Spatial, Temporal, and Modulatory Factors Affecting GasNet Evolvability in a Visually Guided Robotics Task

PHILIP HUSBANDS,[1] ANDREW PHILIPPIDES,[1] PATRICIA VARGAS,[2] CHRISTOPHER L. BUCKLEY,[1] PETER FINE,[1] EZEQUIEL DI PAOLO,[1] AND MICHAEL O'SHEA[1]

[1] School of Informatics and Department of Biology, Centre for Computational Neuroscience and Robotics (CCNR), University of Sussex, Brighton BN1 9QJ, United Kingdom and [2] School of Maths and Computer Science, Heriot-Watt University, Edinburgh

Received 6 May 2010; accepted 27 June 2010

Spatial, temporal, and modulatory factors affecting the evolvability of GasNets — a style of artificial neural network incorporating an analogue of volume signalling — are investigated. The focus of the article is a comparative study of variants of the GasNet, implementing various spatial, temporal, and modulatory constraints, used as control systems in an evolutionary robotics task involving visual discrimination. The results of the study are discussed in the context of related research.

Key Words: GasNets; neuromodulation; evolvability; volume signalling; evolutionary robotics

1. INTRODUCTION

This article describes investigations into the evolvability of a style of artificial neural network strongly inspired by those parts of contemporary neuroscience that emphasize the complex electrochemical nature of biological nervous systems. In particular, they make use of an analogue of volume signalling, whereby gaseous neurotransmitters freely diffuse into a relatively large volume around a nerve cell, potentially affecting many other neurons irrespective of whether or not they are electrically connected [1]. This exotic form of neural signalling does not sit easily with classical connectionist (point-to-point) pictures of brain mechanisms and is forcing a radical re-think of existing theory [2–7].

The class of artificial neural networks developed to explore artificial volume signalling are known as GasNets [8]. They comprise a fairly standard artificial neural network augmented by a chemical signalling system based on a diffusing virtual gas, which can modulate the response of other neurons. A number of GasNet variants, inspired by different aspects of real nervous systems, have been explored as artificial nervous systems in an evolutionary robotics context. They were introduced to explore their potential as robust "minimal" systems for controlling behavior in very noisy

Correspondence to: Philip Husbands, School of Informatics, University of Sussex, Brighton BN1 9QJ, United Kingdom (e-mail: philh@sussex.ac.uk)

DOI 10.1002/cplx.20336
Published online in Wiley Online Library (wileyonlinelibrary.com)

environments. They have been shown to be significantly more evolvable, in terms of speed to a good solution, than other forms of neural networks for a variety of robot tasks and behaviors [8–12]. In two interrelated strands of work, they are being investigated as potentially useful engineering tools, while more detailed modelling work is aimed at gaining helpful insights into biological systems [3, 5].

Reasons for the enhanced evolvability and performance of these networks are explored. In particular, the roles of the spatiotemporal properties of the modulatory processes, and the form of modulation used, are probed through a comparative study of a number of variants of the basic GasNet. The results of this study are discussed in the light of related work before drawing conclusions.

2. THE BASIC GASNET

By analogy with biological neuronal networks, GasNets incorporate two distinct signalling mechanisms, one "electrical" (which henceforth refers to direct point-to-point synaptic connections) and one "chemical" (which henceforth refers to diffusional volume signalling). The underlying "electrical" network is a discrete time step, recurrent neural network with a variable number of nodes. These nodes are connected by either excitatory or inhibitory links with the output, O_i^t, of node i at time step t determined by Eq. (1).

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

where Γ_i is the set of nodes with connections to node i and $w_{ji} = \pm 1$ is a connection weight. I_i^t is the external (sensory) input to node i at time t, and b_i is a genetically set bias. Each node has a genetically set default transfer function gain 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^t (as described later).

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 behavior 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. Spatially, the gas concentration varies as an inverse exponential of the distance from the emitting node with spread restricted by a parameter, r, genetically set for each node [Eq. 2 and Figure 1]. The maximum concentration at the emitting node is 1.0, and the concentration builds up and decays linearly as dictated by the time course function, $T(t)$, defined by Eq. (3).

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

FIGURE 1

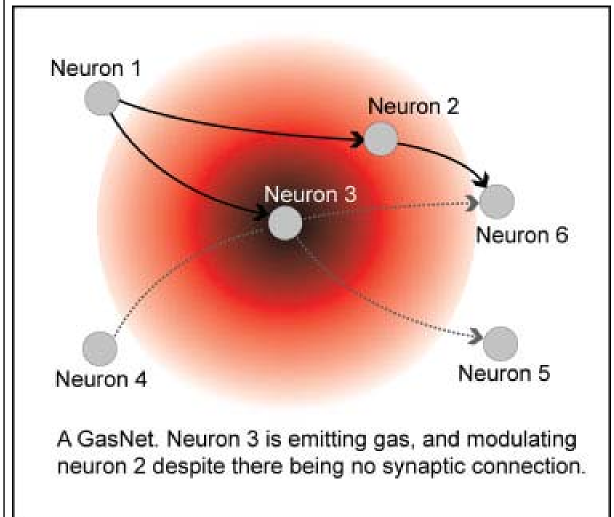

A basic GasNet showing positive (solid) and negative (dashed) "electrical" connections and a diffusing virtual gas creating a "chemical" gradient. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

$$T(t) = \begin{cases} 0 & t = 0 \\ \min\left(1.0, \left(T(t-1) + \frac{1}{s}\right)\right) & \text{emitting} \\ \max\left(0.0, \left(T(t-1) - \frac{1}{s}\right)\right) & \text{not emitting} \end{cases} \tag{3}$$

where $C(d,t)$ is the concentration at a distance d from the emitting node at time t and s (controlling the slope of the function T) is genetically determined for each node. The range of s is such that the gas diffusion timescale can range from $\frac{1}{2}$ to $\frac{1}{11}$ of the timescale of "electrical" transmission (i.e., a little slower to much slower). 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 emitted gases to avoid runaway positive feedback). The diffusion process is modeled in this simple way to provide extreme computational efficiency, allowing arbitrarily large networks to be run very fast (detailed modeling of 2D or 3D diffusion from multiple sources is very expensive [5] and, until very recently, could not be run in real-time without dedicated parallel hardware).

For mathematical convenience, in the original basic GasNet there are two "gases," one whose modulatory effect is to increase the transfer function gain parameter [k_i^t from Eq. (1)] 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

DOI 10.1002/cplx

threshold; note these emission processes provide a coupling between the "electrical" and "chemical" mechanisms). The concentration-dependent modulation is described by Eq. (4), with transfer parameters updated on every time step as the network runs.

$$k_i^t = k_i^0 + \alpha C_1^t - \beta C_2^t \tag{4}$$

where k_i^0 is the genetically set default value for k_i, C_1^t and C_2^t are the concentrations of gas 1 and gas 2, respectively, at node i on time step t, and α and β are constants such that $k_i^t \in [-4, 4]$. 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 general form of the diffusion is based on the properties of a (real) single source neuron as modelled in detail in [3]. The modulation chosen is motivated by what is known of nitric oxide (NO) modulatory effects at synapses [13]. For further details, see [8, 11].

A number of extended forms of GasNet have been developed and will be mentioned later. However, this article is mainly concerned with a study of variants of the basic GasNet which implement a range of constraints on spatio-temporal properties to probe the effect of such factors on evolvability.

3. EVOLVABILITY AND OPERATION OF GASNETS

When they were first introduced, GasNets were demonstrated to be significantly more evolvable than a variety of standard ANNs on some noisy visually guided evolutionary robotics tasks [8]. Typically, the increase in evolvability, in terms of number of fitness evaluations to a reliable good solution, was an order of magnitude or more. This trend was repeated with a wider range of robot tasks in later work [9, 10, 12], where GasNets were found to be more evolvable than many other forms of ANN. The solutions found were often very lean with few nodes and connections, typically far fewer than was needed for other forms of ANN [8]. But the action of the modulatory gases imbued such networks with intricate dynamics: they could not be described as simple. Oscillatory subnetworks based on interacting "electrical" and "gas" feedback mechanisms acting on different timescales were found to be very easy to evolve and cropped up in many forms, from CPG circuits for locomotion [9, 14] to noise filters and timing mechanisms for visual processing [8, 15]. GasNets appeared to be particularly suited to noisy sensorimotor behaviors which could not be solved by simple reactive feedforward systems, and to rhythmical behaviors (on which they often performed as well as or better than CTRNNs [16] with complex internal dynamics).

Two recent extensions of the basic GasNet, the receptor and the plexus models, incorporated further influence from neuroscience [11]. In the receptor model, modulation of a node is now a function of gas concentration and the quantity and type of receptors (if any) at the node. This allows a range of site specific modulations within the same network. In the plexus model, inspired by a type of NO signalling seen in the mammalian cerebral cortex [5], the emitted gas "cloud," which now has a flat concentration, is no longer centred on the node controlling it but at a distance from it. Both these extended forms proved to be significantly more evolvable again than the basic GasNet [11].

3.1. What is the Trick?

The question naturally arises as to why the GasNet and variants are more evolvable. Intriguingly, in a comprehensive study Smith et al. [10], found no explanation for increased GasNet evolvability in terms of fitness landscape properties (neutrality, epistasis, etc.) apart from at high fitness values. There it was argued that the key to understanding the improvement of the GasNet was to analyze its behavior at a higher level of abstraction. In particular, it was shown that the temporal dynamics of the GasNet seemed to make it relatively easy to tune the networks to the time-scales needed in the task [15]. Similar high-level analyses of the spatial structure of successful GasNets and variants, led to the hypothesis that it was the level of coupling between the electrical and gas signalling systems that was key. In particular that successful evolution came through the systems being flexibly coupled: neither independent of each other nor too tightly bound, allowing one system to be "tuned" against the other without causing catastrophic destructive interference [11]. Throughout, however, it was clear that these factors did not act in isolation and that it is the modulatory effect of the gas that lends the networks their adaptivity. This leads to three linked hypotheses on why the GasNets evolve faster:

1. The action of gas over multiple different timescales from the electrical activity introduces rich dynamics which can be exploited.
2. The spatial embedding of the networks serves to (flexibly) couple two interacting signalling systems.
3. The particular modulatory effects are key to evolvability.

In this article, these hypotheses are examined in the light of empirical studies, focusing on a comparison of variants of the basic GasNet formed by imposing various constraints on spatial, temporal, and modulatory properties.

4. COMPARATIVE STUDY

A thorough comparative study of the evolvability of variants of the basic GasNet was carried out on a noisy robotic visual discrimination task. This task has been used before in detailed comparisons between the GasNets and other styles of ANN [8, 11], allowing some comparison to be made across the studies. The study used variants of the basic GasNet, rather than of extended forms of GasNet, to avoid confusing the issue while trying to sift out which factors are most important to evolvability.

DOI 10.1002/cplx

FIGURE 2

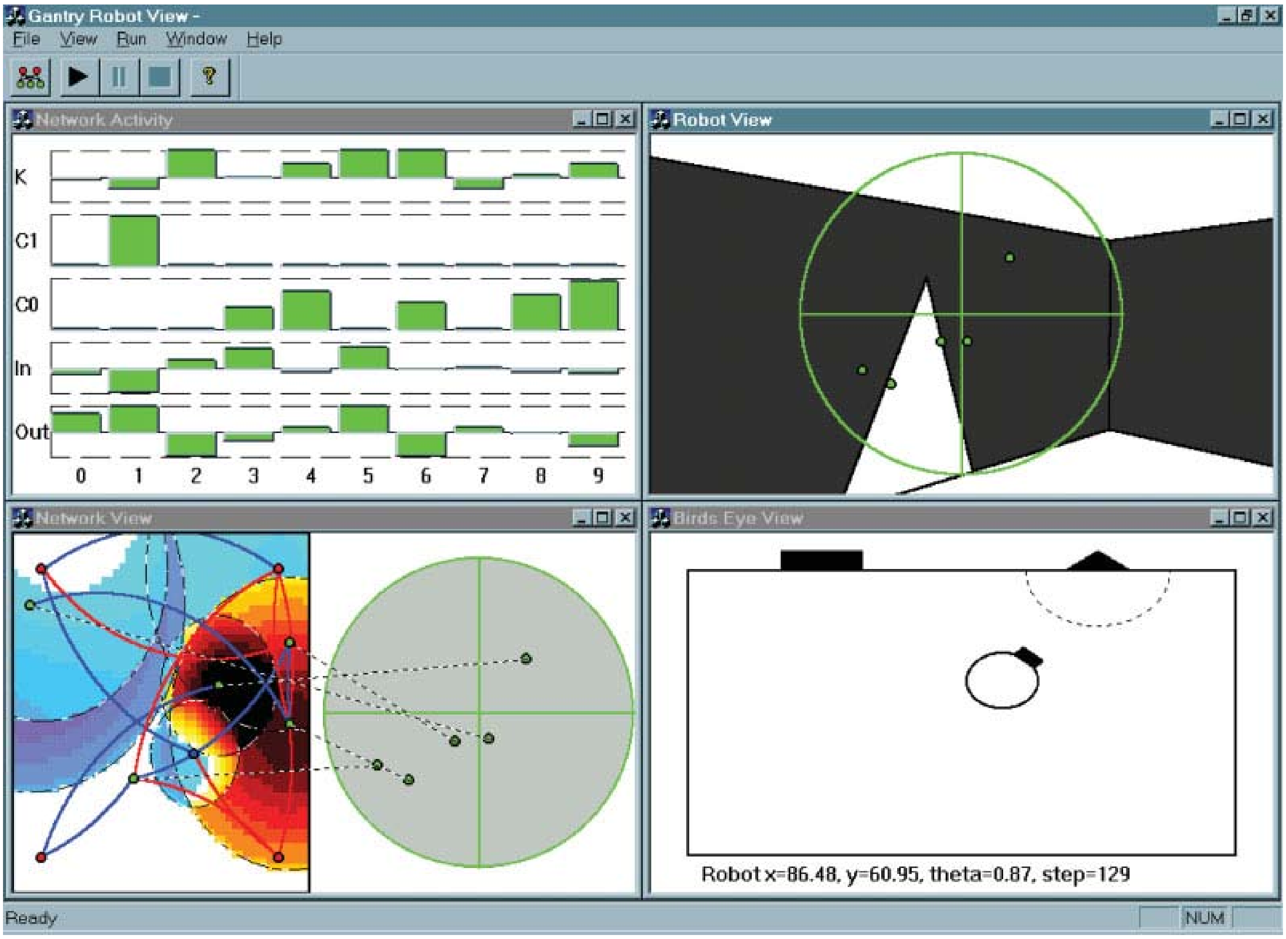

The simulated arena and robot. The bottom right view shows the robot position in the arena with the triangle and rectangle. Fitness is evaluated on how close the robot approaches the triangle. The top right view shows what the robot "sees," along with the pixel positions selected by evolution for visual input. The bottom left view shows how the genetically set pixels are connected into the control network whose gas levels are illustrated. The top left view shows current activity of nodes in GasNet. A validated simulation of the robot shown in Figure 3 was used.

4.1. GasNet Types

Nine variants of the basic GasNet were compared with probe the questions raised in the previous section. Each of these employed different constraints on spatial, temporal, or modulatory properties of the network. Referring to the labels used in the later results section (section 4.3), these were: *gnet*, the basic GasNet as described in section 2; *nchem*, the basic GasNet with all chemical effects turned off, so effectively a simple recurrent connectionist network; *gnetN*, the basic GasNet with no diffusion dynamics [i.e., $T(t) = 1$, $\forall t$, where $T(t)$ is the time course function of Eq. (3)]; *gnetNw*, the same as gnetN but with $T(t) = w$ where $w \in \{0, 1, 2\}$ is a "gas weight" genetically set for each node; *flatR*, same as basic GasNet except the gas concentration within the genetically set radius for each emitter is flat with no gradient [the term $e^{-2d/r}$ in Eq. (2) is replaced by e^{-1}]; *flatRN*, same as flatR except without diffusion dynamics ($T(t) = 1$, $\forall t$); *flatE*, same as flatR except the influence of the gas is not confined to the genetically set radius of influence for a node but now extends everywhere; *flatEN*, same as flatE but without diffusion dynamics; *AddMod*, the most radical variant where the multiplicative modulation of the basic GasNet is replaced by an additive modulation as described by Eq. (5) (i.e., the gas no longer modulates the transfer function gain parameter but instead modulates an additional additive bias term).

$$O_i^t = \tanh\left[k_i^0\left(\sum_{j\in\Gamma_i} w_{ji}O_j^{t-1} + I_i^t\right) + b_i + \gamma_i(C_1^t - C_2^t)\right] \quad (5)$$

DOI 10.1002/cplx

where C_1^t and C_2^t are gas concentrations, $\gamma_i \in [0, 2]$ is a genetically set parameter and all other terms are as in Eq. (1).

4.2. 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 rectangle). The relative position of the shapes varied from trial to trial, as did their size within a variation of 10% in both dimensions. Because of the noise and variation, and limited sensory capabilities, this task is challenging, resulting in a complex rugged search space for all network types studied. Both the robot control network (one or other GasNet variant) and the robot sensor input morphology, i.e., the number and position of the camera pixels used as input and how they were connected into the network, were under evolutionary control. This is illustrated in Figures 2 and 3. Fitness was measured using the function F given in Eq. (6), a weighted sum of values from 16 trials of the controller from different random initial conditions:

$$F = \frac{2}{N(N+1)} \sum_{i=1}^{i=N} i \left(1 - \frac{D_i^F}{D_i^S}\right) \tag{6}$$

where D_i^F is the distance to the triangle at the end of the ith trial, and D_i^S the distance to the triangle at the start of the trial, and the N (= 16) trials are sorted in descending order of $\frac{D^F}{D^S}$. Thus good trials, in which the controller moves some way toward the triangle, receive a smaller weighting than bad trials — this form of fitness punishment encouraging robust behavior on all 16 trials. Success in the task was taken as an evaluated fitness of 1.0 (perfect score) over 10 successive generations of the evolutionary algorithm. Evaluations took place using a validated minimal simulation of the robot as described in [8]. Controllers developed in the simulation successfully transferred to reality generating behaviors 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 [8].

A geographically distributed asynchronous updating evolutionary algorithm was used [8], 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 eight nearest neighbors 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.

The robot controllers were encoded as a variable length string of integers, with each integer lying in the range [0, 99]. Each node in the network was coded for by 19 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 if so, where it comes from in the camera field of view, and all gas diffusion and modulation parameters. Connections were formed using a spatial scheme as described in [8], with each node connecting to nodes lying within two genetically specified connection segments (one for excitatory and one for inhibitory connections). Hence, each genotype consists of N blocks of 19 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).

FIGURE 3

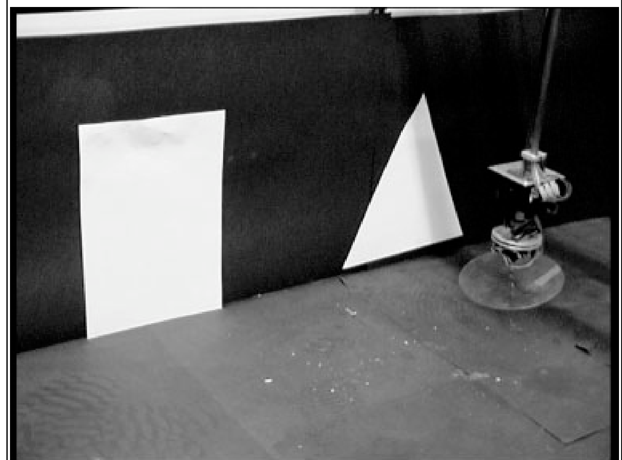

The gantry robot used in this study. A CCD camera head moves at the end of a gantry arm allowing full 3D movement. In this study, 2D movement was used, equivalent to a wheeled robot with a fixed forward pointing camera.

Three mutation operators were applied to solutions during evolution. Each integer in the string had a 10% probability of mutation in a Gaussian distribution around its current value (for certain genes, 20% of its mutation will be random jumps within the full possible range). 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 the 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.

4.3. Results

Results of the comparative study are summarized in Figure 4. A quick glance suggest that the basic gasnet (group 1) is the most consistently evolvable with group 9 (AddMod) clearly the worst (no runs were successful). Group 2 (nchem), in which gas effects are turned off, performs poorly on most runs, although, like most other variants, some runs produce

DOI 10.1002/cplx

FIGURE 4

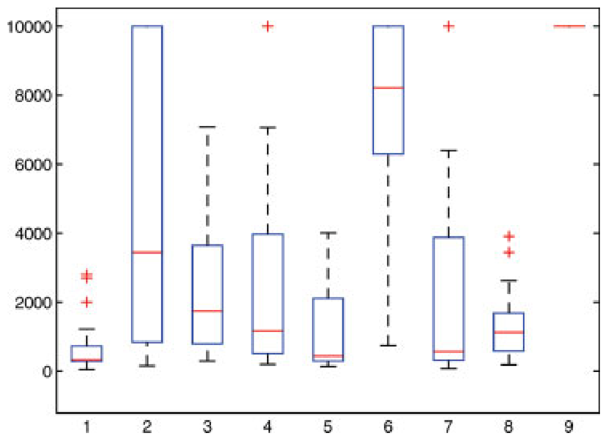

Boxplot summarising results of the comparative study. The *X* axis refers to the network groups as discussed in the text. The *Y* axis shows generations to success as defined by the stopping criteria explained in the text. The horizontal line within each box is the median, the top, and bottom of the box show the 75th and 25th percentiles, respectively, the whiskers extend to extreme points of the data not considered outliers (as defined by Rosner's test), with outliers plotted individually. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

good solutions relatively quickly. Most other network types without diffusion dynamics, thereby robbed of rich temporal properties, including multiple timescales, perform relatively poorly (groups 3, 4, 6). However, the relatively good performance of group 8, without dynamics, especially compared with group 7, which has dynamics, suggests that the story is not quite as simple as it might at first appear. Because it is not possible to assume the data distributions are normal, nonparametric statistical procedures were used to test for significant differences between the network types. A Kruskal-Wallis test performed on the whole data set (all 9 groups) revealed highly significant differences between the distributions ($p < 10^{-14}$). Pair-wise Wilcoxon rank-sum tests, adjusted for multiple comparisons using the Dunn-Sidak procedure for controlling type-1 statistical errors [17], were used to further probe the differences between the distributions. These tests showed that all network types, except group 6 (flatRND), were significantly more evolvable (in terms of generations to consistent success) than group 9 (AddMod). Because the Dunn-Sidak procedure is necessarily conservative and becomes more so as the number of groups increases, pairwise comparisons were recalculated for all network types except AddMod (i.e., groups 1–8). The results of these comparisons are shown in Table 1. Comparisons between smaller independent subgroups are shown in Tables 2 and 3; these reveal further significant differences within the context of a smaller number of comparisons. Further discussion of these results, including insights from other related work, is given in the next section.

5. DISCUSSION

This section draws together conclusions from the comparative study. It proceeds by considering in turn each of the factors identified in section 3, before turning to other phenomena that might help to explain the results.

5.1. Dynamics and Timescales

The figures and tables of section 4.3 reveal the importance of the dynamics conferred by the diffusing gases. The basic GasNet (group 1) is significantly more evolvable than the variant with the gas turned off (group 2) as well as the variants with the gas operating but without dynamics (groups 3 and 4). It is also significantly more evolvable than the variant with the gas operating but with neither a concentration gradient nor dynamics (group 6). However, there is one group

TABLE 1

Summary of Tests for Differences Between Evolvability (Generations to Consistent Success)

Sig Diff?	(1) gnet	(2) nchem	(3) gnetN	(4) gnetNw	(5) flatR	(6) flatRN	(7) flatE	(8) flatEN
(1) gnet	n	**Y**	**Y**	n	n	**Y**	n	n
(2) nchem	**Y**	n	n	n	n	n	n	n
(3) gnetN	**Y**	n	n	n	n	n	n	n
(4) gnetNw	n	n	n	n	n	**Y**	n	n
(5) flatR	n	n	n	n	n	**Y**	n	n
(6) flatRN	**Y**	n	n	**Y**	**Y**	n	**Y**	**Y**
(7) flatE	n	n	n	n	n	**Y**	n	n
(8) flatEN	n	n	n	n	n	**Y**	n	n

Distributions for network types 1–8 were tested against each other using pair-wise Wilcoxon rank-sum tests adjusted for multiple comparisons using the Dunn-Sidak procedure. Cell entries state whether or not there is a significant difference between the two distributions in question ($p < 0.05$).

DOI 10.1002/cplx

TABLE 2

Pairwise Comparison Tests for Subgroup of Standard GasNet and its Simple Variants with no Diffusion Dynamics

Sig diff?	(1) gnet	(3) gnetND	(4) gnetNDw
(1) gnet	n	**Y**	**Y**
(3) gnetND	**Y**	n	n
(4) gnetNDw	**Y**	n	n

Methodology as in Table 1.

without gas dynamic that the basic GasNet is not significantly more evolvable than group 8 (flatEN), of which more later. The version of the GasNet with diffusion dynamics but without a concentration gradient (group 5, flatR) performs fairly well with a low minimum and median, but the fairly high spread of results means that it is not as reliably evolvable as the basic GasNet. There is a similar story for group 7 (flatE) but its reliability is even worse; it should be noted that this restricted form of GasNet has similarities with various network models of neuromodulation that use global modulator signals [18].

Although these results suggest there is more to the GasNet's evolvability than the multiple timescales provided by the gas diffusion dynamics, they do add a certain amount of weight to previous suggestions that their easily tunable dynamics is an important part of their success (as well as to more general claims about the importance of dynamics in the generation of behavior [19]). These suggestion came from Smith et al.'s work [15] on taking GasNet and NoGas (equivalent to nchem, group 2) networks successfully evolved for the same robot visual discrimination task used in this article and re-evolving them in versions of the robot simulator where the natural timescales have been radically altered by making the motor speeds much slower or much faster. It was found that the GasNet was able to successfully re-evolve significantly faster. Analysis of networks showed how the frequency of crucial oscillatory circuits, involving gas diffusion and used to time active visual scanning movements central to the successful behavior, could be easily changed by moving nodes (so the time taken for gas concentration to build up at its target was increased/decreased) and/or changing the slope of the diffusion timecourse function $T(t)$ [Eq. (3)]. These results were backed up by a further study of GasNets on a simple pattern generation task. Four simple test-patterns were used, all consisting of a number of ones followed by a number of zeros, repeating for the entire 200 time-steps of fitness evaluation. Networks were scored according to how closely they produced the required pattern [20]. In this case, a fixed architecture of four fully connected nodes was used; although the positions of the nodes and the spatial aspects of the gas diffusion were under evolutionary control, the architecture of "electrical" connectivity was not (although connection weights were). Again GasNets were shown to be more adaptable to changed timescales. The pattern generation study was only intended as a minor addition to the robot studies, but its computational simplicity has prompted others to use it; although its very simplicity and its architectural constraints do not necessarily make it ideal for studying GasNets, which were originally intended for the generation of sensorimotor behaviors in noisy environments.

5.2. Modulation

Even more obvious is the role of the type of modulation used — additive modulation proved to be useless (group 9). The multiplicative modulation employed in all other variants is able to assert a much more drastic influence on a node, being able to radically change the transfer function by altering the gain k_i^t [Eq. (1)] — for instance flipping the slope from positive to negative or making it flat. These kinds of radical changes were dynamically employed in most successful GasNets and were at the heart of mechanisms, such as oscillators, used to produce stable reliable behavior in the face of significant noise [8, 15]. Additive modulation, which acts at the same level as a node input or bias, could not produce strong enough effects to generate stable behavior. An earlier study on the simple pattern generation task mentioned above also showed that multiplicative modulation was significantly better than additive in a form of GasNet similar to the basic one used here [21]. When GasNets were first introduced [22], an alternative node transfer function was successfully used. This had the form $O_i(S_i) = (S_i^a + S_i^b)/2$, where S_i is the normalized input to the ith node, and the parameters a and b were modulated by the gases. This form of modulation also allowed potentially large alterations to the transfer function, which seems to be necessary for effective evolution. These kinds of (multiplicative or exponential) modulations may well confer evolutionary advantages by allowing network nodes to be sensitive to different ranges of input (internal and sensory) in different contexts. For instance, in one (behavioral) context

TABLE 3

Pairwise Comparison Tests for Subgroup of all GasNet Variants with no Diffusion Dynamics

Sig Diff?	(3) gnetN	(4) gnetNw	(6) flatRN	(8) flatEN
(3) gnetN	n	n	**Y**	n
(4) gnetNw	n	n	**Y**	n
(6) flatRN	**Y**	**Y**	n	**Y**
(8) flatEN	n	n	**Y**	n

Methodology as in Table 1.

DOI 10.1002/cplx

an input node may need to be sensitive to a range of low sensor values, whereas in another, it is required to be sensitive to a range of high values. Changing a node's gain through multiplicative modulation allows its sensitivity to be adjusted in an appropriate way. It has been shown that "calibrating" sensor inputs into suitably exploitable ranges is far from straightforward for most forms of evolved neurocontroller and that additional mechanisms are needed [23]. It seems that GasNets have a triple advantage in that such a mechanism is not only available but it can be applied adaptively throughout the task and is not confined to sensory inputs.

5.3. Spatial Embedding and Coupling

The comparative study does not give such a strong indication of the impact of the spatial properties of the GasNet. The flatR (group 5) variant is identical to the basic GasNet, except the gas concentration within the radial extent of the emitter is flat. Although the evolvability of this group is good (not significantly different from the basic GasNet), the larger spread of generations to success suggests a slight advantage in utilizing a concentration gradient (group 1). The relative performances of flatR and flatE (group 7), although not significantly different, again suggests a slight advantage in constraining the spatial extent of the gas signal. These observations are backed up by results using a nonspatial variant of the basic GasNet. In that case, the network does not reside in a 2D space, and there are explicitly coded "electrical" and "gas" weighted connections between nodes; the gas connections have the same dynamics and modulatory effect as in the basic GasNet [24]. On the simple pattern generation task mentioned above, it was shown that there was no significant difference in evolvability between the basic GasNet and the nonspatial GasNet [24]. A similar independent study came to the same conclusion [21]. Subsequently, it has been demonstrated that a range of robot behaviors can be readily evolved using nonspatial GasNets [12].

These results suggest that the dynamic modulatory processes at play in the GasNet are the important thing rather than the exact details of their implementation. The very abstract computationally cheap diffusion model used could just as easily be implemented as a special kind of direct gas connection. However, analysis of the evolvability of the extended forms of GasNet (receptor and plexus) outlined near the start of section 3 indicate other roles for spatial embedding. It was described earlier how the "electrical" connection encoding relies on an indirect spatial process using segments centred on the node. As the gas emission process is based on circles centred on the node, there is significant room for overlap between "electrical" connections and gas influence between any pair of nodes, potentially causing destructive interference between the two signalling mechanisms. With the receptor model, gas influences between pairs of nodes can easily be altered by the presence or absence of receptors. With the plexus model, overlap between the signalling mechanisms can easily be reduced by moving the (genetically determined) centre of the emitted "cloud" away from the location of the emitting node. In this way, both these forms of GasNet are able to significantly reduce destructive interference between the signalling systems, instead exploiting a loose, flexible coupling between them. This results in significantly improved evolvability in these extended forms of GasNet [11]. In these cases, the spatial embedding is a convenient way to encourage the right kind of coupling.

Reduction in destructive interference is probably also the explanation for the surprisingly good performance of group 8 (flatEN) in the comparative study, despite the lack of dynamics in this variant. Because the gas concentration from any node is flat and extends everywhere in the plane, the gas influence between all pairs of nodes is constant and cannot be changed by moving their positions — space has effectively being taken out of the gas signalling system which is fully connected. Although this does not completely remove potential for destructive interference, it does mean that any changes in "electrical" connectivity act against a constant backdrop of gas "connections" which helps to reduce coupling, leaving evolution free to "tune" the electrical system against the unchanging gas system — a much easier job than tuning two tightly coupled systems against each other. The performance of this variant also demonstrates that although high evolvability is not possible with both gas modulation and diffusion dynamics switched off (nchem, group 2), it is possible, in certain circumstances, without dynamics as long as gas modulation is active.

The initial motivation for using a spatial model, albeit a simple one, was to allow easy visualisation of the networks and their operation, and of course because the natural phenomenon that inspired the work is quintessentially spatiotemporal. It seems that for the simpler forms of GasNet, the spatial aspects are not integral or essential to their success, and they can just as well be implemented in other ways. However, this is unlikely to be the case in extended forms. In ongoing work on more complex networks, involving much more detailed diffusion models, space is not a mere implementation detail — it is essential. In these kinds of networks diffusing chemicals act as excitable media, forming complex spatiotemporal patterns which can potentially greatly increase the power of the systems. Indeed, it has recently been shown that these patterns can themselves act as control systems generating fairly complex memory-based behaviors in simulated agents, even without any interacting neurons at all [25].

5.4. Degeneracy

In biology, degeneracy is the property whereby structurally different elements perform the same function or produce the same output [26]. Because the function of such elements is nearly always context dependent, this phenomenon is distinguished from simple redundancy which refers to identical

DOI 10.1002/cplx

structures performing the same function. As Edelman and Gally [27] have pointed out, most, if not all, biological systems are highly degenerate at all levels of structure and function. They argue that this property, which involves a tradeoff between specificity and generality, confers adaptability and, indeed, is a prerequisite for natural selection to operate effectively. Analysis of GasNet solutions often reveals high levels of degeneracy, with functionally equivalent subnetworks occurring in many different forms, some involving gas and some not [15]. Their genotype to phenotype mapping (where the phenotype is robot behavior) is also highly degenerate with many different ways of achieving the same outcome (e.g., moving node positions, changing gas diffusion parameters or adding new connections can all have the same effect). This is especially true when variable length genotypes are used to efficiently sculpt solutions in a search space of variable dimensions. The levels of degeneracy are generally significantly higher than when using connectionist networks. These properties partly explain the robustness and adaptability of GasNets in noisy environments as well as their evolvability (there are many paths to the same phenotypical outcome with reduced probabilities of lethal mutations) [11].

5.5. Modularity

In a recent study on the interaction of spatial embedding and modularity in neural networks, successfully evolved basic and plexus GasNet solutions for the triangle-rectangle discrimination task were analyzed to test whether or not the different spatial embeddings led to differences in the modularity of the networks [28]. This revealed that all of the plexus topologies consisted of one component (i.e., every node was reachable from every other node), whilst 14 (out of 33) of the original GasNet runs produced a best performing controller with at least two network components. In other words, the plexus GasNets were less modular. It was hypothesized that this lower degree of modularity was linked to the different spatial constraints in operation in the plexus networks and it might explain, in a way that is complementary to coupling, their greater evolvability. Given that the required behavior is decomposable into subtasks but requiring communication between these subparts [15], an alternative explanation might be that nondisconnected networks are desirable. Indeed, one could go further and argue that the basic GasNets were more likely to produce disconnected nets because it was difficult to implement communication between subparts without destructive interference, because of the higher level of coupling, while the looser coupling of the plexus networks allows for all parts to interact.

6. CONCLUSION

The comparative study presented in this article suggests that GasNets employing multiplicative modulation with dynamics are the most evolvable, and it is shown that this finding is backed up by other research. However, the study also revealed that diffusion dynamics are not necessary for high evolvability: in some circumstances "instant" modulation can do the trick, as demonstrated by the flatEN variant (group 8). In extended forms of GasNet, spatial embedding seems to allow easy exploitation of the most beneficial kind of coupling between the "electrical" and "gas" signalling mechanisms, reducing destructive interference between them. A loosening of this coupling also appears to be behind the success of the flatEN variant. In the most successful forms of GasNet, dynamics, modulation, and spatial embedding act in concert to produce highly evolvable degenerate networks. In future work, these issues will be explored with more strongly biologically inspired networks embedded in excitable media formed by reaction-diffusion systems, thus allowing much richer spatiotemporal modulatory patterns to emerge which will enable us to explore functionally and architecturally more complex networks.

Acknowledgments

This research was supported by EPSRC grant EP/C51632X/1 and EU FP7 grant eFlux. Authors thank two anonymous referees for useful comments on an earlier draft.

REFERENCES

1. Gally, J.; Montague, P.; Reeke, G.; Edelman, G. The NO hypothesis: Possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. Proc Natl Acad Sci USA 1990, 87, 3547–3551.
2. Dawson, T.; Snyder, S. Gases as biological messengers: Nitric oxide and carbon monoxide in the brain. J Neurosci 1994, 14, 5147–5159.
3. Philippides, A.; Husbands, P.; O'Shea, M. Four-dimensional neuronal signaling by nitric oxide: A computational analysis. J Neurosci 2000, 20, 1199–1207.
4. Hölscher, C. Nitric oxide, the enigmatic neuronal messenger: Its role in synaptic plasticity. TINS 1997, 20, 298–303.
5. Philippides, A.; Ott, S.; Husbands, P.; Lovick, T.; O'Shea, M. Modeling co-operative volume signaling in a plexus of nitric oxide synthase-expressing neurons. J Neurosci 2005, 25, 6520–6532.
6. Bullock, T.H.; Bennett, M.V.L.; Johnston, D.; Josephson, R.; Marder, E.; Fields, R.D. The neuron doctrine, redux. Science 2005, 310, 791–793.
7. Katz, P., Ed. Beyond Neurotransmission: Neuromodulation and its Importance for Information Processing; Oxford University Press: Oxford, 1999.
8. Husbands, P.; Smith, T.; Jakobi, N.; O'Shea, M. Better living through chemistry: Evolving GasNets for robot control. Connect Sci 1998, 10, 185–210.
9. McHale, G.; Husbands, P. Quadrupedal locomotion: Gasnets, ctrnns and hybrid ctrnn/pnns compared. In: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems (Alife IX); Pollack, J.; Bedau, M.; Husbands, P.; Ikegami, T.; Watson, R., Eds.; MIT Press: Cambridge, MA, 2004; pp. 106–112.

DOI 10.1002/cplx

10. Smith, T.; Husbands, P.; O'Shea, M. Local evolvability of statistically neutral gasnet robot controllers. Biosystems 2003, 69, 223–243.
11. Philippides, A.; Husbands, P.; Smith, T.; O'Shea, M. Flexible couplings: Diffusing neuromodulators and adaptive robotics. Artif Life 2005, 11, 139–160.
12. Vargas, P.; Paolo, E.D.; Husbands, P. Exploring non-spatial gasnets in a delayed response robot task. In: Proceedings Alife XI; Bullock, S. et al., Ed.; MIT Press: Cambridge, MA, 2008; pp. 640–647.
13. Barañano, D.; Ferris, C.; Snyder, S. Atypical neural messengers. Trends Neurosci 2001, 24, 99–106.
14. McHale, G.; Husbands, P. Gasnets and other evovalble neural networks applied to bipedal locomotion. In: Proc. From Animals to Animats 8: Proceedings of the Eigth International Conference on Simulation of Adaptive Behavior (SAB'2004); Schaal, S., Ed.; MIT Press: Cambridge, MA, 2004; pp. 163–172.
15. Smith, T.; Husbands, P.; Philippides, A.; O'Shea, M. Neuronal plasticity and temporal adaptivity: Gasnet robot control networks. Adapt Behav 2002, 10, 161–184.
16. Beer, R.D.; Gallagher, J.C. Evolving dynamical neural networks for adaptive behavior. Adapt Behav 1992, 1, 94–110.
17. Hollander, M.; Wolfe, D. Non-parametric statistical methods; Wiley: New York, 1999.
18. Doya, K. Metalearning and neuromodulation. Neural Netw 2002, 15, 495–506.
19. Port, R.F.; van Gelder, T., Eds. Mind as Motion: Explorations in the Dynamics of Cognition; MIT Press: Cambridge, MA, 1995.
20. Smith, T. The evolvability of artificial neural networks for robot control, PhD Thesis, University of Sussex, Brighton, 2002.
21. Buckley, C. A systemic analysis of the ideas immanent in neuromodulation, PhD Thesis, University of Southampton, Southampton, 2008.
22. Husbands, P. Evolving robot behaviors with diffusing gas networks. In: Evolutionary Robotics: First European Workshop, EvoRobot98, Lecture Notes in Computer Science 1468; Husbands, P.; Meyer, J.-A., Eds.; Springer Verlag: Berlin, 1998; pp. 71–86.
23. Macinnes, I.; DiPaolo, E. The advantages of evolving perceptual cues. Adapt Behav 2006, 14, 147–156.
24. Vargas, P.; Paolo, E.D.; Husbands, P. Preliminary investigations on the evolvability of a non-spatial gasnet model. In: Proceedings ECAL'07, Lecture Notes in Computer Science 4648; e Costa, F. A. et al., Ed.; Springer-Verlag: Berlin, 2007; pp. 966–975.
25. Dale, K.; Husbands, P. The evolution of reaction-diffusion controllers for minimally cognitive agents. Artif Life 2010, 16, 1–19.
26. Tononi, G.; Sporns, O.; Edelman, G. Measures of degeneracy and redundancy in biological networks. Proc Natl Acad Sci USA 1999, 96, 3257–3262.
27. Edelman, G.; Gally, J. Degeneracy and complexity in biological systems. Proc Natl Acad Sci USA 2001, 98, 13763–13768.
28. Fine, P.; Di Paolo, E.; Philippides, A. Spatially constrained networks and the evolution of modular control systems. In: From Animals to Animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behavior; Nolfi, S., Ed.; Springer Verlag: Berlin, 2006; pp. 546–557.

DOI 10.1002/cplx