triangle at the start of the trial, and the N trials are sorted in descending order of $\frac{D_i^F}{D_i^S}$. Thus good trials, in which the controller moves some way towards the triangle, receive a smaller weighting than bad trials, encouraging robust behaviour on all 16 trials. Success in the task was taken as an evaluated fitness of 1.0 over thirty successive generations of the evolutionary algorithm. Evaluations took place using a validated minimal simulation of the robot as described in [29]. Controllers developed in the simulation successfully transferred to reality generating behaviours in the actual physical robot at least as well as in the simulation. The noisy lighting conditions, varying positions and sizes of the shapes, and other properties of the simulation meant that highly robust solutions developed, generalized to the variations experienced during evolution. For further information on the task and robot see [28].

6.1 The Evolutionary Search Algorithm

A geographically distributed asynchronous updating evolutionary algorithm was used, with a population size of 100 arranged on a 10×10 grid. Parents were chosen through rank-based roulette-wheel selection on the mating pool consisting of the 8 nearest neighbours to a randomly chosen grid-point. A mutated copy of the parent was placed back in the mating pool using inverse rank-based roulette-wheel selection. In what follows, a ‘generation’ in such an algorithm occurs every 100 reproduction events. For full details see [28].

6.2 The Solution Representation and Mutation Operators

The robot controllers were encoded as a variable length string of integers, with each integer allowed to lie in the range $[0, 99]$. Each node in the network was coded for by a number of parameters controlling such properties as node positions on the 2D grid in which GasNets operate, ‘electrical’ connectivity, whether or not the node has sensor input, and all gas diffusion and modulation variables. Connections were formed as in figure 8, with each node connecting to nodes lying within one of two connection segments. Hence each genotype consists of N blocks of P integers coding for the properties of the nodes, where N is the number of nodes in a particular network (this varies from genotype to genotype) and P is the number of parameters describing a node (19 for the original GasNet, 21 for the plexus and receptor versions).

Three mutation operators were applied to solutions during evolution. Each *integer* in the string had a probability (4%) of mutation in a Gaussian distribution around its current value. There was also an addition operator, with a 4% chance per *genotype* of adding one neuron to the network by inserting a block of random values describing each of new node’s properties. Finally there was a deletion operator, with a 4% chance per *genotype* of deleting one randomly chosen neuron from the network. In the next section, we describe the speed of evolution results for the three models.

7 Speed of evolution results

Table 1 shows the speed of evolution results for the three GasNet variants. Forty runs were carried out for each version, with runs being terminated once controllers were evolved that achieved 100% fitness over thirty consecutive generations. In all cases successful networks have typically 8-15 nodes with 2-4 visual inputs. GasNet controllers are very lean, in terms of numbers of nodes and connections, in comparison with previously evolved solutions using other styles of ANN [28]. Here we see that the plexus and receptor models evolve good solutions significantly faster than the original GasNet model. The receptor model gives particularly dramatic improvements. During the course of a typical evolutionary run, the original and plexus GasNet fitnesses tended to rise in jumps following periods of stasis, whereas the receptor GasNet fitness tended to rise steadily to a perfect score of 1.0. The question is therefore, what is it about the new features of these two models that mediates this increase? In the remainder of the paper we will explore this question through an analysis of the coupling between the chemical and electrical signalling mechanisms in the networks.

8 Why the difference: flexible loose coupling?

In earlier studies the original GasNet was compared to various styles of recurrent ANNs that did not possess complex intrinsic dynamics at the node level. The much greater evolvability of the GasNets might have been related to the richer internal dynamics of the networks with multiple time scales in operation. However, the significant variation in the evolvability of the 3 GasNet variants studied here, with highly comparable internal dynamics at the node level, suggests that the story is more complex. This is also backed up by recent detailed comparisons of many forms of recurrent networks with rich dynamics, including GasNets and CTRNNs [8], applied to bipedal and quadrupedal locomotion [38, 37], a task that would appear to be more

dynamically demanding than the one studied in this paper. The original GasNet variant was found to be either more evolvable or at least more reliable at finding good solutions quickly, than the other forms of dynamical networks used. Hence the provision of rich internal dynamics is only part of the answer. As mentioned earlier, one of the most significant features of GasNets is the operation of two distinct, yet interacting, signalling mechanisms. An important aspect of the two new variants explored in this paper is the fact that the nature of the coupling between the ‘electrical’ and ‘chemical’ processes is more controllable and flexible than in the original GasNet. We believe this has an important role to play in evolvability.

There is some evidence from evolutionary theory that the degree of coupling between interacting and yet distinct processes might lie at the heart of some important principles for the development of complex systems. To evolve successfully, an organism must satisfy the conflicting pressures of *phenotypic stability* and *genetic instability*, i.e. that the organism be robust to phenotypic change (to not fall off the current adaptive peak), and amenable to genotypic change (to allow movement to a new adaptive peak). (author?) [14] identifies genetic redundancy and multiple weak interaction as possible mechanisms by which these two conflicting pressures can be satisfied.

Such loosely coupled redundant systems contain the potential for genotypic change without phenotypic change; both multiple weak interactions and redundancy allow for gradual, or even neutral, transformation of function through genetic variation [14]. In such systems, phenotypic fitness is likely to be highly correlated across the genotype landscape, either (or both) through significant levels of neutrality, and low levels of ruggedness. Such systems are also robust to phenotypic change; complex systems picked at random are more likely to be stable if the system is characterized by either multiply connected weakly interacting components, or sparsely connected strongly interacting components [23, 35].

By contrast, strongly coupled non-redundant systems are far less amenable to variation; change in one component is more likely to affect the entire system, leading to phenotypic instability. We see this effect clearly in the theoretical NK fitness landscapes, where a higher degree of epistatic connection between the components leads to a less correlated fitness landscape [31]. In other words, even small changes in the genotype in a strongly coupled system lead to large changes in the phenotype. However, in tenably neutral versions of the NK landscapes, high degrees of redundancy compensate in some measure for the strong coupling, allowing genetic variation without massive phenotypic variation [6, 39].

We hypothesize that systems involving distinct yet coupled processes are highly evolvable

when the coupling is flexible (i.e. it is relatively easy for evolution to change the degree of coupling in the system) with a bias towards a loose coupling; this allows the possibility of ‘tuning’ one process against the other without destructive interference. Some evidence that GasNets are this kind of system can be found in **(author?)** [49] where it is shown that GasNets can re-evolve to deal with a changed timescale much faster than more standard ‘electrical connection only’ networks. Here we present further evidence of a different kind.

As described earlier, the plexus model allows network nodes to emit gas from anywhere in the grid. This partly separates the gas diffusion and the electrical synaptic activity mechanisms; synaptic connections are formed from the current node position, while gas diffusion connections are formed from the gas emission position. Thus gas connections in the grid can be changed through modifying the gas emission position, while synaptic connections can be altered through moving the node itself¹. Similarly, the addition of receptors to mediate the modulatory affects of the virtual gas potentially allows even more independence between electrical and chemical signalling. This is because, as explained in Section 5.5, the receptor GasNet used in the comparative experiments had only one type of receptor; its presence or absence at a node in the network essentially acts as a switching mechanism turning on and off modulation at the node.

There is no simple way of calculating the degree of coupling in the three forms of network; in principle one can measure the degree of ruggedness through correlation lengths or similar methods. However, **(author?)** [48] shows that these types of measures do not discriminate well between highly heterogenous problem spaces such as those found here. In this Section, we introduce a number of simple, and admittedly incomplete, measures of the degree of coupling between the gas diffusion and electrical synapse mechanisms.

The first of these involves calculating the two connectivity matrices (electrical and chemical) for a given successful GasNet, and computing the coupling as the number of overlapping connections, i.e the number of elements which are non-zero in both connectivity matrices. Table 2 shows this coupling between the electrical and gas diffusion processes for the three models. The number of electrical synaptic connections and number of gas diffusion connections (averaged per neuron), and the percentage of overlapping connections, are shown for each GasNet variant. The values shown were calculated by averaging over the best evolved controllers from each of the successful runs. Three points can be made. First, there are no significant differences between the numbers of electrical synaptic connections across the three models. Second, the percentages of overlapping connections in the GasNet are significantly higher than those in the receptor and

¹Note that the two mechanisms are not entirely separated: both act on the actual position of the destination nodes.

plexus models; thus indicating that coupling between the electrical and diffusion processes is far stronger in the GasNet model than in the receptor and plexus models and may provide part of the reason for the faster evolutionary search. Third, the number of diffusion connections is significantly lower in the receptor GasNets. This last fact suggests that the receptor mechanism is being used to shape leaner networks with fewer, more specific, gas connections.

However, this method of measuring the coupling of the two signalling systems does not take account of the actual action of the network. For instance, a gas or electrical connection is registered regardless of whether the neuron emitted gas or was electrically active, respectively. To remedy this, we examined the *on-line* coupling by measuring the overlap between the electrical and gaseous activity matrices (denoted by E and G respectively) at each time-step during a run of the robot. These matrices are generated by setting G_{ij} (or E_{ij}) = 1 if the amount of gas or total electrical activity (i.e. excitatory activity minus inhibitory activity, including self-recurrency) at node i due to node j at a given time-step t , is non-zero. G_{ij} (or E_{ij}) = 0 otherwise. In addition, for the gas matrix only, the value G_{ij} is set to zero if there is no electrical activity at node j at that time-step, since, in its absence, the gas would be ineffectual in terms of influencing electrical activity at the node. For the receptor GasNet there must also be a non-zero quantity of receptors associated with node j for G_{ij} to be set to one. The amount of overlap between the two matrices at each time-step, i.e. the number of corresponding non-zero elements, is a measure of the on-line coupling between the electrical and chemical mechanisms. This value is averaged over all time-steps during 16 runs of the robot with random starting conditions to give the on-line coupling value for a given network.

As can be seen from Table 3, where the measure is compared for the three GasNet variants, the on-line coupling gives a lower value to that generated from the static connectivity matrices (Table 2). This measure is lower for the plexus GasNet than for the original model, but higher for the receptor model. This may indicate that the receptor GasNets, with fewer gas connections, have more highly tuned architectures where the connections are used more frequently, giving rise to the higher value. The best that can be said is that this measure is rather inconclusive.

While the on-line connectivity coupling measure described above gives an indication of the functional coupling of a network, we have hypothesized that flexible loose coupling is useful in an evolutionary context. To investigate this further we must, therefore, look at the mutants of evolved networks rather than just the networks themselves. In particular, if we are to suppose that a loose coupling between the two systems is beneficial, then we might expect that the distribution of fitnesses of mutations which affected either chemical or electrical signalling systems

independently would be different. This is indeed the case as can be seen in Figures 9 and 10 where we have compared, for 10 successful GasNets of each type, the fitnesses of all possible one-point mutants in which either the gas connectivity or electrical connectivity matrices have been changed independently, or both matrices have changed. Mutants were created by generating each possible single-point mutation of the genomes of the successful GasNets. What is clear is that for each style of network, mutations which change the gaseous connections only have a superior average fitness (Figure 9) and a distribution most biased towards high fitnesses (Figure 10), while those mutants that change both types of connectivity are most detrimental, indicating destructive interference between the two mechanisms.

Given that the *properties* of mutants which affect the signalling systems independently differ between the three network search spaces, in light of the difference in evolutionary speeds, the natural question to ask next is whether the *distribution* of these independent mutations is different for the spaces. Examining the average number of mutations as a fraction of the total number of mutations for all the one-point mutants which changed either one or other or both connectivity matrices, we do see differences, as shown in Figure 11. While both the total proportion of mutants which affect connectivity matrices at all, and the proportion affecting electrical connections only, are very similar for all three spaces, the proportion affecting the gaseous signalling system only (which have higher average fitness) is significantly greater for the plexus and receptor GasNet spaces *relative to* the proportion of destructive mutations which affect both systems. In the original GasNet the proportion of these two types of mutants is the same. This means that the receptor and plexus variants are less likely to suffer deleterious mutations. When we add to this the fact that receptor GasNet ‘gas only’ mutants have a higher average fitness than their plexus GasNet counterparts (Figure 9), and have their distribution of fitnesses biased more heavily towards high scores (Figure 10), we can start to see why, under the guidance of a selection mechanism that favours fitter individuals, the receptor GasNets are the most evolvable. Selection is able to exploit the greater proportion of fitter (or at least less deleterious) mutants.

These measure all point towards the benefit of systems with loosely coupled distinct processes that can be (flexibly) ‘tuned’ against each other. Here the picture emerging is of tuning the gas mechanisms against the electrical networks which can be done more easily in a non-destructive beneficial way in the less coupled plexus and receptor GasNets.

9 Conclusions

The results on degree of coupling and speed of evolution presented in this paper support our view that systems involving distinct yet flexibly coupled processes are highly evolvable when there is a bias towards a loose coupling between the processes; this allows the possibility of evolution ‘tuning’ one against the other without destructive interference. The receptor model, in which the search process arguably has the most direct control over the degree of coupling, is seen to be by far the most evolvable in terms of consistent speed to very good solutions. Indeed, preliminary experiments on extended versions of the shape discrimination task, involving more shapes and two stage discriminations, have proved highly successful with the receptor model. The one-point mutant studies only really address evolvability at high fitnesses, since the mutants were generated from already successful individuals. A fuller (and very time consuming!) study will repeat this kind of analysis at intervals throughout the entire evolutionary history.

This paper marks a first step in our attempts to gain deeper insights into the importance, or otherwise, of the coupling issue. As well as exploring the use of our artificial nervous systems for generating more complex behaviours, we are trying to build a formal framework to extend our theoretical understandings. In this we have some considerable distance to go. Although we have concentrated on changes (to plastic systems) over an evolutionary timescale, very similar issues are likely to be important at the timescale of the plastic changes themselves.

Of course there are many other forms of plasticity in real neuronal networks and many other potential principles to be abstracted [17, 34], but we believe that explicitly dealing with the electrochemical nature of nervous systems is likely to be an increasingly fruitful area of research that will likely force us to broaden our notions of what behaviour generating mechanisms might look like [53].

Acknowledgements: The authors would like to thank Inman Harvey, Ezequiel Di Paolo, Seth Bullock, Michael Wheeler and all the members of the CCNR (<http://www.cogs.susx.ac.uk/ccnr/>) for constructive discussion. They also thank the anonymous reviewers for their constructive comments on an earlier draft of this paper.

References

- [1] *Nobel Lectures: Physiology or Medicine 1922-1941*. Elsevier, 1965.
- [2] E.D. Adrian. *The Basis of Sensation*. Christophers, London, 1928.

- [3] Michael L. Anderson. Embodied cognition: A field guide. *Artificial Intelligence*, 149(1):91–130, 2003.
- [4] W. Ross Ashby. *Design for a brain*. Chapman and Hall, London, 1952.
- [5] D. Barañano, C. Ferris, and S. Snyder. Atypical neural messengers. *Trends in Neuroscience*, 24(2):99–106, 2001.
- [6] L. Barnett. Ruggedness and neutrality: The NKp family of fitness landscapes. In C. Adami, R.K. Belew, H. Kitano, and C.E. Taylor, editors, *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*, pages 18–27. MIT Press, Cambridge, Massachusetts, 1998.
- [7] M. Bedau. Four puzzles about life. *Artificial Life*, 4:125–140, 1998.
- [8] R.D. Beer. A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72:173–215, 1995.
- [9] M.A. Boden, editor. *The Philosophy of Artificial Life*. OUP, Oxford, 1996.
- [10] M.A.B Brazier. *A History of the Electrical Activity of the Brain*. Pitman, London, 1961.
- [11] R. A. Brooks. *Cambrian Intelligence: The early history of the new AI*. MIT Press, 1999.
- [12] J.-P Changeux. Chemical signalling in the brain. *Sci. Am.*, 269(5):58–62, 1993.
- [13] Andy Clark. *Being There*. MIT Press, Cambridge, MA, 1996.
- [14] M. Conrad. The geometry of evolution. *BioSystems*, 24:61–81, 1990.
- [15] T.M Dawson and S.N. Snyder. Gases as biological messengers: nitric oxide and carbon monoxide in the brain. *Journal of Neuroscience*, 14(9):5147–5159, 1994.
- [16] Peter Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, Mass., 2001.
- [17] E. DiPaolo. Evolving spike-timing-dependent plasticity for single-trial learning in robots. *Phil. Trans. R. Soc. Lond. A*, 361:2299–2319, 2003.
- [18] G.M. Edelman and J.A. Gally. Nitric oxide: Linking space and time in the brain. *Proceedings of the National Academy of Sciences, USA*, 89:11651–11652, 1992.

- [19] M.R. Elphick, L. Williams, and M. O'Shea. New features of the locust optic lobe: evidence of a role for nitric oxide in insect vision. *Journal of Experimental Biology*, 199:2395–2407, 1996.
- [20] C. Estrada and J. DeFilipe. Nitric oxide-producing neurons in the neocortex: Morphological and functional relationship with intraparenchymal microvasculature. *Cerebral Cortex*, 8:193–203, 1998.
- [21] S. Finger. *Origins of Neuroscience: A History of Explorations into Brain Function*. OUP, 2001.
- [22] J.A. Gally, P.R. Montague, G.N. Reeke, and G.M. Edelman. The NO hypothesis: Possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. *Proceedings of the National Academy of Sciences, USA*, 87:3547–3551, 1990.
- [23] M.R. Gardner and W.R. Ashby. Connectance of large dynamic (cybernetic) systems: Critical values for stability. *Nature*, 228:784, 1970.
- [24] J. Garthwaite and C.L. Boulton. Nitric oxide signaling in the central nervous system. *Annu Rev Physiol*, 57:683–706, 1995.
- [25] S. Grand. Creatures: An exercise in creation. *IEEE Intelligent Systems Magazine*, 12(4):4–6, 1997.
- [26] C. Hölscher. Nitric Oxide, the enigmatic neuronal messenger: Its role in synaptic plasticity. *Trends in Neurosciences*, 20:298–303, 1997.
- [27] P. Husbands, A. Philippides, T. Smith, and M. O'Shea. Volume signalling in real and robot nervous systems. *Theory in Biosciences*, 120:251–268, 2001.
- [28] P. Husbands, T. Smith, N. Jakobi, and M. O'Shea. Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 10(4):185–210, 1998.
- [29] N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behaviour*, 6:325–368, 1998.
- [30] B. Katz. *The Release of Neural Transmitter Substances*. Liverpool University Press, Liverpool, UK, 1969.

- [31] S.A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Oxford, UK, 1993.
- [32] T. Kondo, A. Ishiguro, Y. Uchikawa, and P. Eggenberger. Autonomous robot control by a neural network with dynamically-rearranging function. In *Fourth International Symposium on Artificial Life and Robotics: AROB99*, pages 324–329. Oita, Japan, 1999.
- [33] J. Lancaster. Diffusion of free nitric oxide. *Methods in Enzymology*, 268:31–50, 1996.
- [34] B. Mathayomchan and R.D. Beer. Center-crossing recurrent neural networks for the evolution of rhythmic behavior. *Neural Computation*, 14:2043–2051, 2002.
- [35] R.M. May. Will a large complex system be stable? *Nature*, 238:413–414, 1972.
- [36] W.S. McCulloch and W.H. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133, 1943.
- [37] G. McHale and P. Husbands. Gasnets and other evolvable neural networks applied to bipedal locomotion. In S. Schaal, editor, *Proc. From Animals to Animats 8: Proceedings of the Eighth International Conference on Simulation of Adaptive Behaviour (SAB'2004)*, page in press. MIT Press, 2004.
- [38] G. McHale and P. Husbands. Quadrupedal locomotion: Gasnets, ctrnns and hybrid ctrnn/pnns compared. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. Watson, editors, *Proc. 9th Int. Conference on the Simulation and Synthesis of Living Systems (Alife IX)*, page in press. MIT Press, 2004.
- [39] M.E.J. Newman and R. Engelhardt. Effects of selective neutrality on the evolution of molecular species. *Proceedings of the Royal Society of London, B*, 265:1333–1338, 1998.
- [40] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*. MIT Press, Cambridge, Massachusetts, 2000.
- [41] Rolf Pfeiffer and Christian Scheier. *Understanding Intelligence*. MIT Press, Cambridge, Mass., 1999.
- [42] A. Philippides. *Modelling the diffusion of nitric oxide in brains*. PhD thesis, University of Sussex, Brighton, UK, 2001.
- [43] A.O. Philippides, P. Husbands, and M. O’Shea. Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *Journal of Neuroscience*, 20(3):1199–1207, 2000.

- [44] A.O. Philippides, P. Husbands, T. Smith, and M. O’Shea. Structure based models of NO diffusion in the nervous system. In J. Feng, editor, *Computational Neuroscience: a Comprehensive Approach*, pages 97–130. Chapman and Hall/CRC Press, 2003.
- [45] D. Purves. *Neuroscience*. Sinauer, 1997.
- [46] D.E. Rumelhart, J.L. McClelland, and the PDP Research Group, editors. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations*. MIT Press, Cambridge, Massachusetts, 1986.
- [47] C. Sherrington. *Man on his Nature*. OUP, Oxford, UK, 1940.
- [48] T.M.C. Smith, P. Husbands, and M. O’Shea. Local evolvability, neutrality, and search difficulty in evolutionary robotics. *Biosystems*, 69:223–243, 2003.
- [49] T.M.C. Smith, P. Husbands, A. Philippides, and M. O’Shea. Neuronal plasticity and temporal adaptivity: Gasnet robot control networks. *Adaptive Behavior*, 10(3&4):161–184, 2002.
- [50] S. Snyder and C. Ferris. Novel neurotransmitters and their neuropsychiatric relevance. *American Journal of Psychiatry*, 157:1738–1751, 2000.
- [51] M. Vaughn, L. Kuo, and J. Liao. Estimation of nitric oxide production and reaction rates in tissue by use of a mathematical model. *Am J Physiol*, pages 2163–2176, 1998.
- [52] S. R. Vincent. Nitric oxide: a radical neurotransmitter in the central nervous system. *Progress in Neurobiology*, 42:129–160, 1994.
- [53] M. Wheeler. Explaining the evolved: Homunculi, modules, and internal representation. In P. Husbands and J.-A. Meyer, editors, *Evolutionary Robotics: First European Workshop, EvoRobot98*, pages 87–107. Springer, Berlin, April 1998.
- [54] J. Wood and J. Garthwaite. Models of the diffusional spread of nitric oxide: Implications for neural nitric oxide signalling and its pharmacological properties. *Neuropharmacology*, 33:1235–1244, 1994.

	Original	Plexus	Receptor
Num Runs	40	40	40
Mean (sd)	3042 (3681)	1579 (2609)	82(102)
Median	1201	512	47
Best	136	101	13
Worst	> 10000	> 10000	512

Table 1: Number of generations before consistent success is achieved for the three models described earlier. Results for both plexus and receptor GasNets were significantly better than for the original GasNet. In particular, the receptor variant achieves success an order of magnitude faster. NB runs not achieving consistent success by generation 10000 were terminated.

	Original	Plexus	Receptor
Num successful runs	33	37	40
Synaptic connections (sd)	1.89 (0.52)	1.72 (0.41)	1.85 (0.42)
Diffusion connections (sd)	2.27 (0.93)	2.78 (0.84)	1.02 (0.56)
Overlapping connection coupling (sd)	40.5% (13.2%)	10.8% (8.1%)	12.32% (5.26%)

Table 2: A simple measure of coupling in the original GasNet, plexus and receptor models. For each of the successfully evolved controllers, the number of electrical synaptic connections and number of gas diffusion connections are shown (averaged per neuron). The percentage of connections which overlap, i.e. that connect the same neurons, are also shown. Standard deviations in brackets. See text for further details.

	Original	Plexus	Receptor
Num successful runs	33	37	40
Synaptic connections (sd)	1.89 (0.52)	1.72 (0.41)	1.85 (0.42)
Diffusion connections (sd)	2.27 (0.93)	2.78 (0.84)	1.02 (0.56)
on-line connection coupling (sd)	5% (7.5%)	1.3% (3.1%)	6.12% (5.39%)

Table 3: A measure of on-line coupling in the original GasNet and plexus models. For each of the successfully evolved controllers, the number of electrical synaptic connections and number of gas diffusion connections are shown (averaged per neuron). The percentage of on-line connections which overlap, are also shown. Standard deviations in brackets. See text for further details.

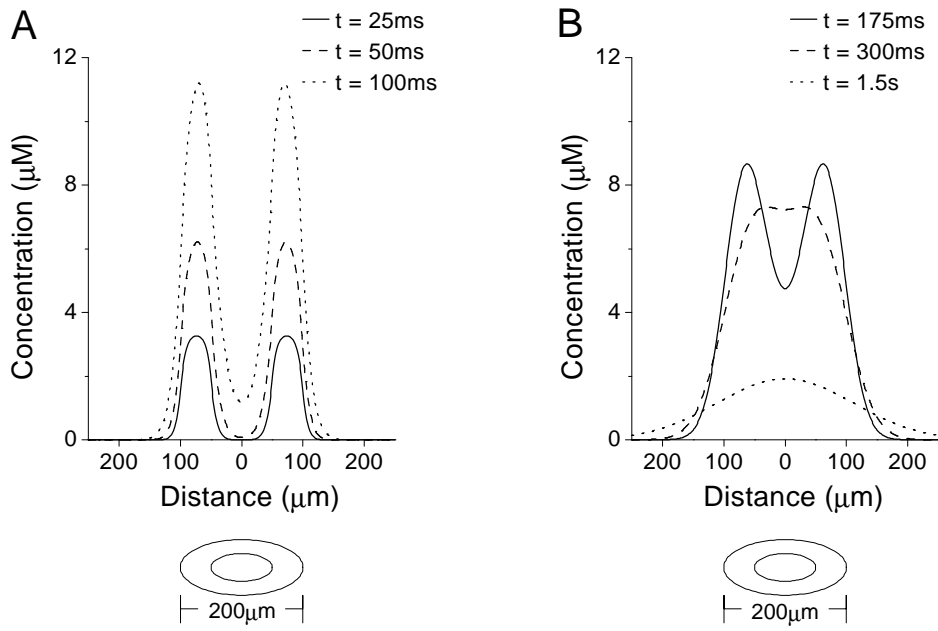


Figure 1: Concentration of NO plotted against distance from the centre of a hollow spherical source of inner radius $50\mu\text{m}$ and outer radius $100\mu\text{m}$ for a 100ms burst of synthesis starting at time $t = 0$. The graphics underneath each plot depict the structure. A. Concentration of NO at times $t = 25, 50$ and 100ms , two time points during and one at the end of synthesis. B. Concentration of NO after synthesis at times $t = 175, 300$ and 1.5s . The reservoir effect following the end of synthesis is clearly seen as the centrally accumulated NO is trapped by the higher surrounding concentrations. See test for further details.




Figure 2: Volume over threshold ($0.25\mu\text{M}$) generated by NO synthesizing fibres of $2\mu\text{m}$ diameter organized in ordered arrays separated by $10\mu\text{m}$ after 1 second of synthesis. The fibre dimensions and spacing have been chosen so as to approximate the arrangement of the nNOS-expressing fibres in the optic lobe of the locust (Elphick et al. 1996). The upper graph shows the volume over threshold per unit length of the fibres. The lower four graphs show the concentrations of NO (dark = low, light = high) in a two-dimensional slice through the fibres which project out of the page. Here we see how NO from several sources can combine to produce above threshold concentrations (areas inside the white boundaries) which extend away from the synthesizing region. Scale bar is $50\mu\text{m}$.

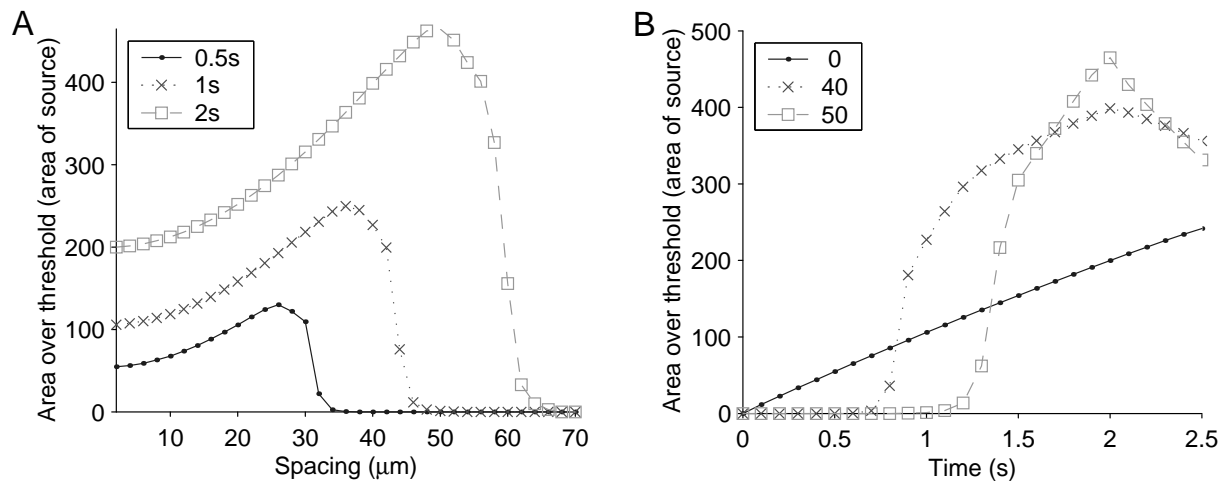


Figure 3: Area over threshold as a function of the total cross-sectional area of source for different numbers of evenly spaced fibres of diameter $2\mu\text{m}$. A. Affected area against spacing for 100 fibres for NO synthesis of length 0.5, 1 and 2 seconds. B. Affected area over time due to 100 fibres arranged as a single source (spacing = 2) or separated by 40 or $50\mu\text{m}$ for 2 seconds of NO synthesis. Note the delay till effective co-operation of the separated sources.

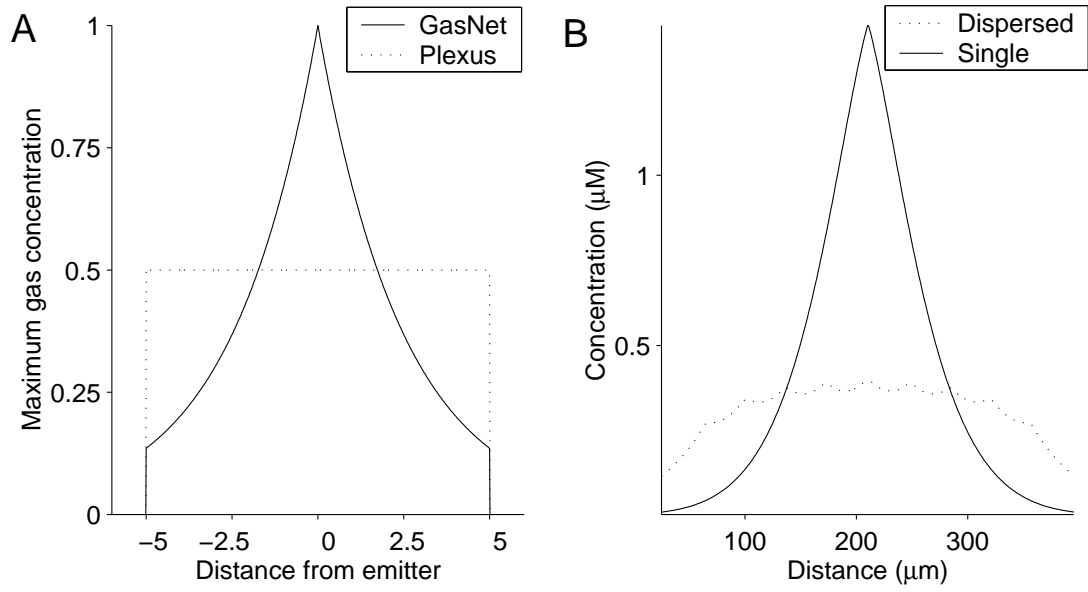


Figure 4: A. The spatial distributions of gas concentration for the different GasNet models. The solid line denotes the spatial distribution for the GasNet model, while the dotted line shows the spatial distribution for the plexus model. Units are the (internally consistent) ones used in the implementation. See text for further details. B. The spatial distributions of gas concentration outside the emitting (real) neuron for a single source (solid line) and dispersed sources (dotted line).

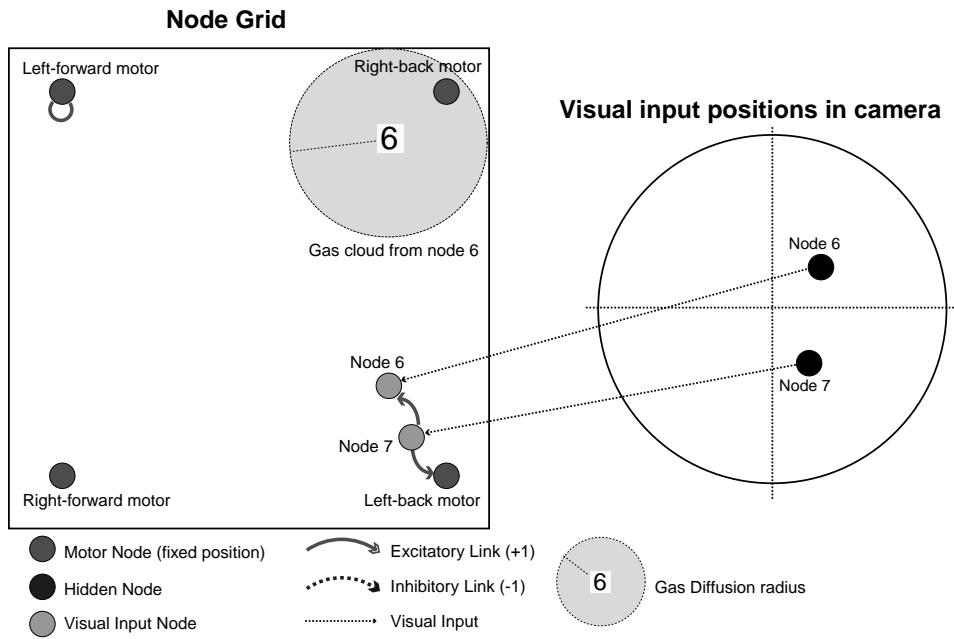


Figure 5: An example plexus architecture network. The node plane is shown on the left, with the positions and connections of the network nodes, while the camera on the right shows the position of the visual inputs (see section 6). Node 6 has a dispersed gas cloud centre in the top-right of the network plane, with a uniform concentration over the area of effect, illustrating the effects of the plexus model. Thus node 6 can easily affect nodes which are far from its position in the node plane.




Figure 6: The simulated arena and robot. The left-hand view (A) shows the robot position in the arena with the triangle and square. Fitness is evaluated on how close the robot approaches the triangle. The right-hand view (B) shows what the robot ‘sees’, along with the pixel positions selected by evolution for visual input. A validated simulation of the robot shown in Figure 7 was used.




Figure 7: The gantry robot used in this study. A CCD camera head moves at the end of a gantry arm allowing full 3D movement. In this study 2D movement was used, equivalent to a wheeled robot with a fixed forward pointing camera. A validated simulation was used: controllers developed in the simulation work at least as well on the real robot.

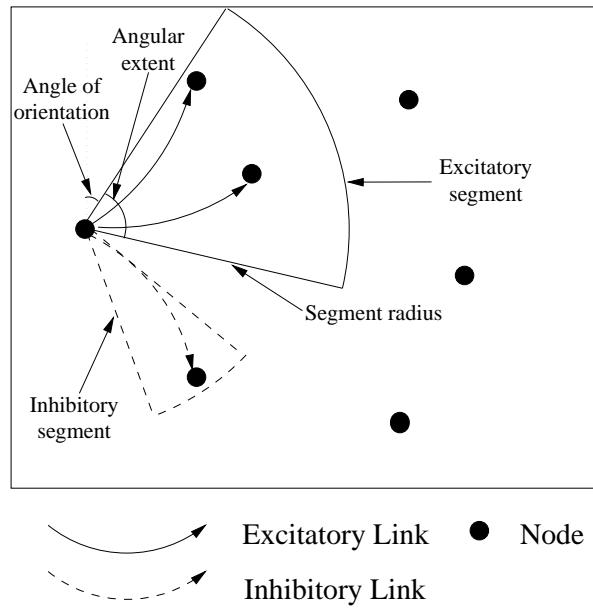


Figure 8: The connectivity of the network is defined by positive (excitatory) and negative (inhibitory) segments emanating from the nodes. Segment orientations, width and radii are controlled by genetically determined parameters as illustrated in the diagram. Networks develop and function on a 2D plane. Negative connections are made to any node situated in an inhibitory segment, positive connections to those situated in an excitatory segment. See text for further details.




Figure 9: Mean fitnesses of one-point mutants of 10 evolved Original, Plexus and Receptor GasNets. For all styles of network, all possible one-point mutants of the original successful evolved genotypes were exhaustively generated. The mutants were then divided into 3 distinct groups: those whose electrical connectivity matrices are altered by the mutation, relative to the original genotype, but whose gaseous connectivity matrices are unchanged (labelled ‘Electrical only’); those whose gaseous connectivity matrices are altered but whose electrical connectivity matrices are unchanged (labelled ‘Gas only’); those in which both electrical and gaseous connectivity matrices are altered (labelled ‘Both’). Mutants where there is no change in either connectivity matrix are ignored for the purposes of this analysis. Electrical connectivity matrices are generated in the standard way; elements of gaseous connectivity matrices G_{ij} are generated by calculating the maximum concentration that could be experienced at node j due to node i . Each mutant network was then evaluated 5 times to account for noise and the mean mutant fitnesses for the sub-groups of each evolved network calculated. The figure shows the means over the 10 networks, with error bars showing standard deviations. Note the repeated pattern across all network variants: high mean fitness of ‘Gas only’ mutants and the low mean fitness of mutants where both connectivity matrices are changed.

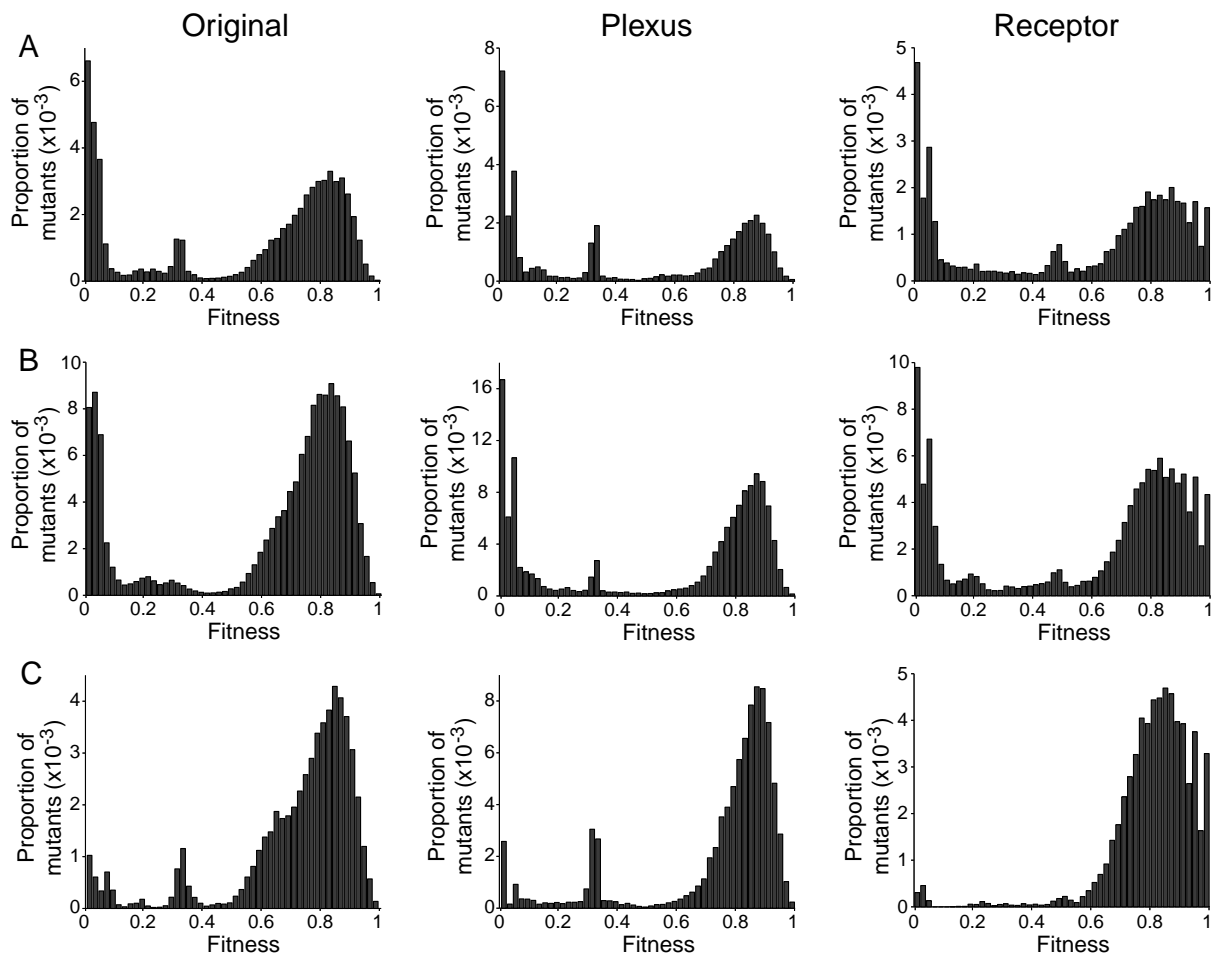


Figure 10: Distributions of fitnesses of one-point mutants of 10 evolved Original, Plexus and Receptor GasNets. Each row compares distributions for different mutant sub-classes of the three styles of network. (A) Distribution of fitnesses of mutants in which the mutation altered **both** gaseous and electrical connectivity matrices. (B) Distribution of fitnesses of mutants in which the mutation altered **electrical** connectivity matrices only. (C) Distribution of fitnesses of mutants in which the mutation altered **gaseous** connectivity matrices only. See figure 9 caption for more details. Distributions shown as proportions of total number of one-point mutants. For all styles of network, mutations which affect gas connectivity only have a distribution which has a lighter tail (less low fitnesses) than the other groups. This is particularly marked in the receptor case where there are practically no mutants with a fitness less than 0.6.




Figure 11: Mean proportions of all one-point mutants of 10 evolved Original, Plexus and Receptor GasNets where the mutation has altered either the mutant’s electrical connectivity matrix only (‘Electrical only’), its gaseous connectivity matrix only (‘Gas only’), or both gaseous and electrical connectivity matrices (‘Both’). See Figure 9 caption for more details. Error bars show standard deviations. Here we see that plexus and receptor networks have relatively more of the ‘Gas only’ mutants, which have a higher mean fitness than the other groups, than the low mean fitness ‘Both’ mutants. The original GasNet has the same relative proportion of each. See text for further details.