

A study of GasNet spatial embedding in a delayed-response task

Patricia A. Vargas, Ezequiel A. Di Paolo and Phil Husbands

Centre for Computational Neuroscience and Robotics (CCNR)
Department of Informatics
University of Sussex
Falmer, Brighton
BN1 9QH, United Kingdom
{p.vargas, ezequiel, p.husbands}@sussex.ac.uk

Abstract

GasNet artificial neural networks can be used as complex neurocontrollers involving virtual chemical neuromodulation as well as synaptic interaction. The aim of this paper is to further explore the role of space in GasNet models on a delayed-response robot task. Comparative results demonstrate that the use of spatial constraints is not a prerequisite for a good performance of the original model in terms of speed of evolution.

Introduction

Evolutionary robotics allows us to explore complex dynamical neural processes and architectures that connect to interesting issues in neuroscience (Nolfi and Floreano, 2004). The GasNet models can be considered as examples of such complex neurocontrollers involving chemical neuromodulation as well as synaptic interaction (Husbands, 1998). A recently devised non-spatial GasNet model named NSGasNet (Vargas et al., 2007) follows the same principles and had already been successfully applied as a robot controller where the task did not require the controller to have a non-reactive response (Moioli et al., 2008). This work attempts to further explore this novel model in a delayed-response robot task, in addition to compare it with the original GasNet model in terms of evolvability. In essence, our aim is to investigate whether the space embedding present in the original GasNet is the main explanation for its success when applied to more elaborate robot tasks.

We will present a comparison which follows the investigation started by Vargas and collaborators (Vargas et al., 2007). In that work, the NSGasNet model has proven to have higher evolvability with respect to the original model on a central pattern generator task (CPG).

However, it is unclear whether the conclusions obtained for the CPG task will carry over to more complex situations, especially to cases involving an embodied agent. For this reason, we decided to perform a comparative study on a well-researched delayed-response task involving a T-Maze (Husbands, 1998; Jakobi, 1993, 1997; Ulbricht, 1996).

Our results corroborate the fact that the use of spatial embedding is not a prerequisite for better performance either

in terms of speed of evolution or in robustness. This might also indicate that the success demonstrated by GasNet models so far (Husbands et al., 1998; McHale and Husbands, 2004; Philippides et al., 2005) are not related to the spatial embedding of nodes but maybe to the temporal dynamics promoted by the gaseous diffusion amongst them.

We will start by briefly describing the original GasNet plus the novel model, together with a summary of the previous results on a CPG task. Thereafter, we will describe our experiment in detail including the respective network architecture and genetic encoding, together with the evolutionary regime. After the results section we will provide a discussion and propose future work.

Non-Spatial GasNet: NSGasNet

Since the introduction in 1943 of the first artificial neuron model proposed by McCulloch and Pitts (McCulloch and Pitts, 1943) most of the subsequent classical artificial neural networks (ANNs) architectures have employed numerical synaptic interaction between their neurons. However, recent findings in neuroscience have suggested the existence of chemical signaling by gases that would play the role of neurotransmitters (Gally et al., 1990). By drawing inspiration from these latest discoveries, the GasNet model was introduced by Husbands (1998) in an attempt to create a novel recurrent artificial neural network, which seeks to combine the electrical and chemical signaling onto a single network.

In the original GasNet model, the classical sigmoided output function $y = \tanh(x)$ of each neuron at each time step is modulated by a transfer function parameter k which will define which curve from the family of eleven sigmoids ($x = [-4, 4]$) will be employed during the network's operation. The value of k is controlled by the concentration of diffusing transmitter gas at a node following the network dynamics dictated by the network's equations as described in Husbands (1998).

Almost all GasNet parameters and variables are under evolutionary control. The use of evolutionary computation techniques to evolve ANNs is of fairly recent origin (Signals et al., 1990; Whitley et al., 1990; Yao and Liu, 1997; Yao,

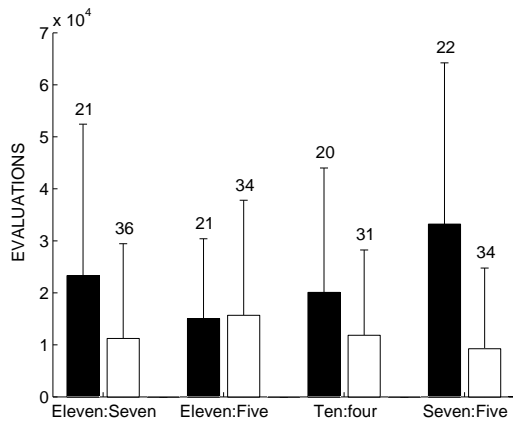


Figure 1: Mean and standard deviations (error bars) of fitness evaluations required to evolve successful networks for each CPG pattern, Eleven-Seven, Eleven-Five, Ten-Four and Seven-Five. Black bar shows original mean data and white bar shows NSGasNet mean data. The numbers above each error bar represent the total number of successfully evolved networks within 50 runs (adapted from (Vargas et al., 2007))

1999). Following this initiative, GasNet models were particularly designed to "evolve" for every task addressed. Hence, the network size, topology and almost all its parameters are under unconstrained evolutionary control.

Normally, depending on the task, the network is composed of a variable number of nodes. Thus, a network is encoded on a variable-sized genotype, where each gene represents a network node. A gene consists of an array of integer variables lying in the range [0, 99] (each variable occupies a gene locus). The decoding from genotype to phenotype obeys simple laws for continuous values and for nominal values (Husbands et al., 1998).

Vargas et al. (2007) introduced a novel spatially unconstrained GasNet named NSGasNet, in which the nodes do not have a location in a Euclidean space. Reminiscent of how the gas neurotransmitter NO normally diffuses once released (Gally et al., 1990; Wood and Garthwaite, 1996, 1994), in the NSGasNet model all emitted gases can spread freely among neurons.

NSGasNet is a discrete time recurrent neural network, which could be fully or partially connected with fixed or variable number of nodes. This full or partial connectivity refers to the synaptic connections. The gaseous connections are defined in terms of sensitivity limits, which impose to each network node a filter that regulates the strength of gas modulation (Vargas et al., 2007). Thus, each node has its set of sensitivity limits lying in the range [0, 1] of which values correspond to each node within the network. Although the NSGasNet has a bias that modulates the concentration of the gas at each node, the rules for how and when the gas

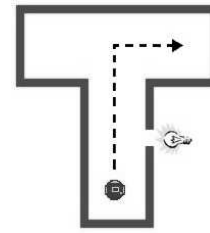


Figure 2: Schematic drawing of the robot and the T-Maze environment with two corridors. The robot is represented by the small circle and it is positioned in the bottom of the first corridor facing north. On the right-hand side there is a beam of light.

is emitted are the same as the original GasNet (Husbands, 1998).

In a previous work by (Vargas et al., 2007), this non-spatial model has been successfully applied to a CPG task where the network should evolve to generate a sequence of cyclic output values from the set 0,1. Four patterns were tested and in all of them the NSGasNet was demonstrated to outperform the original spatially constrained GasNet Model in terms of speed of evolution (Figure 1). Some preliminary statistical analysis around mutants was performed to investigate the possible reasons for the best performance hypothesising about the role of the fitness landscape smoothness. A more profound analysis has been carried out in another work using further statistical correlation analysis between both models and the results will be submitted to publication soon. This work on the other hand intends to apply both network models to a more elaborate robot task to further assess the role of space in the performance of GasNet models.

Methods: T-Maze and Evolutionary Regime

T-Maze with light task

The experiment is a delayed response task in which a robot must learn to negotiate a T-Maze turning at the junction in the correct direction after passing a beam of light located in the first corridor either to the left or the right (Figure 2). Therefore, the robot must 'remember' the position of the light in order to successfully accomplish the task. This task, and similar ones, have been used by various researchers to endow artificial agents with minimal memory mechanisms (Husbands, 1998; Jakobi, 1993, 1997; Lanzi, 1998; Ulbricht, 1996; Webb et al., 2003); in this context it is interesting to note that it is still not well understood how biological memory works (Wilson, 1994; De Zeeuw, 2005; Levenson, 2006).

For this task we make use of a dedicated 2D robot simulator (Figure 3) of a Khepera II robot for the evolution of the GasNet models. The Khepera II robot has two wheels and two separate motors, 8 infra-red distance sensors (6 on the

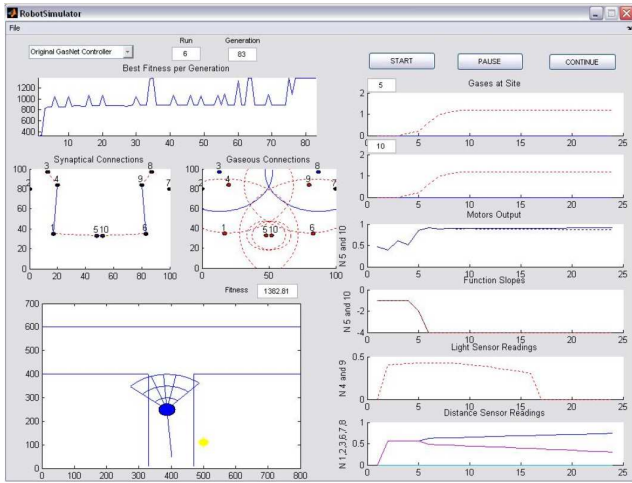


Figure 3: GUI of the robot simulator especially designed for the T-Maze delayed response experiment. On the left, from top to bottom, the interface shows a list to choose the GasNet model of interest (e.g. original or NSGasNet) and the specific run and generation; the best fitness per generation, the network architecture in terms of synaptic and gaseous connections and the robot within its arena. The right side of the interface shows in time from top to bottom: the values of the gas at site for two chosen nodes together with their function slopes, the values of the motor outputs, the values of the light sensor reading, and the values of the distance sensor reading

front and 2 on the rear) and 8 infra-red light sensors (6 on the front and 2 on the rear) .

The robot implemented in our simulator is a simplified model of a Khepera II robot and it was employed to avoid the overloading of graphical encoding in order to speed-up the simulations. It has 5 front distance sensors and 2 almost diametrically opposite light sensors (Figure 4(a)). While implementing the simulator, the two original front-most distance sensors were coupled (Figure 4(b)); hence, both sensor readings have the same value during the simulation. This was due to observations made during the design phase of the simulator where both sensors readings presented the same values most of the time.

Network Architecture, Genetic Encoding and Evolutionary Regime

Both GasNet models, original and NSGasNet, were implemented with a fixed number of nodes (total of 10 nodes). The networks are partially connected in addition to having genetically determined recurrent connections (Figure 5). Nodes 1, 2, 3, 6, 7 and 8 have input from the robot distance sensors S1, S2, S3, S5, S4 and S3, respectively. Nodes 4 and 9 have input from the left (L1) and right (L2) light sensors, respectively (Figure 8). Nodes 5 and 10 are responsible

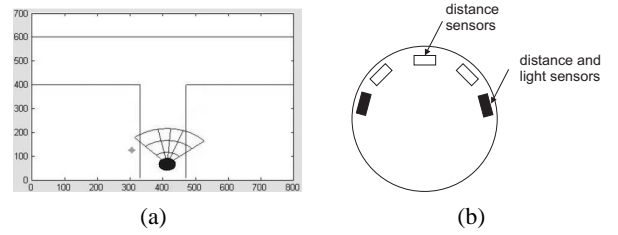


Figure 4: (a) Zoom of the T-Maze arena and the simulated robot (black round shape) localized at the bottom of the first corridor, facing north, and its five distance sensors stressing their range. Distance sensors were numbered from left to right: S1, S2, S3, S4 and S5 and light sensors: L1 - on the left side and L2 - on the right side. The arena is composed of two corridors forming a T-Maze and there is a beam of light (shaded star shape) shining from the left side of the robot. (b) Presents a schematic of the same robot illustrating the disposition of the distance and light sensors.

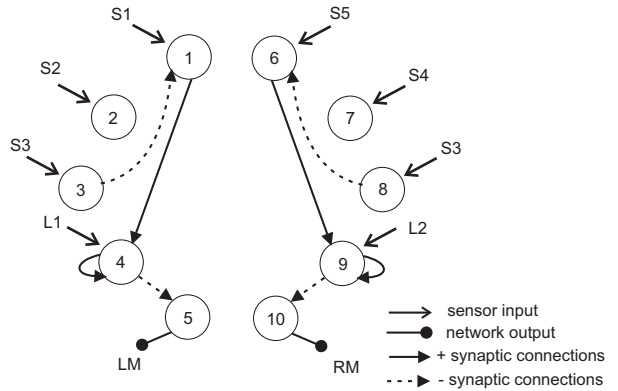


Figure 5: Pictorial example of a symmetrical partially connected ANN for the T-Maze task with ten nodes. The network receives external input from the sensors and supplies output to the motors.

for the output to the robot motors (left (LM) and right (RM) wheels, respectively).

Both networks have a symmetrical architecture meaning that for the genetic encoding we will only have to evolve half of the network. Hence, the original GasNet gene will have 65 parameters for the entire network, i.e. 13 parameters *times* 5 nodes. Each node is coded as follows: $\langle gene \rangle = \langle node \rangle = \langle x \rangle, \langle y \rangle, \langle x1 \rangle, \langle y1 \rangle, \langle x2 \rangle, \langle y2 \rangle, \langle rec \rangle, \langle Es \rangle, \langle Gt \rangle, \langle s \rangle, \langle Gr \rangle, \langle k0 \rangle$ and $\langle bias \rangle$, where $\langle x \rangle$ and $\langle y \rangle$ are the node coordinates on the plane; $\langle x1 \rangle, \langle y1 \rangle, \langle x2 \rangle, \langle y2 \rangle$ specify the center of two circles on the network plane defining the node spatial electrical connectivity; $\langle rec \rangle$ is the recurrent status; $\langle Es \rangle$ is the emitting status; $\langle Gt \rangle$ is the gas type; $\langle s \rangle$ is the build up/decay rate; $\langle Gr \rangle$ is the gas maximum radius of

Parameter	T-Maze task
Mutation rate	8%
Fitness function	$Fitness_{T-Maze} = d1 + d2 + bonus$
Number of runs	40
Maximum number of generations	150
Population size	100
Genotype size	65 (Original) 100 (NSGasNet)
Trials	10
Number of evaluations per trial	[70, 100]

Table 1: Evolutionary regime parameters employed on the T-Maze task.

emission, $\langle k0 \rangle$ is the transfer function default value and $\langle bias \rangle$ is the bias value (Husbands, 1998).

The NSGasNet genotype does not have parameters related to node coordinates, spatial electrical connectivity and maximum radius of emission, thus each NSGasNet gene will have 6 parameters for each node plus 10 values for the NSGasNet sensitivity limits (10 nodes), which makes the totality of 80 parameters for the entire network, plus 2 times the maximum number of allowed synaptic connections per node (to include the node number and the synaptic connection weight). For instance, if the maximum allowed number of synaptic connections per node is 2, then the NSGasNet genotype will have $80 + 5(2(2)) = 100$ variables.

The choice of partially interconnected networks for this task follows from previous works (Psujek et al., 2006; Williams and Noble, 2006) and also from the preliminary experiments on the T-Maze task where it was observed that full connectivity produced a negative impact on the evolvability of the networks for this particular task. The fully connected networks were too sensitive to genetic operations and initial conditions (e.g., the starting angle of direction) during the evolutionary process; therefore, a successful controller from one evaluation could hardly repeat its performance on the next fitness evaluation.

We employed a distributed steady-state genetic algorithm as described in (Husbands et al., 1998), who developed the idea from an early work using distributed populations (Hillis, 1990). The current population is updated steadily during the evolutionary process, i.e. each offspring is placed immediately into the current population (Whitley et al., 1990), instead of an entirely new population being generated and replacing the current population at a single time. Offspring were created through mutation operators (no recombination was used) with a probability of (8%) for each gene locus following a Gaussian distribution around its value for non-nominal values and a random value for nominal values. Non-nominal values refer to variables that have continuous values and nominal for discrete values.

In order to gather statistics 40 runs were performed for each model. One evolutionary run is composed of a maxi-

mum of 150 generations, or until successful genotypes are produced. Each generation comprises of 100 reproduction events or fitness evaluations.

The robot is tested for ten trials. Each trial is divided into two phases, following Jakobi’s experiments set-up (Jakobi, 1997). The fitness value for phase 1 accounts for the distance $d1$ traveled by the robot in the first corridor ($d1_{max} = 200$) and the fitness for the phase 2 is composed of the distance $d2$ traveled in the second corridor ($d2_{max} = 180$) plus a bonus if the robot turns to the correct direction. The total fitness is the sum of the fitness at each trial divided by the total number of trials. The only difference from Jakobi’s fitness calculation is the bonus value, which is computed as follows during the trials:

- 200 if the robot has turned to the correct side once;
- 500 if the robot has turned an equal number of times to both sides, plus:
 - +200 if the robot has turned four times to one side and four times to the other side
 - +500 if the robot has turned five times to one side and five times to the other side

Therefore the maximum fitness has a value around 1,380 according to 1. This new bonus scheme was devised for it was observed that the evolution was very sensitive to the bonus criteria which imposes a selection pressure. Possibly this change was due to not implementing Jakobi’s minimal simulations schema. Basically, this schema encompasses the addition of a controlled degree of noise and uncertainty during the evolution which will lead the robot to an improved robust behaviour when transferred to the reality. However, in these first robot experiments we are not concerned with the reality gap but with the measure of the evolvability of each GasNet model under noiseless circumstances. Therefore, we do not add noise to our simulations, just the start directional angle of the robot varies from trial to trial.

$$Fitness_{T-Maze} = d1 + d2 + bonus \quad (1)$$

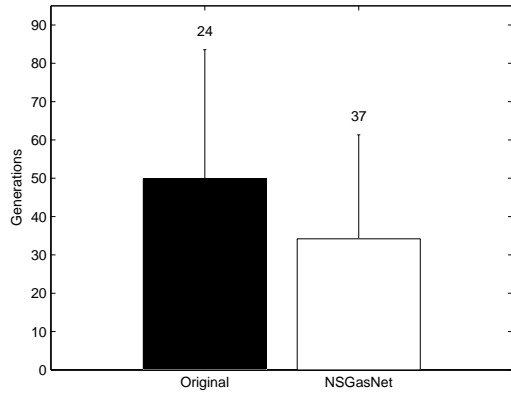


Figure 6: Mean and standard deviations (error bars) of generations required to evolve successful controllers for T-Maze with light task. Black bar shows original mean data and white bar shows NSGasNet mean data. The numbers above each error bar represent the total number of successfully evolved controllers within 40 runs.

Table 1 summarizes the parameter settings implemented within the evolutionary regime.

The successful evolution of a controller is considered if the robot obtains a fitness value that is greater than a threshold of 1,260 over seven subsequent trials. A robot with such fitness value has received the maximum *bonus* = 1,000 for having turned correctly in all 10 trials, plus the minimum distances traveled in both corridors, which when added may vary between [260, 380]. Runs that exceed 150 generations were aborted.

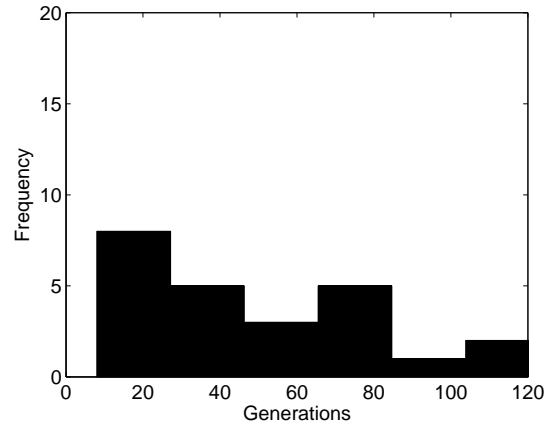
Results

The statistical results over 40 runs for each model are graphically illustrated in Figure 6. Black bar shows original mean data and white bar shows NSGasNet mean data. The numbers above each error bar represent the total number of successfully evolved controllers within 40 runs.

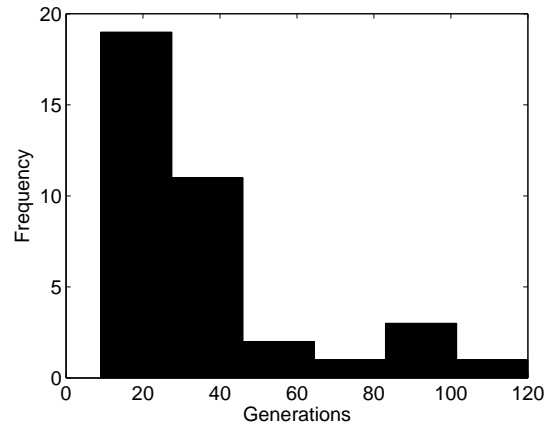
The NSGasNet outperforms the original GasNet model in terms of number of successful runs. The frequency histograms portrayed at Figure 7 show that both distributions are skewed to the right, thus not symmetric, the difference between the mean and the median tend to spot a similar performance in terms of speed of evolution for the robot task between both models. However, the percentage of successfully evolved networks for the NSGasNet ($37/40 = 92\%$) is greater than the original ($24/40 = 60\%$).

Concerning the network architecture, in contrast to the NSGasNet model, the original model evolved less synaptic connections and more gaseous connections (Figure 8).

Many NSGasNet networks and some original ones did not make use of gases in the final evolved solution. When gases



(a)



(b)

Figure 7: Frequency histograms comparison between the original (a) and the NSGasNet (b) models over the number of generations for the T-Maze task.

were at play, normally the nodes connected to the robot sensor lights had a coupled gaseous connection. Therefore, both nodes were making explicit use of gases to control the dynamics of each other and/or of other network nodes in response to environment changes, e.g. source of light (Figure 9).

An analysis of the behaviour of the robot shows that some of the successfully evolved controllers developed a reactive response to the task. For instance, the robot starts to follow the wall after passing the beam of light and thus, the robot is using the wall as an external memory, instead of creating an internal memory based on its internal state (Braitenberg, 1986; Nolfi, 2002). Naturally, this observation does not invalidate our evolvability results of the GasNet models. It only sheds some light on the potential requisite for an improved way to assess the robot behaviour during evolution, possibly in terms of a more elaborated fitness function.

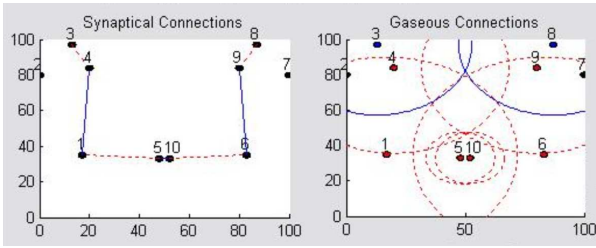


Figure 8: Picture of the simulation of an original GasNet successfully evolved controller, highlighting its synaptic and gaseous connections. There are few synaptic connections and intricate gaseous connections.

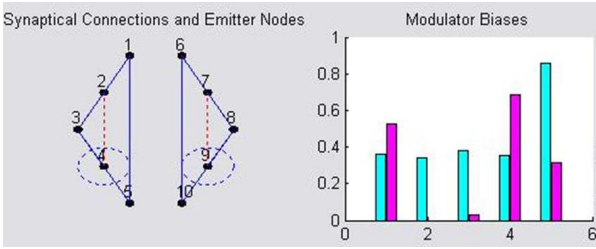


Figure 9: Screenshot of the simulation of a NSGasNet successfully evolved controller, highlighting its synaptic connections on the left and NSGasNet bias values for nodes 1, 2, 3, 4, and 5 on the right (the bars refer to nodes 4 and 9, respectively). Only nodes 4 and 9 are gas emitters. Remember that these are the nodes directly connected to the light sensors.

Discussion

This paper is a further step on the investigation of a novel non-spatial GasNet model (NSGasNet) in an attempt to uncover the role of space within this neural network paradigm. The performance of the original and the NSGasNet model was explored on a memory robot task. Unlike the previous results on a CPG task, the comparison between both models showed little difference in terms of speed of evolution. Although the evolvability values are quite similar they differ in the percentage of evolved controllers meaning that the NSGasNet has a higher success rate. Nonetheless, further analysis should be carried out in order to further assess this better performance.

Additional remarks could be made from the experiments. For instance, the use of partial connection between nodes was adopted in both models for the fully connected networks were too sensitive to genetic operations and initial conditions (e.g. the starting angle of direction) during the evolutionary process. Therefore, a successful controller from one evaluation could hardly repeat its performance on the next fitness evaluation, thus compromising its speed of evolution. One may argue that the problem might be the elevated muta-

tion rate adopted. However, many mutation rates were tested and no improvement was observed. Thus, in order to make a compromise between evolvability and good performance, apart from the partial connection, we adopted 8% for the mutation rate.

It was observed that after evolution, some nodes either had their synaptic weights set to zero or there were no gaseous connections whatsoever. This fact shows the ability of the evolutionary process to find simple solutions to the problem and it also indicates that the introduction of metadynamics could improve the results. Metadynamics in this context means exploring a variety of network's dimensions during the evolutionary process. Therefore, in a future work we envisage using not only partially connected networks, but also exploring the network metadynamics. In our opinion, which is shared by others (Pszujek et al., 2006), this coupling might lead to superior results.

According to (Strogatz, 2001) realistic networks have both nontrivial node dynamics and specific but irregular connection topologies. Moreover, highly distributed and non-hierarchical neural circuits had been identified in neuroscience investigations of simple organisms as pointed out by (Altman and Kien, 1990) and stressed by (Beer, 1995). Likewise, an analysis of the resulting network architectures for the T-Maze task has demonstrated a huge variety of topologies of connections (synaptic and gaseous) among the evolved controllers. This enormous variety was also verified by Vargas et al. (2007) for the CPG task. In both cases, it was impossible to identify a predominant pattern of connections and/or of spatial location of the nodes (in the case of the original GasNet model).

Internal state is not a pre-requisite for the agent to perform sophisticated interactions with the environment, as pointed out by (Izquierdo and Di Paolo, 2005; Nolfi, 2002; Stanley and Miikkulainen, 2002; Ziemke and Thieme, 2002). Accordingly, the fact that some of the robots presented a reactive response to the T-Maze task seems to indicate that the chosen task does not require a non-reactive response in order to be successfully accomplished.

In conclusion, the results obtained on this work together with the first investigations presented by (Vargas et al., 2007) seem to indicate that the explicit use of spatial constraints and a spatially embedded diffusion process is not necessary to explain the success of GasNet models. Rather, the interplay between two distinct processes (electrical signals and gas modulation) acting on different timescales, and the multiplicative modulation effect of the gases appear to be the important factors (Philippides et al., 2005).

In order to fully clarify the role of space within GasNet models, future work should include an analysis of the performance of both models in other tasks that require networks with higher dimension.

Acknowledgements

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

References

- Altman, J. S. and Kien, J. (1990). Highlighting aplysia's networks. *Trends in Neuroscience*, 13 (3): 81-82.
- Beer, R. D. (1995). A dynamical system perspective on agent-environment interaction. *Artificial Intelligence*, 72:173-215.
- Braitenberg, V. (1986). *Vehicles, Experiments in Synthetic Psychology*. MIT Press.
- De Zeeuw, C. I. (2005). Time and tide in cerebellar memory formation. *Current Opinion on Neurobiology*, 15: 667-674.
- Gally, J. A., Montague, P. R., Reeke, G. N., and Edelman, G. M. (1990). The no hypothesis: Possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. In *Proceedings of the National Academy of Sciences of the USA*, pages 87, 3547-3551.
- Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42:228-234.
- Husbands, P. (1998). Evolving robot behaviours with diffusing gas networks. In *Evolutionary Robotics: First European Workshop, EvoRobot98*, pages 71-86. Springer-Verlag.
- Husbands, P., Smith, T., Jakobi, N., and O Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 10:185-210.
- Izquierdo, E. and Di Paolo, E. (2005). Is an embodied system ever purely reactive? In M. Capcarrere et al (Eds.) *Proceedings of the 8th European Conference on Artificial Life*. pages 252-261. Springer-Verlag. LNAI 3630.
- Jakobi, N. (1993). Half-baked, ad-hoc, and noisy: Minimal simulations for evolutionary robotics. In Husbands, P., and Harvey, I., (Eds.) *Proc, Fourth European Conference on Artificial Life*, MIT Press. Kaufmann, S. A. (1993). *The Origins of Order*. Oxford University Press., pages 348-357.
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adapt. Behav.*, 6(2):325-368.
- Lanzi, P. L. (1998). Adding memory to xcs. In *Proceedings of the IEEE Conference on Evolutionary Computation (ICEC 98)*, pages 609-614. IEEE Press.
- Levenson, J. M. (2006). Epigenetic mechanisms: a common theme in vertebrate and invertebrate memory formation. *Cellular and Molecular Life Sciences*, 63:1009-1016.
- McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. 265:115-133.
- McHale, G. and Husbands, P. (2004). Quadrupedal locomotion: Gasnets, ctrnns and hybrid ctrnn/ pnnns compared. In In J. Pollack, M. Bedau, P. Husbands, T. Ikegami and R. Watson, editors, *Alife X Proc. of the 9th Int. Conf. on the Simulation and Synthesis of Living Systems*, pages 106-112. MIT Press.
- Moioli, R. C., Vargas, P. A., Zuben, F. J. V., and Husbands, P. (2008). Towards the evolution of an artificial homeostatic system. In *World Conference on Computational Intelligence 2008, WCCI'2008, Hong Kong, China*.
- Nolfi, S. (2002). Power and limits of reactive agents. *Neurocomputing*, 42 (1): 119-145.
- Nolfi, S. and Floreano, D. (2004). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Bradford Book.
- Philippides, A., Husbands, P., Smith, T., and O'Shea, M. (2005). Flexible couplings: Diffusing neuromodulators and adaptive robotics. *Artificial Life*, 11:139-160.
- Psujek, S., Ames, J., and Beer, R. D. (2006). Connection and coordination: The interplay between architecture and dynamics in evolved model pattern generators. *Neural Comput.*, 18(3):729-747.
- Signals, M., Fogel, S., Fogel, D., and Porto, L. (1990). Evolving neural networks. *Biological Cybernetics*, 63, pages 487-493.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99-127.
- Strogatz, S. H. (2001). Exploring complex networks. *Nature*, 510:268-276.
- Ulbricht, C. (1996). Handling time-warped sequences with neural networks. In *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, MIT Press, 1996.*, pages 180-189.
- Vargas, P. A., Di Paolo, E. A., and Husbands, P. (2007). Preliminary investigations on the evolvability of a non-spatial GasNet model. In *Proc. of the 9th European Conf. on Artificial life ECAL 2007*, pages 966-975. Springer-Verlag.
- Webb, A., Hart, E., Ross, P., and Lawson, A. (2003). Controlling a simulated khepera with an xcs classifier system with memory. In *ECAL*, pages 885-892.
- Whitley, D., Starkweather, T., and Bogart, C. (1990). Genetic algorithms and neural networks: Optimizing connections and connectivity. *Parallel Comput.*, 14:347-361.
- Williams, H. and Noble, J. (2006). Homeostatic plasticity improves signal propagation in continuous-time recurrent neural networks. *Biosystems*, September 9:252-259.
- Wilson, M. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265.
- Wood, J. and Garthwaite, J. (1994). Models of the diffusional spread of nitric oxide: Implications for neural nitric oxide signaling and its pharmacological properties. *Neuropharmacology*, 33:1235-1244.
- Wood, J. and Garthwaite, J. (1996). Diffusion of free nitric oxide. *Methods in Enzymology*, 268:31-50.
- Yao, X. (1999). Evolving artificial neural networks. *Proc. of the IEEE*, 87:1423-1447.

Yao, X. and Liu, Y. (1997). A new evolutionary system for evolving artificial neural networks. *IEEE Transactions on Neural Networks*, 8(3):694–713.

Ziemke, T. and Thieme, M. (2002). Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, Vol. 10 No.3/4:175–99.