

Spatially Constrained Networks and the Evolution of Modular Control Systems

Peter Fine, Ezequiel Di Paolo, and Andrew Philippides

Centre for Computational Neuroscience and Robotics (CCNR),
Department of Informatics, University of Sussex,
Brighton, BN1 9QG, UK
{p.a.fine, ezequiel, andrewop}@sussex.ac.uk

Abstract. This paper investigates the relationship between spatially embedded neural network models and modularity. It is hypothesised that spatial constraints lead to a greater chance of evolving modular structures. Firstly, this is tested in a minimally modular task/controller scenario. Spatial networks were shown to possess the ability to generate modular controllers which were not found in standard, non-spatial forms of network connectivity. We then apply this insight to examine the effect of varying degrees of spatial constraint on the modularity of a controller operating in a more complex, situated and embodied simulated environment. We conclude that a bias towards modularity is perhaps not always a desirable property for a control system paradigm to possess.

1 Introduction

Modular control systems can be used to enhance our ability to generate adaptive behaviour. Based on the knowledge that brains are often shown to contain highly distinct regions of processing activity, artificial mechanisms which promote modularity are often added to neural networks in order to build similarly decomposable structures. Yamauchi & Beer [18] use separate neural networks, each evolved to perform dedicated functions and later combined, to carry out a sequential learning task that could not otherwise be satisfied. Often, however, such an a priori decomposition of a task into modules is not available. In order to benefit from modularity in a more generic way, work such as Calabretta et al. [1] examines modular neurocontrollers which emerge from the evolutionary optimization of their networks, rather than from any hardwired structural constraint. They achieve this using a genetic operator which duplicates existing modules, thus allowing the topology to develop on its own towards the best performing modular architectures.

Often in evolutionary robotics, a neurocontroller is generated by fully connecting each of an agent's sensors to each of its interneurons, which are in turn fully linked to its motors. If recurrent connections are allowed, interneurons will be connected to every other interneuron. The functional structure of the network can then emerge by variations in the *strength* of those connections, some of which may drop to near 0 and impose an effectively topological constraint on the network's dynamics. Alternatively, it is possible to structure the connectivity of the network, perhaps using

a genetically mediated mechanism to specify which nodes are connected. Lindenmeyer systems, for example, can generate repeating, modular architectures using a compact genetic representation [10].

Another method of configuring neural network architectures involves embedding the neurons and connections within an abstract ‘space’, where they interact directly with only their nearby neighbouring components. It is apparent that the use of such a spatial constraint has important implications regarding the presence of modularity, without requiring any explicit modularity-generating mechanisms. For this reason, this paper concentrates on exploring modular control systems, within spatially embedded networks. We study the nature of the relationship between modularity and spatiality, and exploit the discovered inherent modularity found in spatially embedded systems to look at whether a bias towards modularity is necessarily a positive feature. Prior uses of spatially embedded networks are discussed in section 2. Since we will need to quantify the level of modularity related to a given form of spatial constraint, methods for measuring modularity are also discussed. This uncovers further distinctions regarding the nature and role of modularity in complex systems.

Section 3 then investigates the supposed relationship in a simple, necessarily modular task/controller scenario, using an evolutionary algorithm to compare the ability of spatial and non-spatial networks to generate modular solutions. Having found a correlation, we wished to examine its potential relevance to our understanding of a more realistic scenario. Section 4 uses data from a more complex, simulated robot task with no apparent necessity for a modular controller. The role of differing degrees of spatial embeddedness is investigated using these experiments, suggesting that in some cases, less modularity may be a positive attribute. Finally, Section 5 discusses the limitations of our approach whilst offering our conclusions.

2 Spatially Embedded Networks and Modularity

Cliff and Miller [2] embed the nodes in their neural networks within a 2-D planar surface. Neurons can then connect to other nodes which are nearby. They find that this provides a natural framework within which to map a control system onto the actuators and sensors of a simulated agent. This can also be seen as biologically inspired, since real neurons within the brain emit axons to connect to their neighbours (see [16] for a review). More recently, Philippides et al. [12] have studied more complex features of spatial neural networks by including a model of neurotransmitter diffusion across the plane, allowing neurons to interact both electrically and chemically with other local neurons.

This embedding feature has implications regarding the presence of modularity within network topologies. The notion of a neuron’s local neighbourhood may encourage the division of network activity into distinct regions. One part of the spatial plane could implement a particular aspect of the control system, without interfering with (or being affected by) the other regions. Whilst a similar division of labour could potentially emerge from a fully connected network, it would be more difficult to generate, since without a spatial constraint a neuron is equally likely to interact with every other neuron in the network.

2.1 Measuring Modularity

The varied attempts to provide a metric for the quantification of modularity demonstrate different ways to conceive of its role in complex systems. One recent contribution by Newman & Girvan [11] searches for highly-interconnected communities within networks by first compiling them to abstract graphs of nodes and edges. This structural measure ignores the dynamics of the system under consideration, but nevertheless proved successful when applied to the modular architectures of a number of social interaction networks. Watson & Pollack [17] offer a measure which investigates the dynamical correlations between the node's behaviour. They claim that ignoring dynamics is liable to mislead, since some physical connections could be irrelevant to a network's overall behaviour. Conversely, even sparse connections between two topological clusters of nodes could play a dynamically highly significant role, undermining the distinctiveness of the supposed modules. Also, Polani et al. [14] present an information theoretic metric, which looks at the degree of 'mutual information' which is shared between subsystems.

Neurons often exhibit nonlinear interactions, and so in neural networks some connections will be likely ignored due to quiescent synapses. Others, however, may be exaggerated through phenomena such as bursting or hyper-excitability [8]. Measures which account for dynamical relationships will therefore no doubt be important in future studies of neural modularity. However we have reason to believe that a purely structural metric can still provide useful insights into the behaviour of the two cases studied in Sections 3 and 4 of this paper. In the first case, the network is made up of a simple spiking network which omits many of the more complex nonlinear features of real neurons. Combined with the fact that all weights were set to either 0 or 1, it is likely that neurons will influence each other in a quite homogenous manner across the network. A quantification of *structural* modularity may therefore tell us something about the modularity of the network's dynamics, as proved to be the case in Section 3.

The second case study involves networks comprising of two subsystems (gaseous and electrical, as described in Section 4). An analysis was undertaken to compare the structural relationships between these two systems with the level of overlap in their dynamical activity (presented in Philippides et al., [12]). They concluded that in this respect, the physical arrangement and dynamical behaviour was very similar despite some of the physical connections being inactive during the experimental trials. Whilst this is somewhat indirect and limited evidence towards a correspondence between the network's topology and its actual function, it suggests that if we restrict our analysis to structural organization of these neural networks, our conclusions may be informative to our understanding of a controller's behaviour. The validity of this approach forms an additional research question.

With this in mind, we used the structural measure of Newman & Girvan [11]. Since the spatial constraint operates by imposing restrictions on a network's physical connectivity, this measure provides the most direct way of studying how such constraints affect the development of the network. Additionally, it is clear that a network's physical topology will play a role in constraining its dynamics.

The measure proceeds by removing connections from a network, one by one. At each stage, the node with the highest degree of ‘betweenness centrality’ (a measure of the number of shortest paths between any two nodes which a given connection is a part of) is deleted. It is claimed that if a connection is frequently part of the shortest route between pairs of nodes, it is likely to be one of the sparse ‘bridges’ between different clusters of highly interconnected nodes. Eventually, as nodes are removed, the network breaks down into increasingly small and numerous separate graph components. Each discovered division into these component ‘modules’ is then examined and compared to see which provides the most appropriate division of nodes into clusters. The ‘best’ cluster division is determined in terms of the likelihood of a given node connecting to another member of its own module (rather than a different one). The proportion of intra-module connections versus inter-module links of the best rated division becomes a network’s measured ‘modularity rating’, which should be used comparatively to determine whether a network is more or less modular than an alternative structure [6].

3 Minimally Modular Control Systems

As detailed above, we have reason to believe that spatially constrained networks offer the potential for clustered architectures, without requiring any mechanisms which ‘create’ or impose modularity. To investigate whether this is true in practice, a genetic algorithm was used to evolve controllers for a simple task. The emergence of modular topologies could thus be investigated without explicitly biasing the controller towards them, allowing the conditions under which they arise to be discussed.

For this, we required a minimal task/neurocontroller combination which necessitated a modular structure for successful completion, whilst being simple enough to analyse. In the course of previous investigations into spiking network dynamics, this task was found to require two distinct modules to be present within the controller. Note that the task only requires a modular controller when paired with the spiking neuron model presented; we do not claim that it represents a necessarily modular *task* by itself. The justification for the claim that success here requires modularity is presented after the description of the methods used.

3.1 Experimental Method

A simple simulated ‘agent’ was used to evolve a memory task, where it had to move to either the left or the right of its ‘world’, depending on its prior (since extinguished) sensory input. The agent was placed on a 1-dimensional line (200 units in length), where it could move to the left or right based on the output of two network nodes which were designated as outputs (the location $x=0$ corresponds to the far left, and $x=200$ the far right). Two different nodes were used as inputs, each receiving signals from one of two externally controlled sensors (labelled the left and the right sensor). For each evaluation of a genotype’s fitness, two trials were run. In each, just one of the sensors was enabled (set to output a value of 1) for 2 ‘seconds’, whilst the agent was held in its initial location at the middle of the line ($x=100$). The input was set to 0, and only then were the agent’s motor neuron values used to move its location. High

fitness was achieved if the agent moved to the left of the line when its left sensor had been previously enabled, and the right of the line in the right sensor trials.

$$F_l = \frac{\Delta t \sum_{t=0}^T \left(\frac{x_t}{\hat{X}} \right)}{T}, \quad F_r = \frac{\Delta t \sum_{t=0}^T \left(\frac{\hat{X} - x_t}{\hat{X}} \right)}{T}. \quad (1)$$

Equation (1) shows the fitness measures F_l and F_r for the left and right sensor trials respectively, which were averaged to give an agent's fitness. Δt corresponds to the timestep of integration, T is the length of the trial, x_t the agent's position at time t , and \hat{X} the width of the line on which the agent resides.

A leaky integrate-and-fire neuron model was used (after [5], see [3] for use in evolutionary robotics setting). Sensor values were translated into input neuron spike trains using a Poisson process, with a probability of firing dependent on the sensor value plus noise. In addition to the 2 sensor and 2 motor neurons, 9 interneurons were included. The network activity is governed by:

$$\tau_m \frac{dV}{dt} = V_{rest} - V + g_{ex}(t)(E_{ex} - V) + g_{in}(t)(E_{in} - V). \quad (2)$$

$$g_{type}(t) \rightarrow g_{type}(t) + w_{ij}. \quad (3)$$

$$\tau_{ex} \frac{dg_{ex}}{dt} = -g_{ex}, \quad \tau_{in} \frac{dg_{in}}{dt} = -g_{in}. \quad (4)$$

$$\tau_{out} \frac{dM_x}{dt} = -M_x + M_G \left(\sum \delta(t_{now} - t_x) \right). \quad (5)$$

All parameters were evolved, unless otherwise specified. In Equation 2, τ_m is the membrane time constant [10ms, 40ms], V is the membrane potential, $V_{rest} = -70\text{mV}$ and the excitatory and inhibitory reversal potentials (E_{ex} and E_{in}) are 0mV and -80mV respectively. V 'spikes' when it reaches a normally distributed noisy threshold (V_{thresh} , mean [-60mV, -50mV] with deviation 1mV), and is followed by a random refractory period [2ms, 4ms] after which $V \rightarrow V_{rest}$. Equation 3 is applied when a neuron's inhibitory or excitatory synapse receives a spike (type = *in* or *ex* respectively, and w_{ij} is the connection strength from the presynaptic neuron). Equation 4 describes the decay of the input conductances, with τ values [2ms, 4ms]. Finally, Equation 5 shows the motor output integrator (with $x = \text{left or right}$), capped at 1, where t_x is the time of the last spike in the corresponding motor neuron, M_G is the motor gain (scaled exponentially between [0.1, 50]) and $\tau_{out} = [40\text{ms}, 100\text{ms}]$.

The connectivity of the network was generated by embedding each node in a 2 dimensional plane using a genetically specified x and y coordinate. Each neuron maintained a number of connections n (evolved per neuron between 0 and 5), which were made to the n nearest neighbouring neurons (excluding self-connections, which were not permitted). To aid analysis, all connections were made with a weight of 1. A

non-spatially embedded network was used as a control. Each neuron maintained an evolvable parameter in the range $[0, 1]$ for every other neuron in the network. They were then connected to the n nodes for which this value was highest. Other control experiments were run using a more typical, fully connected architecture with evolvable weights, producing similar results to the first control network (results not shown). Note that the spatiality of the experimental condition was only exploited to generate the network architecture; once a trial had begun the network operated as in the control condition.

All controllers were optimized using a rank based genetic algorithm with elitism. The non-elite genotypes were mutated by adding a vector of random numbers of total length 2.5, taken from a Gaussian distribution around 0. Any mutations which took a gene beyond its 0 to 1 limits were discarded, and a different random mutation selected for that gene until the boundaries were satisfied.

The fact that this task requires a modular architecture stems from the way in which the neurons reset their internal state value after every spike with a fast timescale (order of milliseconds). In order to retain activity for longer (allowing the agent to move in relation to the extinguished sensory input at a timescale extending over seconds), more than one neuron would thus be required. In fact, an exhaustive search of each possible two and three neuron topology (with weights constrained to $-1, 0$ or 1) provided no architecture capable of showing reliable persistent activity after the input was removed. Only when four neuron circuits were considered was an adequate ‘switch’ behaviour found, which could be triggered by an input to one of the nodes, and maintained regular spiking activity through its recurrent dynamics after that input was removed.

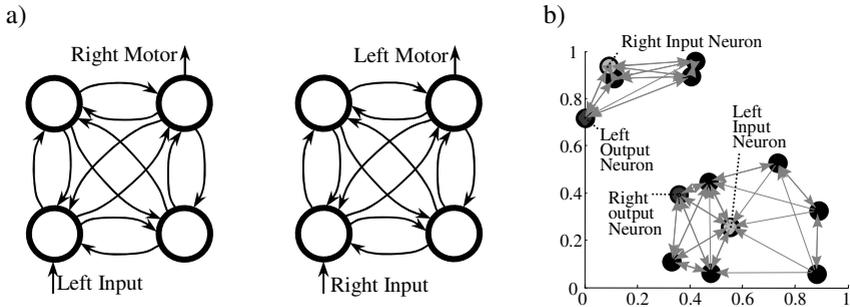


Fig. 1. a) Hand designed topology of spiking neurons which completes the task successfully. b) An example spatially embedded network, with a fitness of 0.825.

In order to complete the task, the agent would require two of these ‘switches’, one pairing the left sensor with the right motor (enabling the agent to move towards the left), and vice-versa. An example of this, using two four-node fully connected switches, is shown in Figure 1a. Whilst other mechanisms may exist, none were found throughout the extensive evolutionary simulations undertaken. The four neuron switches were somewhat robust to the removal of a limited number of connections. However the addition of a single connection linking any pair of nodes across different switches (thus breaking the perfect modularity) prevented the network from

succeeding. Activity would ‘leak’ across that connection, triggering both switches (and therefore both motors) equally, preventing the agent from moving.

3.2 Results

A set of 24 evolutionary runs of each condition were undertaken, each with a population of 50 for 800 generations. Every run of the spatially constrained network reached a maximum fitness of 0.825, which was the same as the hand designed version (a score of 1 is not possible because the agent cannot move at the start of each trial). The control network, however, could only reach 0.67. These networks were only able to move to one side of the line when one particular sensor was enabled, and did not move at all in the other case. An analysis of the spiking dynamics of one of these control networks showed that one sensor and one motor neuron had been disconnected from the network, and all of the other neurons acted as a switch enabling movement in the direction mediated by the remaining connected motor neuron. Only when the spatial constraint was added were two distinct switches able to evolve, as shown in the example (Figure 1b), which was taken from generation 800 of one of the spatially embedded runs. This is corroborated by the results from Newman & Girvan’s modularity measure, with the highest scoring spatial network from each run producing a mean score of 0.48, whilst the control network mean was 0.14 (significantly less modular, $p = 5.41 \times 10^{-15}$).

The above results demonstrate that even an apparently simple task, such as using a spiking network to ‘remember’ which direction an agent should move in, can necessitate a highly modular architecture. Crucially, this form of topology was not evolvable with a standard network generation technique. Only when embedded in space did the requisite modularity emerge, and it did so with apparent ease (in 100% of the runs carried out).

4 Controller Modularity in an Embodied, Active Vision Task

The previous results suggested a strong relationship between modularity and spatiality. However the use of an *a priori* modular task/controller, and the simplicity of the experiments leave open questions about the wider validity of this trend. The results of a complex scenario (which used a spatially embedded controller) were therefore analysed, looking at the modularity of the best performing controllers. Their use of a *partially* spatial controller (in addition to a fully embedded version) also allows for an interesting comparison to be made.

The experiments analysed in this section were carried out by Philippides et al. [12], and were an extension of the GasNet model (see reference for further details). GasNets are spatially embedded controllers, combining a standard neural network with a model of diffusing neurotransmitters. In addition to forming the topology of the network based on spatial factors (similarly to the spiking network experiments in Section 3), each node may act as a point source of gas. This diffuses across the plane, affecting nearby neurons by modulating their transfer function.

As in the minimal model presented in Section 3, each node in a GasNet maintains an evolvable x and y coordinate. Rather than connecting to its n nearest neighbours, it

maintains two genetically specified ‘cones’ emanating from the neuron’s position, for a limited distance across the plane. Positive connections are made to any neuron falling within the first of these, and negative connections to the second. The activity of these *electrical* connections are modelled as a discrete time, recurrent network which maintains a gain function k . Each neuron is also capable of producing one of two gases, either when its electrical output crosses a given threshold, or when gas in the vicinity reaches a certain level. Each of these parameters is genetically specified. The gases diffuse and disperse automatically. The summed concentration of the each gas is measured at each neuron’s location. The level of the first gas is used to increase the gain of the transfer function applied to the neuron’s electrical inputs, and the second gas decreases this value.

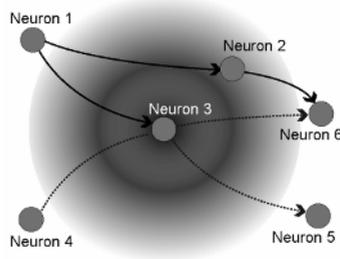


Fig. 2. An example GasNet electrical topology (not from experimental data), showing just the gas emitted by Neuron 3. The dashed lines denote negative connections.

The topologies were generated from the results of the experiment detailed in [12]. It involved a gantry robot starting from an arbitrary position and orientation in a black-walled arena. Equipped with a minimal vision system, taken from a forward-facing camera, it must navigate under extremely variable lighting conditions towards one shape (a white triangle), while ignoring the second shape (a white square).

Philippides et al. [12] also provide a modification of standard GasNets, taking inspiration from the way in which neurons emit Nitric Oxide (NO) in the mammalian cortex where it is generated by a network, or plexus, of fine fibres [13]. Rather than emanating from the neuron’s location, the gas forms a uniform cloud, targeted over a potentially distant area, away from the source node. This was termed the *plexus* model, and used two extra evolvable parameters per neuron, determining the x and y coordinates of the centre of that neuron’s gas cloud. This essentially means that whilst the electrical topology of the network is spatially constrained exactly as in the original GasNet, the gaseous connectivity (i.e. whether a pair of neurons influence each other through the production of gas) is not limited in this sense. Space still plays a role in this connectivity, because the gas from each neuron falls over a defined circular region of the plane. However this is a much more subtle constraint than in the original model, because that cloud can be located anywhere, instead of being restricted to being over nodes in the source neuron’s locality.

The plexus GasNets were found to be significantly more evolvable for this task, both in terms of the chance of a given run producing a high fitness controller, and the number of generations required to do so. The authors characterize the reasons for the

difference in terms of the lesser degree of coupling between the gaseous and electrical parts of the network's operation. They describe how during evolution changes in the gaseous structure of the network have a looser effect on the electrical systems (and vice versa). This is said to allow greater flexibility in 'tuning' each system against the other, with a corresponding greater level of evolvability.

Without wishing to question this factor, we propose a different, perhaps complementary hypothesis for the observed performance differences. In the plexus model, the spatiality of the electrical connections is preserved, but the gaseous links are no longer spatial (as described above). In light of the relationship between spatial constraints and modularity (see Section 3), it may be that the plexus networks are prone to a lower degree of modularity. With only a part of the structure restricted to neighbourhood-only connections, potentially parts of the overall network will be less modular and this may aid the development of a successful controller.

To determine whether this was the case, the best performing networks from 40 original GasNet runs and 40 plexus runs were compared. Firstly, the 7 original and 2 plexus runs which failed to reach the maximum fitness scores were removed, so as to only consider those networks which satisfied the task. Then, the excitatory and inhibitory electrical connectivity matrices were combined, so the resulting electrical matrix simply records whether a connection of either type was present or not (nodes connected by both positive *and* negative connections were assumed to be not connected, since the effect on the target neuron would be cancelled out). Also, the gaseous connectivity matrix was thresholded so that any non-zero level of influence between two nodes was regarded as a full connection, due to the Newman & Girvan metric's use of unweighted graphs. All connections were assumed to be undirected and self-connections were deleted, also due to further restrictions of the modularity measure.

4.1 Results

It is apparent from Figure 3 that the original GasNet produced overall topologies which were considerably more modular than the plexus model (confirmed by a T-test, $p = 6.24 \times 10^{-17}$). Visualisations of the combined topologies were used to ensure that

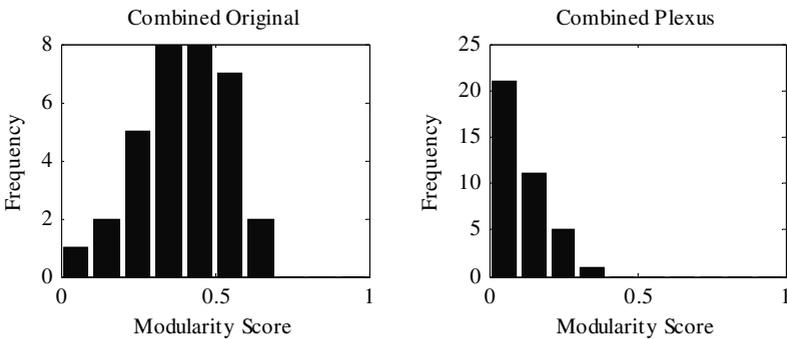


Fig. 3. Distribution of modularity scores for the combined electrical and gaseous topologies for the original and plexus models

the results of the modularity test did appeal to our own notions of modularity. For example, it was noticed 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. The examples shown in Figure 4 show a typical, best performing topology from each controller type.

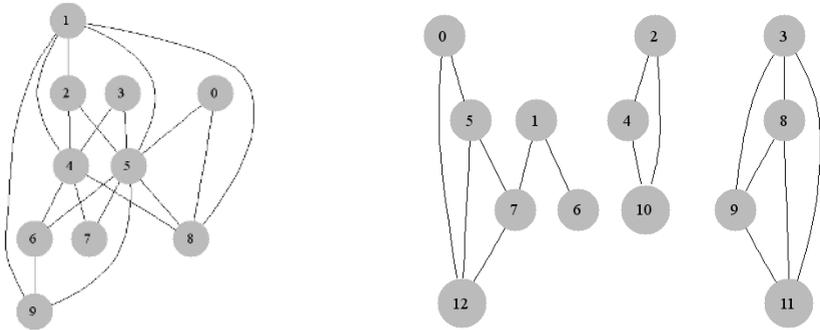


Fig. 4. Example combined topologies from the Plexus model (left), and the original GasNet (right). Generated by the GraphViz package [4].

Additionally, a statistically significant decrease in modularity for the plexus case was found when considering the gaseous interaction topologies alone, whereas the electrical connections were found to be statistically similar. This is to be expected, since the electrical systems were the same in both cases. The gaseous connectivity, however, was only spatially constrained in the original GasNet, which is likely to have led to the increased level of gaseous modularity in that case. It would seem that networks which are *partially* spatially embedded have evolved topologies with a lower degree of modularity than fully embedded networks. These more interconnected networks were considerably more adept at evolving a suitable controller to complete the triangle/square detection task.

5 Discussion

Section 3 has demonstrated that spatially embedded neural networks are considerably more able to generate a modular solution to a minimally modular scenario than non-spatial neurocontrollers. Spatially constrained networks could thus provide a useful substrate within which to investigate the emergence of modularity, both in neural systems and potentially other forms of network. The form taken by the modules is, in some respects, free of assumptions imposed by the researcher. The use of pre-wired modules or specific, ‘modularity-generating’ mechanisms entails that the resultant modular properties of a network are liable to show a degree of bias in their configuration. Within a spatial network, however, the structure of any modules which ensue emerge from the same mechanisms which generate all aspects of the network’s

connectivity, and so could be used to study the emergence and form of modular architectures in general.

That said, the modules which form a spatially embedded network are manifestly structural constructions, rather than dynamical in nature. Whilst the physical topology of a network will undoubtedly influence its dynamical properties, structure does not tell the whole story. It is possible that a structurally modular connectivity could prove to be less flexible than a modularity which is determined by dynamics. One such example (Izhikevich, [9]) examines groups of neurons which are defined by stable, repeatable patterns of firing in large, randomly structured networks. Neurons can be a part of more than one dynamically determined ‘module’, and many more modules can occur than the number of neurons in a network. Structural modularity is nevertheless clearly an important feature to investigate, and biologically relevant, relating to cortical column work [7] amongst others. However future work should additionally investigate dynamical properties, possibly in relation to structural constraints. A measure of *dynamical* modularity would aid such a study.

Section 4 investigated the association between spatiality and modularity within a more complex scenario, including features such as a variable-length genotype and a requirement for embodied, situated perception and action in a simulated environment. In one experiment, this used a network with a similar level of constraint found in the spiking model. Another version used a looser kind of spatial constraint. Firstly, it was apparent that within this more realistic simulation, the fully embedded controllers produced highly modular networks. This was demonstrable using a measure of structural modularity, the results of which matched our intuitions gained from counting the number of components found in each network structure.

It turned out, however, that these highly modular solutions were considerably less evolvable than those of the other, partially embedded neurocontroller. Whilst modularity is clearly a desirable property in many cases, it seems that here too much inherent modularity could be a negative influence on the development and behaviour of the networks. The ‘plexus’ model may correspond to a compromise between spatially modular, and flexibly non-modular structural architectures. Of course, the reduction in modularity found in these networks could impact on the plexus model’s capacity to evolve modular properties when highly modular controllers are required to satisfy a task. It would be interesting to determine whether they could, in fact, produce more modular topologies in a different experimental scenario which benefited from such an architecture.

This compromise could appeal to notions of near-decomposability (Simon, [15]), which discusses modularity in terms of subsystems which behave in distinct ways, but whose interdependencies are also of importance in describing a system’s behaviour. Simon [15] describes these different levels of influence in terms of different timescales of interaction, which is not something that can be examined using the purely structural metric employed in this paper. More suitable tools would be required to make progress understanding in more general terms how different degrees of modularity arise from systems partially embedded in space.

Acknowledgements. This work has been supported by the Spatially Embedded Complex Systems Engineering (SECSE) project, EPSRC grant no EP/C51632X/1.

References

1. Calabretta R., Nolfi S., Parisi D. & Wagner G.P. (1998). Emergence of functional modularity in robots. From Animals to Animats V: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior. MIT Press, pp. 497-504.
2. Cliff, D., and Miller, G. F. (1995). Co-evolution of pursuit and evasion II: Simulation methods and results. From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior Cambridge, MA: MIT Press. pp. 506-515.
3. Di Paolo, E. (2003). Evolving spike-timing-dependent plasticity for single-trial learning in robots. Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences, 361(1811), pp. 2299-2319.
4. Gansner, E. R. & North, S. C. (1999). An Open Graph Visualization System and its Applications to Software Engineering. Software. Practice and Experience, 00(S1), 1-5.
5. Gerstner, W., Kreiter, A. K., Markram, H. & Herz, A. V. M. (1997). Neural codes: ring rates and beyond. Proc. Natl Acad. Sci. USA 94, pp. 12740-12741.
6. Guimera, R., Sales-Pardo, M. and Amaral, L. A. N. (2004). Modularity from Fluctuations in Random Graphs and Complex Networks. Phys. Rev. E 70, 025101.
7. Hubel, D. H., Wiesel, T. N. (1977). Functional architecture of macaque visual cortex. Proc. Roy. Soc. (London) 198B, pp. 1-59.
8. Izhikevich E. M. (2004). Which Model to Use for Cortical Spiking Neurons? IEEE Transactions on Neural Networks, 15. pp. 1063-1070.
9. Izhikevich, E. M. (2006). Polychronization: Computation With Spikes. Neural Computation 18, pp. 245-282.
10. Kitano, H. (1990), Designing neural networks using genetic algorithms with graph generation system. Complex Systems Vol. 4. pp. 461-476.
11. Newman, M. E. J. & Girvan, M. (2004). Finding and evaluating community structure in networks. Physical Review E 69, 026113.
12. Philippides, A., Husbands, P., Smith, T. and O'Shea, M. (2005). Flexible Couplings: Diffusing Neuromodulators and Adaptive Robotics. Artificial Life, 11(1&2):139-160.
13. Philippides, A., Ott, S., Husbands, P., Lovick, T. and O'Shea, M. (2005). Modeling co-operative volume signaling in a plexus of nitric oxide synthase-expressing neurons. Journal of Neuroscience 25(28). pp. 6520-6532.
14. Polani, D., Dauscher, P. and Uthmann, T. (2005). On a Quantitative Measure for Modularity Based on Information Theory. Advances in Artificial Life: 8th European Conference, ECAL 2005, Canterbury, UK. pp. 393. Springer Berlin / Heidelberg
15. Simon, H.A. (1969). The Sciences of the Artificial, Cambridge, MA. MIT Press.
16. Tessier-Lavigne, M. & Goodman, C. S. (1996). The molecular biology of axon guidance. Science. Nov 15;274(5290). pp. 1123-33.
17. Watson, R. A. and Pollack, J. B. (2005). Modular Interdependency in Complex Dynamical Systems. Artificial Life 11(4). pp. 445-457.
18. Yamauchi, B. and Beer, R.D. (1994). Integrating reactive, sequential and learning behavior using dynamical neural networks. From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior. pp. 382-391. MIT Press.