

Fast and Loose: Biologically Inspired Couplings

Andy Philippides¹, Phil Husbands², Tom Smith¹, and Michael O’Shea¹

Centre for Computational Neuroscience and Robotics (CCNR)

¹School of Biological Sciences, ²School of Cognitive and Computing Sciences

University of Sussex, Brighton, UK

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

Abstract

Recent years have seen the discovery of gaseous transmitters in biological nervous systems. An ANN inspired by such gaseous signalling, the GasNet, has previously been shown to be more evolvable than traditional ANNs. Here we present 2 new versions of the GasNet which take further inspiration from the properties of gaseous signalling. The plexus model is inspired by the cortical nNOS plexus and the properties of the 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 preliminary results suggesting that the reasons for the increase in evolvability is the loose coupling of distinct signalling mechanisms. Issues surrounding the degree of coupling between these mechanisms, one ‘chemical’ and one ‘electrical’, are discussed.

Introduction

Beyond Connectionism

A connectionist model of neuron-to-neuron communication involving a tight coupling between electrical and chemical signalling has provided the basic biological inspiration for many forms of artificial neural networks (ANNs). This has led to ANNs that have been successfully used in many applications, including as artificial nervous systems generating behaviour in autonomous robots. But the model represents an incomplete view of neuronal signalling which we know today operates over a very wide range of temporal and spatial scales. It might therefore be productive to augment standard ANNs to reflect this. But just how does the current view of neuronal signalling differ from the one that inspired conventional ANNs?

In the traditional model, neurons generate brief electrical signals (action potentials) which propagate along wire-like axons terminating at highly localised 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. 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. In this type of indirect chemical signalling, a released neurotransmitter initiates long-lasting changes in excitability, thereby modulating a neuron’s subsequent response to other electrical signals. 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 (Changeux 1993).

In addition to this, the spatial scale over which neurons can communicate is extended by the recent discovery of non-synaptic chemical signalling (Snyder & Ferris 2000). 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 (Edelman & Gally 1992). 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 (Philippides, Husbands, & O’Shea 2000). 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 (Barañano, Ferris, & Snyder 2001; Snyder & Ferris 2000). So not only can *NO* operate over a large region, it can also mediate long-lasting 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 localised 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.

In our work we have attempted to incorporate into ANNs, in an abstracted form, some of the richness and complexity that characterises the temporal and spatial dynamics of real neuronal signalling, especially chemical signalling by gaseous transmitters. As these systems operate on different temporal and spatial scales to electrical signalling, it suggests that it might be useful to develop models in which electrical and indirect chemical signalling are controlled in ANNs by separate processes. Thus, we developed the GasNet, a standard ANN augmented by a diffusing gas which can modulate the response of other neurons. We have used the methods of evolutionary robotics to explore the suitability of this class of networks for generating a range of behaviours in a variety of autonomous robots (Husbands *et al.* 1998; 2001). Since as yet we have no deep formal theory for such systems, we have found the use of stochastic search methods (such as evolutionary algorithms) to be a very helpful tool in this exploration. Of course our ultimate aim is to gain a better understanding of autonomous behaviour generating mechanisms in real and artificial systems.

While the GasNet has been seen to be more evolvable (in terms of speed of evolution) on a number of robotics tasks (Smith, Husbands, & O’Shea 2002), it is a first version using a very abstract model of non-classical signalling. We therefore wanted to see if adding in more features of gaseous signalling could further improve the evolvability. To do this, however, one must first overcome a problem familiar to anyone attempting to simplify biological systems: which of the panoply of complex features of neuronal signalling seen does one take? Does one simply use intuition about what may be useful or are there perhaps some guiding principles that one should keep in mind? We decided to take inspiration from the type of gaseous signalling seen in the cerebral cortex and from receptor mechanisms to generate two new versions of the GasNet, the *plexus* and *receptor* models. We present the new models in the next section along with details of the original GasNet, and then go on to show that they significantly further improve the evolvability of the GasNet.

Coupling of signalling mechanisms

The question of why these particular new models worked while others failed, remains, and may have implications for the types of network we should be trying to develop in the future. One answer for this might be found in an analysis of the connectivity patterns of successfully

evolved controllers. In the original GasNet model, the genotype to phenotype mapping meant that the coupling between ‘electrical’ and ‘chemical’ signalling was tighter than would be seen in the biological systems that inspired it. That is, if neurons were electrically linked to other neurons they tended to be chemically linked to those neurons as well. Here (section 4) we show that the two new models reduce this level of coupling significantly. Is this perhaps a guiding principle in generating future models, enabling us to reduce the vast space of features from which inspiration can be taken? When distinct processes involved in the functioning of the nervous system are loosely coupled it may be much easier for evolutionary mechanisms to effectively ‘tune’ one against the other without destructive interference (Gardner & Ashby 1970). This will help in satisfying the conflicting pressures of phenotypic stability and genetic instability needed for successful evolution (Conrad 1990). These issues are discussed further in the final section of the paper. In the next section we describe the GasNet, plexus and receptor models.

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 (Husbands *et al.* 1998). In previous work the networks have been used in a variety of evolutionary robotics tasks, 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 (Husbands *et al.* 1998). A number of related studies have investigated the nature of the GasNet fitness landscapes (Smith, Husbands, & O’Shea 2001; 2002) 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 (Kondo *et al.* 1999; Grand 1997) 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 the two new variants (plexus and receptor) together with details of their biological inspiration.

The GasNet model

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:

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

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).

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 2 and figure 3). The maximum concentration at the emitting node is 1.0 and the concentration builds up and decays from this value linearly as defined by equations (equation 3 and 4) at a rate determined by s .

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

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

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

where $C(d, t)$ 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).

Modulation by the Gases

For mathematical convenience there are two ‘gases’, one whose modulatory effect is to increase the transfer function gain parameter 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 5, with transfer parameters updated on every time step as the network runs.

$$k_i^n = k_i^0 + \alpha C_1^n - \beta C_2^n \quad (5)$$

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 and can, for instance, change the output from being positive to being zero or negative even though the input remains constant.

Extensions to the basic GasNet I: The plexus model

Philippides (2001) introduces a number of variants to the original GasNet model, based on research into the diffusion of NO in real brains (Philippides, Husbands, & O’Shea 2000). In this section we investigate one such variant, the plexus model, directly inspired by the type of signalling seen in the mammalian cerebral cortex. Here activity in a neuron is translated via a plexus of exceedingly fine nNOS-expressing fibres into a volume signal in a different part of the network.

Nerve fibres constituting the nNOS plexus in the cerebral cortex have a maximum diameter of about 5µm but the overwhelming majority are a fraction of a micron in diameter. They are therefore well below the critical size required for a volume signal, and cannot produce significant concentrations of NO unless many small sources effectively combine their production. An illustration of how such co-operation can occur is provided by figure 1 which shows the spatial extent of the NO signal generated by a single fibre of 2µm diameter and by arrays of four, nine and sixteen identical sources separated by 10µm. 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




Figure 1: A. Volume of tissue above threshold per unit length of fibre for 4 fibres of diameter $2\mu\text{m}$ spaced $10\mu\text{m}$ apart for NO synthesis of length 2s plotted against time after 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. B-C. Graphical representations of the NO concentration (dark = low, light = high) due to 1 (B) and 4 (C) fibres at times $T = 0.125, 0.5, 0.75, 1, 1.25$ and 2s after synthesis showing the build up of NO over time. The black line on each of the plots shows the $0.25\mu\text{M}$ (the threshold concentration) contour and thus indicates which regions can be affected by NO. Scalebar = $10\mu\text{m}$. Here 4 fibres can achieve an effective volume signal over the whole synthesising region but the single fibre cannot. In addition, note that inside the synthesising region the concentration is relatively even.

derived from small and well-separated individual sources can summate to produce an effective NO cloud.

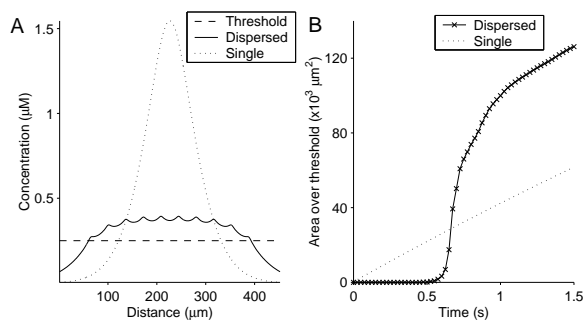


Figure 2: NO concentrations generated by a 10×10 ordered array of 100 NO synthesising fibres of diameter $2\mu\text{m}$ after 1s of NO synthesis. A. Concentration of NO due to 100 fibres separated by $36\mu\text{m}$ (solid line) and $0\mu\text{m}$, i.e. arranged as one single source, (dotted line) plotted against distance from the centre of the array. The dashed line shows the threshold concentration. B. Area over threshold due to 100 fibres separated by $36\mu\text{m}$ (solid line) and $0\mu\text{m}$ (dotted line) plotted against time after synthesis.

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 gener-

ated by a single neuron of the same size (figure 2). The summation of NO from several separated fibres means that the concentration in and around them is, in a sense, averaged and hence smoothed. Thus due to the dynamics of diffusion one tends to get a relatively even concentration within the synthesising region with small peaks around the fibres themselves (figure 1 and figure 2 A), resulting in a much greater region being above threshold (figure 2 B). In conjunction with the use of a threshold concentration, this means that there will come a point when the concentration in a region around the fibres is just sub-threshold and a small increase in the general level of NO will result in large areas rising above threshold. Thus we see a delay before a steep rise in the volume affected which is characteristic of signalling by dispersed sources.

What, though, if anything, can these features do for evolutionary robotics? In an attempt to answer this question we developed the plexus model, a variant of the GasNet, whose diffusion properties are modified so as to produce an abstraction of the type of signal seen in the cortex. Firstly, we changed the spatial distribution of gas concentration. In the original GasNet this was modelled as an exponentially decaying function (equation 2) which is loosely based on the type of spatial distribution

of NO one would see outside a single neuron (figure 3). For the plexus model this has been modified to a uniform distribution over the volume of effect, with a peak concentration half that of the original (illustrated in figure 3):

$$C(d, t) = \begin{cases} 0.5 \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (6)$$

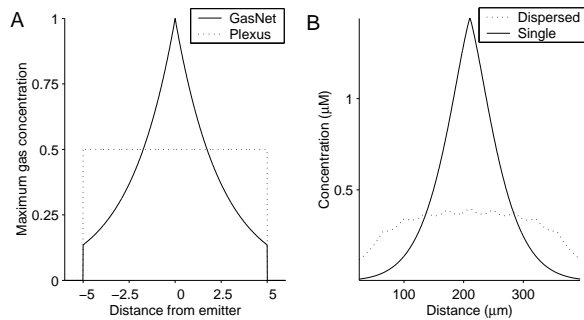


Figure 3: A. The spatial distributions of gas concentration for the different GasNet models. The solid line denotes the spatial distribution for the GasNet model, while the dotted line shows the spatial distribution for the plexus model. See text for further details. B. The spatial distributions of gas concentration outside the emitting (real) neuron for a single source (solid line) and dispersed sources (dotted line).

The second change is to allow the centre of this gas diffusion cloud to lie anywhere within the space, not just at the emitting node position. 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 the area of effect, with this area centred anywhere in the space (figure 4). All other details of the models are identical to the original GasNet model, as described earlier.

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 (Changeux 1993; Purves 1997). 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




Figure 4: 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. 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.

ion channels but affect them by the activation of intermediate G-proteins (Purves 1997). 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, been 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 7. 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.

$$\Delta M_j^n = \rho_i C_i^n R_j \quad (7)$$

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
- Action of recetor4: if above a threshold switch transfer function of node for sustained period

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 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.

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). Both the robot control network, one or other form of GasNet, and the robot sensor input morphology, i.e. the position of the input pixels on the visual array, were under evolutionary control. 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:

$$F = \frac{2}{N(N+1)} \sum_{i=1}^{i=N} i \left(1 - \frac{D_i^F}{D_i^S}\right) \quad (8)$$

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. For further information on the task and robot see (Husbands *et al.* 1998).

The Evolutionary Search Algorithm

A geographically distributed asynchronous updating evolutionary algorithm was used, with a *PopSize* 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. For full details see (Husbands *et al.* 1998).

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 nineteen parameters (21 for the receptor and plexus models), controlling such properties as node positions, connectivity, sensor input, and all gas diffusion and modulation variables. Connections were formed as in figure 5, with each node connecting to nodes lying within one of two connection arcs.




Figure 5: The connectivity of the network is defined by positive and negative arcs ($T=\theta$, $Tw=\theta_{width}$). Networks develop and function on a 2D plane. See text for further details.

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. 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.

Speed of evolution results

Table 1 shows the speed of evolution results for the three GasNet variants. Forty runs were carried out with each model, with runs being terminated once controllers were evolved that achieved 100% fitness over thirty consecutive generations. Here we see that the plexus and receptor models evolve good solutions significantly faster (T-test analyses were carried out to confirm this) than the GasNet model. The receptor model gives particularly dramatic improvements. 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 gas diffusion and electrical synaptic mechanisms in the networks.

Coupling

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 , 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, Smith *et al.* (2001) shows that these types of measures do not discriminate between highly heterogenous problem spaces such as those found here. In this section, we introduce a very simple, and admittedly incomplete, description

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

to measure the degree of coupling between the gas diffusion and electrical synapse mechanisms. We calculate the two connectivity matrices (electrical and chemical) for a given successful GasNet, and calculate 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 of the successfully evolved controllers, over all models. Two 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 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. In the final section we discuss this hypothesis further.

Discussion

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). Conrad (1990) 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 (Conrad 1990). 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 characterised by either multiply connected weakly interacting components, or sparsely connected strongly interacting components (Gardner & Ashby 1970; May 1972).

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

	Original	Plexus	Receptor
Num Runs	40	40	40
Mean (S.D.)	3042 (3681)	1579 (2609)	260 (161)
Median	1201	512	158
Best	136	101	46
Worst	> 10000	> 10000	840

Table 1: Number of generations before consistent success is achieved, for the three models described earlier. Both plexus and receptor models results were significantly better than the original GasNet results. NB runs not achieving consistent success by generation 10000 were terminated.

	Original	Plexus	Receptor
Num successful runs	33	37	40
Synaptic connections (S.D.)	1.89 (0.52)	1.72 (0.41)	1.61 (0.3)
Diffusion connections (S.D.)	2.27 (0.93)	2.78 (0.84)	2.1 (0.7)
Overlapping connection coupling (S.D.)	40.5% (13.2%)	10.8% (8.1%)	11.4 (5.5%)

Table 2: 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 connections which overlap, i.e. that connect the same neurons, are also shown. See text for further details.

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 (Kauffman 1993). 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 (Barnett 1998; Newman & Engelhardt 1998).

The results on degree of coupling and speed of evolution presented in this paper support our view that systems involving distinct yet coupled processes are highly evolvable when there is a bias towards a loose coupling between the processes; this allows the possibility of ‘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. 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.

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 them-

selves.

Acknowledgements: The authors would like to thank all the members of the CCNR (<http://www.cogs.susx.ac.uk/ccnr/>) for constructive discussion, as well as the anonymous referees. We would also like to thank the Sussex High Performance Computing Initiative (<http://www.hpc.sussex.ac.uk/>) for computing support. TS and AP are funded by BT Exact sponsored CASE studentships (<http://innovate.bt.com/projects/ftg.htm>).

References

- Barañano, D.; Ferris, C.; and Snyder, S. 2001. Atypical neural messengers. *Trends in Neuroscience* 24(2):99–106.
- Barnett, L. 1998. Ruggedness and neutrality: The NKp family of fitness landscapes. In Adami, C.; Belew, R.; Kitano, H.; and Taylor, C., eds., *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*. Cambridge, Massachusetts: MIT Press. 18–27.
- Changeux, J.-P. 1993. Chemical signalling in the brain. *Sci. Am.* 269(5):58–62.
- Conrad, M. 1990. The geometry of evolution. *BioSystems* 24:61–81.
- Edelman, G., and Gally, J. 1992. Nitric oxide: Linking space and time in the brain. *Proceedings of the National Academy of Sciences, USA* 89:11651–11652.
- Gardner, M., and Ashby, W. 1970. Connectance of large dynamic (cybernetic) systems: Critical values for stability. *Nature* 228:784.
- Grand, S. 1997. Creatures: An exercise in creation. *IEEE Intelligent Systems Magazine* July.
- Husbands, P.; Smith, T.; Jakobi, N.; and O’Shea, M.

1998. Better living through chemistry: Evolving Gas-Nets for robot control. *Connection Science* 10(3-4):185–210.
- Husbands, P.; Philippides, A.; Smith, T.; and O’Shea, M. 2001. Volume signalling in real and robot nervous systems. *Theory in Biosciences* 120:251–268.
- Kauffman, S. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford, UK: Oxford University Press.
- Kondo, T.; Ishiguro, A.; Uchikawa, Y.; and Eggenberger, P. 1999. Autonomous robot control by a neural network with dynamically-rearranging function. In *Fourth International Symposium on Artificial Life and Robotics: AROB99*.
- May, R. 1972. Will a large complex system be stable? *Nature* 238:413–414.
- Newman, M., and Engelhardt, R. 1998. Effects of selective neutrality on the evolution of molecular species. *Proceedings of the Royal Society of London, B* 265:1333–1338.
- Philippides, A. O.; Husbands, P.; and O’Shea, M. 2000. Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *J. Neurosci.* 20(3):1199–1207.
- Philippides, A. 2001. *Modelling the Diffusion of Nitric Oxide in Brains*. Ph.D. Dissertation, U. Sussex.
- Purves, D. 1997. *Neuroscience*. Sinauer.
- Smith, T.; Husbands, P.; and O’Shea, M. 2001. Not measuring evolvability: Initial exploration of an evolutionary robotics search space. In *Proceedings of the 2001 Congress on Evolutionary Computation: CEC2001*. Piscataway, New Jersey: IEEE Press. 9–16.
- Smith, T.; Husbands, P.; and O’Shea, M. 2002. Local evolvability, neutrality, and search difficulty in evolutionary robotics. *Biosystems*. In press.
- Snyder, S., and Ferris, C. 2000. Novel neurotransmitters and their neuropsychiatric relevance. *American Journal of Psychiatry* 157:1738–1751.