Preliminary Investigations on the Evolvability of a Non-spatial GasNet Model

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

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

Abstract. This paper addresses the role of space in evolving a novel Non-Spatial GasNet model. It illustrates that this particular neural network model which make use of modulatory effects of diffusing gases has its evolvability improved when its neurons are not constrained to a Euclidean space. The results show that successful behaviour is achieved in fewer evaluations for the novel unconstrained GasNet than for the original model.

1 Introduction

Our aim in this work is to explore and analyze the role of spatial organization and processes in a particular class of non-classical artificial neural network (ANN) called GasNets. We embark on this endeavour by proposing and investigating the evolvability [1] of a novel spatially unconstrained version of this network model, which we call NSGasNet (Non-Spatial GasNet). We expect that this investigation may have wider implications for understanding the interactions between space, neuro-modulation, temporal dynamics and evolvability in several systems apart from GasNet models.

The original GasNet neurocontroller uses a simple model of a diffusing gas that may affect the transfer function of nodes within a delimited spatial range. Such a model has proved significantly more evolvable for certain robotic tasks than non-gaseous neurocontrollers [2, 3]. Further developments of this model have led to even higher evolvability [4, 5, 6]. Hypotheses concerning this increased evolvability have been discussed, including the combined role of slow and fast temporal timescales and the flexible coupling between distinct processes in the networks, but still no clear answer has been found to explain the success of this model.

There are good reasons to believe that the inherent spatiality of the original model may play a significant role in enhancing evolvability. Spatiality introduces a topology where variations produced by mutation events may be smoother and it may enhance the exploration of modular architectures [7]. The effects, however, are far from linear or straightforward. For instance, the Plexus GasNet model [4, 5] has introduced a spatial decoupling between synaptic and gaseous interaction, and the freeing up of this constraint has resulted in enhanced evolvability [6].

The objectives of the present investigation are to begin a series of systematic statistical explorations into the role of space in GasNets. We will introduce some comparisons between a spatial and a non-spatial version of GasNets for a central pattern generator task (a CPG task). The CPG was chosen as they avoid very simple regularities

F. Almeida e Costa et al. (Eds.): ECAL 2007, LNAI 4648, pp. 966 – 975, 2007.

and have significant timescale properties. Some generalization to tasks evolved under similar conditions and of similar complexity might be expected.

We start by giving a brief account of existing versions of GasNets plus a novel model in Section 2. In Section 3 we describe our experiments in detail including the respective network architecture and genetic encoding. Section 4 highlights the evolutionary regime. The results for the task are presented in Section 5 and in Section 6 we provide a discussion of the main findings and future directions.

2 Non-classical Artificial Neural Networks

At the beginning of the past decade, with the advent of remarkable discoveries of neuronal modulation and non-standard cell-to-cell signaling in the biological nervous system [8], a novel ANN model, namely GasNet, was proposed by Husbands [2]. This particular architecture could be considered non-classical as it involves non-synaptic chemical signaling as well as synaptic interactions. The network is conceptualized as operating on a 2D Euclidean plane, thus both types of interaction are constrained by a spatial relation in the sense that the synaptic connections and neuronal modulation by diffusing gases are restricted to a spatial range. Hence, all neural connections depend upon a spatial organization. In this work, a new GasNet model is devised in which there is no spatial relation among neurons in order to help us investigate the role played by space.

2.1 Spatially Constrained GasNet Models

The rationale behind this gaseous artificial neural network model is to mimic the production and release of nitric oxide (NO) by real neurons, in order to affect long and short-term modulations of the behaviour of other neurons in (spatial) range. In fact, the original GasNet model is a discrete-time, recurrent neural network with a variable number of nodes. These nodes can be connected in terms of synapses by either excitatory (with a weight of +1) or inhibitory (with a weight of -1) links and, in terms of dynamic gas modulation depending on their spatial relation [3].

In the GasNet model, the classical sigmoid output function $y = \tanh(x)$ of each neuron at each time step is modulated by a transfer function parameter which will define which curve from the family of eleven sigmoids will be employed during the network's operation. Almost all GasNet parameters and variables are under evolutionary control.

2.2 Non-spatial GasNet: The NSGasNet Model

We have devised a novel spatially unconstrained GasNet model (NSGasNet). In this model nodes do not have a location in a Euclidean space. In the absence of a spatial relation, all emitted gases can spread freely among neurons, thus there is no notion of a gas cloud anymore.

In this new scenario, all nodes can in principle be affected by the gas emitted from any other node. Therefore it was envisaged that a *sensitivity limit* should be imposed to each network node in order to regulate the strength of modulation. The *sensitivity limits* are under evolutionary control lying in the range [0, 1] and are specific to each other emitting node. This can be understood as if there are 'gas' connections between nodes of a strength defined by the respective *sensitivity limit*.

In the NSGasNet model, the *sensitivity limit* was named *Mbias* (modulator bias) and its product with the amount of gas emitted $T(t)$ will now determine the gas concentration at the node. Each node will have a modulator bias lying in the range [0,1] for every emitting node. Therefore, given an emitting node, any network node could "decide" whether it will be affected (*Mbias* > 0), or not (*Mbias* = 0) by the gas emitted, without the requirement of being within its gas cloud limits.

The NSGasNet network genetic encoding differs from the original model [3] in the absence of variables that deal with spatial parameters, such as node *coordinates*, *spatial electrical connectivity* and *maximum radius of emission*. Similar to the original GasNet model, the network size, topology and almost all its parameters are also under unconstrained evolutionary control. Nonetheless, each node has a list of modulator biases (one modulator bias for each emitting node).

Depending on the task, the network is encoded on a variable-sized genotype, where each gene represents a network node. A genotype 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 the same simple laws for continuous values and for nominal values adopted for the original model [3]. Apart from the task-dependent parameters, the NSGasNet model has 6 variables associated with each node plus 1 modulator bias for each node. Suppose the network is composed of 4 nodes, than the NSGasNet genotype will have 6+4+task dependent parameters.

In the following sections we will provide a set of comparisons between the original spatially constrained GasNet model and the novel non-spatial NSGasNet model in an attempt to address the role of space in the evolvability of both networks.

3 Comparative Experiment

CPGs are widely known neuronal circuits found in almost all nervous systems from invertebrates [9] to vertebrates [10, 11]. They are used in many biological functions including the production of rhythmic movement found in locomotion in many species [12]. These circuits have been under investigation for a long time [13, 14, 15, 16] in a hope to fully uncover how locomotion patterns, such as swimming, walking and flying are internally organized and coordinated.

The original GasNet model has already been successfully applied to CPG problems [4, 17]. In fact, the experiment described in our CPG task was first suggested by Smith [17]. It is comprised of different target patterns representing sequences of bits from the set {0, 1}. Networks are evolved to generate the required patterns. Four patterns were tested (Table 1).

As a CPG produces a regular cyclic pattern of movement resulting in the animal behaviour usually cycling among a set of states, our idea here is to mimic these states dictated by a set of synthetic patterns to be learned by the ANN.

In this experiment, both GasNet models, original and NSGasNet, were designed as fully connected ANNs (including self-connections) with four nodes (Fig. 1). We decided to start with simple systems in order to try to fully uncover them, furthermore, the task does not require more complex networks and that there is still an important

Table 1. CPG patterns for the CPG task adapted from Smith [17].

Ten:Four	Eleven:Five	Eleven:Seven	Seven:Five
1111111111:0000	11111111111:00000	11111111111:0000000	1111111:00000

number of node arrangements modulated by a set of possible distances that may make the spatial factor an important one.

Apart from each model's particularities, both genetic encodings only differ from the basic GasNet in that the synaptic weights are also under evolutionary control lying in the range [-1, 1], and there are no other electrical connectivity parameters.

Hence, the original GasNet genotype will have 9 parameters for each node, which makes a total of 36 parameters for the entire network plus 6 parameters for the synaptic connection weights. Each NSGasNet gene will have 6 parameters for each node plus 4 parameters for the modulator bias (4 nodes), giving a total of 40 parameters for the entire network plus 6 parameters for the synaptic connection weights.

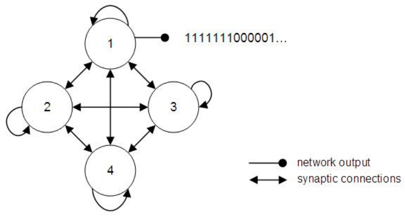

Fig. 1. Pictorial example of a fully-connected ANN for the CPG task with four nodes. The network does not receive external input and the first network neuron output determines the network output.

The choice of fully interconnected networks for this task follows from previous experiments on another CPG task [18, 19]. Nonetheless, in a more recent work, Psujek, Ames and Beer [20] investigated, among other issues, the number of connections required to achieve high performance in a CPG for walking in a simple legged body and stated that "far sparser than fully interconnected circuits are sufficient to achieve high performance on the walking task…".

In this primary investigation, we will adopt a full connectivity, however in future work we envisage using not only partially connected networks, but will also explore the network *metadynamics*, i.e. exploring a variety of network dimensions (including number of nodes) during the evolutionary process, both of which in our opinion might lead to superior results for our CPG task.

For the evolutionary regime, we employ a distributed steady-state genetic algorithm as described in Smith [17]. In order to gather statistics, fifty runs were performed for

Table 2. Statistics on the evolvability in terms of number of fitness evaluations for each GasNet model, original and NSGasNet, over 50 runs for each CPG pattern. The values are presented in the following order: mean, number of successful evolved networks, standard deviation and median.

Pattern	Eleven-Seven	Eleven-Five	Ten-Four	Seven-Five
Original				
Mean/n	23310/21	15048/21	20085/20	33218/22
(Std) Median	(29123) 12200	(15343) 7600	(23924) 7350	(31001) 15500
NSGasNet				
Mean/n	11231/36	15691/34	11845/31	9252/34
(Std) Median	(18209) 3050	(22105) 6050	(16395) 4000	(15523) 3200

the CPG. One evolutionary run is composed of a maximum of 1,000 generations for the CPG task or until successful genotypes are produced. Each generation comprises 100 reproduction events or fitness evaluations.

Each genotype is evaluated over 200 steps. The network output is "1" if the activation of the output node is greater than zero, and "0" otherwise. The network output is compared to the pattern output at each step and a value of one is added to its fitness each time the value is the same. A weighted sum is performed at the end of the steps in order to account for the different numbers of ones and zeros of each test pattern. The fitness is than scaled to [0, 1].

4 Results

Here we present the statistical measures of evolvability for the original GasNet and the NSGasNet in the CPG task. The statistics on the number of fitness evaluations for the original and NSGasNet models for each CPG pattern are shown in Table 2. The first column presents the CPG patterns: Eleven-Seven, Eleven-Five, Ten-Four and Seven-Five (Section 3). The second column presents the statistical measures: mean, number of successful evolved networks, standard deviation and median. The third and fourth columns present the values for each measure for the original and the NSGasNet, respectively.

All values displayed in Table 2 are graphically illustrated in Fig. 2. Frequency histograms comparisons between the two GasNet models over the number of fitness evaluations for each pattern are illustrated in Fig. 3.

A quantitative analysis of the data shown in Fig. 2 demonstrates that the distribution is not symmetric, actually all distributions are skewed to the right. Therefore the mean is always greater than the median (Table 2). Moreover, there is also the presence of outliers (a), (b), (c), (d), (e) and (h) and the large values of standard deviations show that the data has a great variability (Table 2). The median is less sensitive to outliers and therefore we will use the median instead of the mean in order to compare evolvability between the GasNet models for all patterns.

From Table 2, it is clearly observed that the NSGasNet model outperforms the original model, not only in terms of the median as a measure of evolvability but also in terms of the total number of successfully evolved networks.

As the evolved networks were fully connected, the synaptic connections architecture might not have had a great impact on the networks evolvability and dynamics.

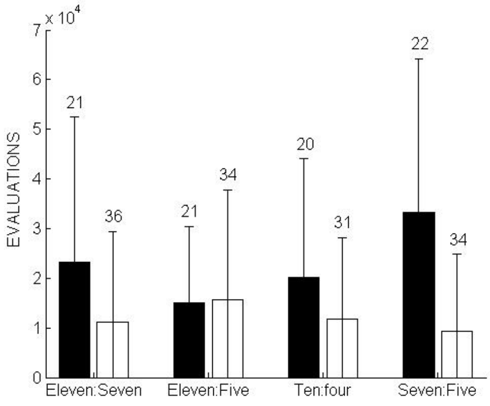

Fig. 2. Mean and standard deviations (error bars) of fitness evaluations required to evolve successful networks for each CPG pattern, Eleven-Seven, Eleven-Five, Ten-Four and Seven-Five. Black bar shows original mean data and white bar shows NSGasNet mean data. The numbers above each error bar represent the total number of successfully evolved networks within 50 runs.

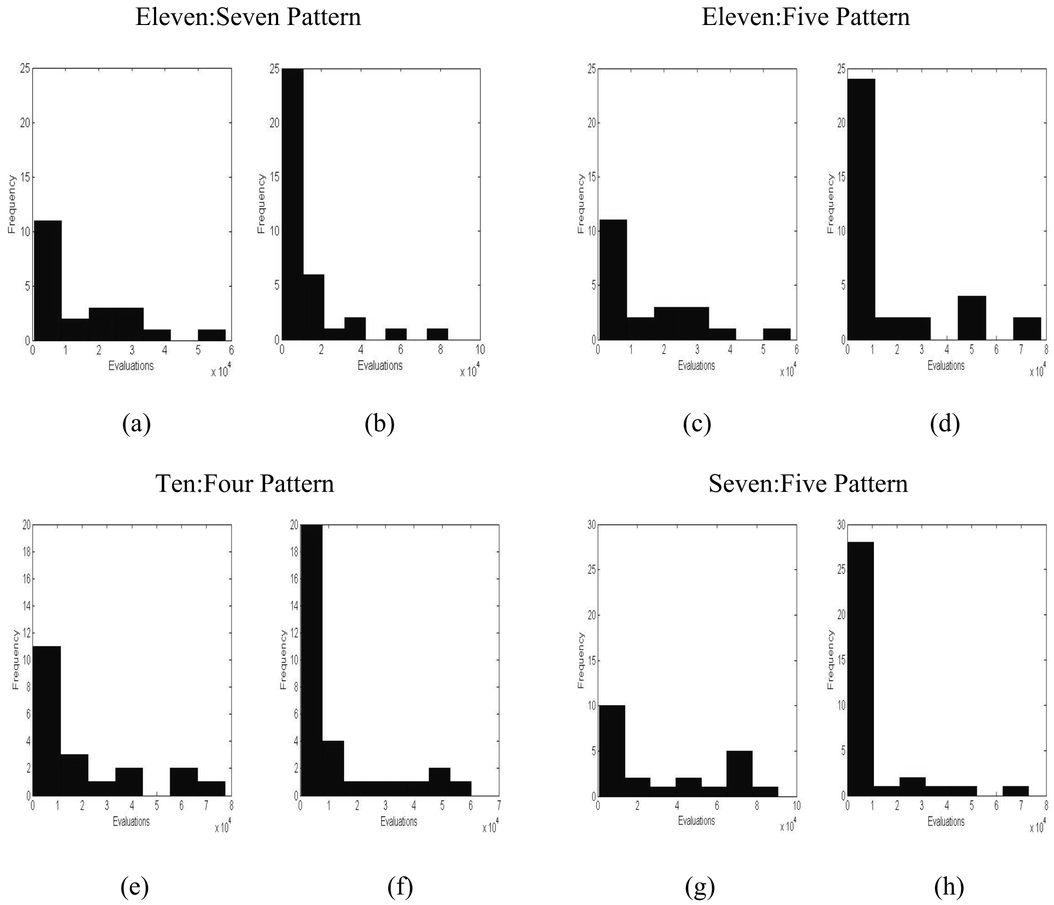

Fig. 3. Frequency histograms comparison between the original (graphics: (a), (c), (e) and (g)) and the NSGasNet (graphics: (b), (d), (f) and (h)) models over the number of fitness evaluations for each pattern. The distribution is conditional on the number of successfully evolved networks within 50 runs for each case (Fig. 2).

Hence, the gaseous connections were mostly responsible for the performance of successfully evolved networks, especially for the NSGasNet network, where the absence of spatial constraints seemed to have contributed to a better performance in terms of speed of evolution. For instance, some evolved NSGasNet networks made explicit use of the modulator bias in a scenario where the presence of a spatial limitation would not have allowed such coupling between the nodes involved.

In both models, after evolution some nodes had no participation on the network dynamics, either for their synaptic weights were set to zero or for there were no gaseous connections whatsoever.

5 Analyzing the Fitness Landscape

In order to investigate the influence of the fitness landscape on evolvability, we have decided to carry out a series of systematic explorations of the fitness landscape around the best networks for each case, thus testing the smoothness of the landscape. This analysis will be performed under the following policy, adopted from Philippides *et al.* [6]. Take the 20 best genotypes of each model and generate the electrical and the gas connection matrices. For each genotype, mutate it and measure the new fitness. Assemble the results under three groups, namely, *Both*, *Elec* and *Gas*. The *Both* group means that the mutant has altered both electrical and gas matrices, the *Elec* group means that the mutant has altered only the electrical matrix, and the *Gas* means that the mutant has altered only the gas matrix. Mutants that do not alter the matrices are discarded. Repeat the process for the entire genotype, mutating each locus. To generate statistical measures the entire policy is performed 100 times for each genotype.

Fig. 4 plots the mean fitness and standard deviations of one-point mutants for the 20 best networks for each model on the Eleven-Seven CPG pattern, original and NSGasNet, respectively. Each bar represents a group in the following order: *Both*, *Elec* and *Gas*.

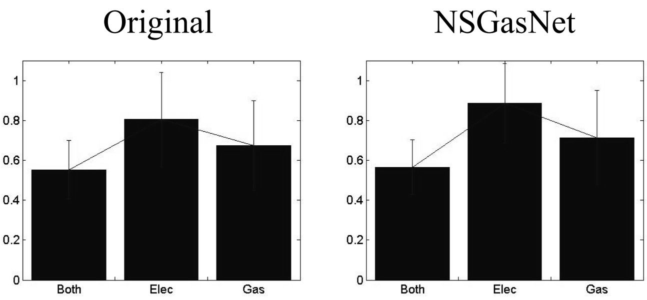

Fig. 4. Mean and standard deviations (error bars) of fitness evaluations of one-point mutants of 20 evolved original and NSGasNet models for the Eleven-Seven CPG Pattern. The bars represent the groups: Both, Elec and Gas, respectively.

Observe that the mutants for the *Elec* group present a higher fitness when compared to the other groups. Also, the mutants that have altered both matrices were the most deleterious ones, thus scoring the lowest fitness. To better access the statistical results, Fig. 5 shows the fitness distribution of the fitness for each group. The first row shows the results for the original model and the second row the NSGasNet model. The histogram depicts the number of fitness evaluations for each group distributed along the fitness value. The first column shows the results for the *Both* group, the second column for the *Elec* group and the third column for the *Gas* Group. The distribution is conditional on the number of fitness evaluations for all one-point mutants within 100 iterations for each successfully evolved network.

From Fig. 5 it is possible to notice that the fitness landscape is not smooth for it seems that many mutations take the individuals to a very low fitness value. This is represented by the high number of mutants with low fitness, especially for the *Gas* group (3rd column). It is now necessary to verify whether the *Gas* group is the best representative of the results. Therefore in Fig. 6, we show the result of the proportion, i.e. the number of one-point mutants divided by the total number of mutants for each group and for each model. It is clear that the number of *Gas* type mutants is much higher than the other groups.

This investigation shows that the landscapes are both equally rugged for the spatial and non-spatial versions. However, spatial constraints affect gas connections in one case and not in the other and this is the set of mutations that have the strongest effect on fitness (Fig. 6), so it is possible that removing spatial constraints from this type of interaction has an overall positive effect on the search. However, this needs further investigation with a more detailed analysis of the evolutionary dynamics.

It is important to highlight that the same analysis was performed for the other CPG patterns and the results are similar.

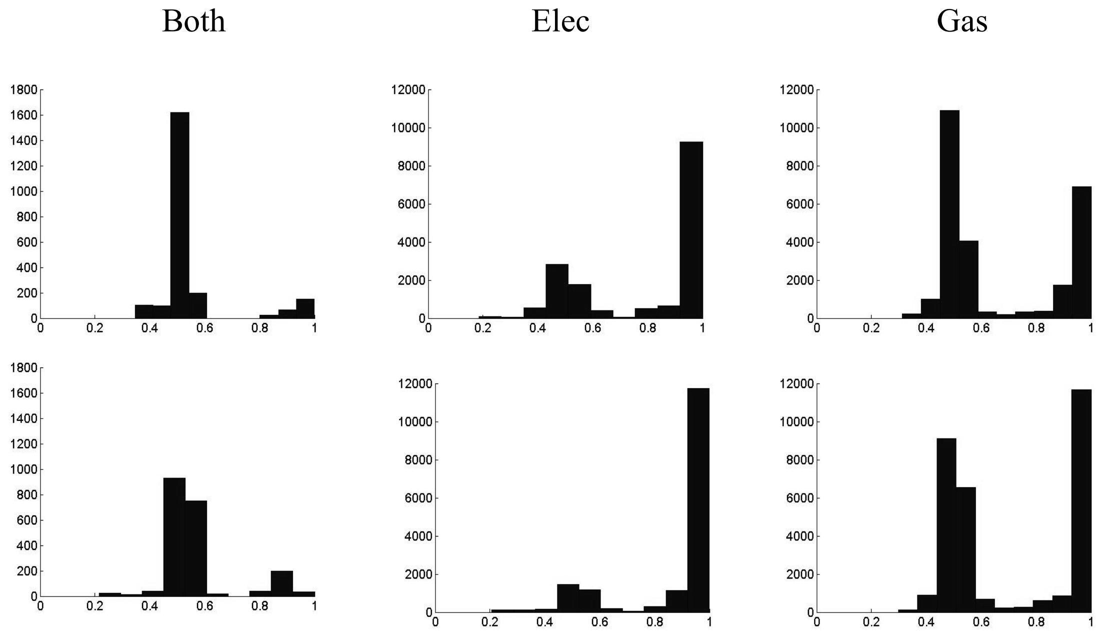

Fig. 5. Frequency histograms comparison between the original (first row) and the NSGasNet (second row) models over the number of fitness evaluations for each group. *Both* group (first column), *Elec* group (second column) and *Gas* Group (third column). The distribution is conditional on the number of fitness evaluations for all one-point mutants within 100 iterations for each successfully evolved network.

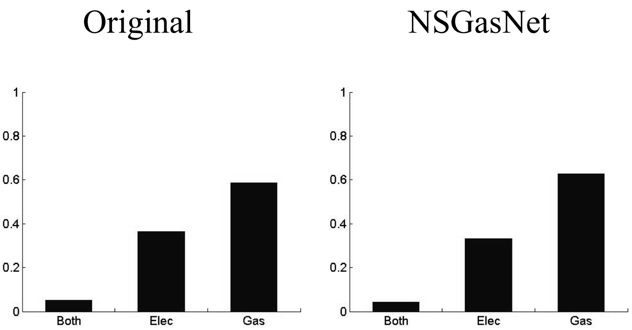

Fig. 6. Proportion of one-point mutants for each group and each model, original and NSGasNet respectively for the Eleven-Seven CPG pattern.

6 Discussion

In this paper, we have attempted to address the role of space in evolving GasNet models by illustrating that this particular neural network model has its evolvability improved when its neurons are not constrained to a Euclidean space. A quantitative analysis of the statistical data was performed and the most relevant findings were reported.

The adoption of a scheme whereby the strength of neuromodulation is regulated by a genetically-determined sensitivity allows us to break from the constraints of Euclidean space. It is still questionable whether some other form of spatiality is introduced by this scheme (where sensitivities would be analogous to some function of distance). The scheme, however, is general and does not comply with basic properties of a distance measure (i.e., symmetry and triagle inequality) so that the model is properly non-spatial. A future comparative study might include sensitivity limits also for the spatial version in order to further assess their role.

With regard to the evolved network architectures, it was impossible to identify a predominant pattern of connections and/or of spatial location of the nodes (original GasNet). This enormous variety could also be explained by the fact that the controllers were evolved instead of being designed by an engineer [21].

It seems that the performance in terms of evolvability of the GasNet models is influenced by the particularities of the task under consideration. For instance, in the first experiment (CPG task) the difference between the numbers of required generations to evolve a successful network is relatively large. Furthermore, networks have low dimensions, i.e. few nodes. This could imply a lack of genetic pressure during evolution; it is possible that the difference in the size of the genotype between both networks could be a matter of concern for larger networks, e.g. hundreds or thousands nodes. Therefore, to fully validate the novel GasNet model it might be necessary to compare both networks in a more complex task, which would require higher dimension networks.

The statistical results show that successful behaviour is achieved in fewer fitness evaluations for the novel unconstrained GasNet than the original model, supporting a tendency investigated by previous GasNet models which had effectively introduced a sort of spatial decoupling between synaptic and gaseous interactions [6].

Finally, we intuit that the role of space might be directed linked to the smoothness of the fitness landscape. An initial analysis of the fitness landscape appeared to demonstrate that the less smooth the landscape the faster is the NSGasNet in terms of evolvability. Certainly, further investigations into landscape smoothness shall be performed in order to fully validate this hypothesis.

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

References

1. Yao, X.: Evolving Artificial Neural Networks. Proc. of the IEEE 87(9), 1423–1447 (1999)
2. Husbands, P.: Evolving Robot Behaviours with Diffusing Gas Networks", in Husbands, P. In: Husbands, P. (ed.) Evolutionary Robotics. LNCS, vol. 1468, pp. 71–86. Springer, Heidelberg (1998)
3. Husbands, P., Smith, T., Jakobi, N., O'Shea, M.: Better living through chemistry: Evolving GasNets for robot control. Connection Science 10(3-4), 185–210 (1998)
4. Philippides, A,: Modelling the Diffusion of Nitric Oxide in Brains. PhD thesis, School of Biological Sciences, University of Sussex, UK (2001)
5. Philippides, A., Ott, S.R., Husbands, P., Lovick, T., O'Shea, M.: Modeling Cooperative Volume Signaling in a Plexus of Nitric Oxide Synthase-Expressing Neurons. The Journal of Neuroscience 25(28), 6520–6532 (2005a)
6. Philippides, A., Husbands, P., Smith, T., O'Shea, M.: Flexible Couplings: Diffusing Neuromodulators and Adaptive Robotics. Artificial Life 11, 139–160 (2005b)
7. Fine, P., Di Paolo, E.A., Philippides, A.: Spatially Constrained Networks and the Evolution of Modular Control Systems. In: Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J.C.T., Marocco, D., Meyer, J.-A., Miglino, O., Parisi, D. (eds.) SAB 2006. LNCS (LNAI), vol. 4095, Springer, Heidelberg (2006)

8. Garthwaite, J.: Glutamate, nitric-oxide and cell-cell signaling in the nervous system. Trends in Neuroscience 14, 60–67 (1991)
9. Getting, P.A.: Emerging Principles Governing the Operation Of Neural Networks. Annual Review of Neuroscience 12, 185–204 (1989)
10. Grillner, S., Wallén, P.: Central Pattern Generators for Locomotion, with Special Reference to Vertebrates. Annual Review of Neuroscience 8, 233–261 (1985)
11. Grillner, S,: Neural Networks for Vertebrate Locomotion. Scientific American, 48-53 (January 1996)
12. Delcomyn, F.: Neural Basis for Rhythmic Behaviour in Animals. Science 210, 492–498 (1980)
13. Katz, P.S., Getting, P.A., Frost, W.N.: Dynamic Neuromodulation of Synaptic Strength Intrinsic to a Central Pattern Generator Circuit. Nature 367, 729–731 (1994)
14. Marder, E., Calabrese, R.L.: Principles of Rhythmic Motor Pattern Generation. Physiological Reviews 76, 687–717 (1996)
15. Roberts, P.D.: Classification of Temporal Patterns in Dynamic Biological Networks. Neural Computation 10, 1831–1846 (1998)
16. Katz, P.S., Fickbohm, D.J., Lynn-bullock, C.P.: Evidence that the Central Pattern Generator for Swimming in Tritonia Arose from a Non-Rhythmic Neuromodulatory Arousal System: Implications for the Evolution of Specialized Behavior. Amer. Zool. 41, 962–975 (2001)
17. Smith, T. M. C,: The Evolvability of Artificial Neural Networks for Robot Control. PhD thesis, CCNR, Department of Informatics, University of Sussex, UK (2002)
18. Chiel, H.J., Beer, R.D., Gallagher, J.C.: Evolution and Analysis of Model CPGs for Walking I. Dynamical Modules. Journal of Computational Neuroscience 7, 99–118 (1999)
19. Beer, R.D., Chiel, H.J., Gallagher, J.C.: Evolution and Analysis of Model CPGs for Walking II. General Principles and Individual Variability. J. of Comp. Neuroscience 7, 119–147 (1999)
20. Psujek, S., Ames, J., Beer, R.D.: Connection and Coordination: The Interplay Between Architecture and Dynamics in Evolved Model Pattern Generators. Neural Comp. 18, 729–747 (2006)
21. Beer, R.D.: A Dynamical System Perspective on Agent-Environment Interaction. Artificial Intelligence 72, 173–215 (1995)