

# Neutrality and ruggedness in robot landscapes

Tom Smith<sup>\*1</sup>, Andy Philippides<sup>1</sup>, Phil Husbands<sup>2</sup> and Michael O’Shea<sup>1</sup>

Centre for Computational Neuroscience and Robotics (CCNR)

<sup>1</sup>School of Biological Sciences, <sup>2</sup>School of Cognitive and Computing Sciences

University of Sussex, Brighton, UK

\* `toms@cogs.susx.ac.uk`

**Abstract-** The twin fitness landscape properties of neutrality and ruggedness are crucial to the dynamics of evolutionary optimisation. In this paper, we investigate the interplay between these two properties in a complex evolutionary robotics fitness landscape, through the introduction of four robot controller architecture models; the *GasNet*, *uniform*, *dispersed* and *plexus* models. We show that in isolation, neither added neutrality or decreased ruggedness (coupling) in the models produces increase in speed of evolution. However, both effects in conjunction produce a significant increase in the speed of evolution.

## 1 Introduction

The twin fitness landscape properties of neutrality and ruggedness are crucial to the dynamics of evolutionary optimisation. However, what little research has been carried out investigating the interaction between the two properties, has typically been based on abstract mathematical fitness landscapes. In this paper, we describe a series of experiments investigating this interaction in a complex evolutionary robotics search space. The experiments described here should be seen as a first step towards a detailed analysis of the interplay between the two landscape properties, providing some preliminary conclusions on the nature of such interaction.

We introduce two variants on the *GasNet* artificial neural network (ANN) architecture [5]. The *uniform* model is designed to introduce extra redundancy into the system, while the *dispersed* model weakens the interaction between two different inter-neuron signalling mechanisms, i.e. reduces ruggedness in the system. In a series of evolutionary runs, we show that both the uniform and dispersed models show no increase in speed of evolution of good solutions when compared with the original *GasNet* model. However, a third variant on the *GasNet* model, the *plexus* model, incorporating both changes from the uniform and dispersed models, does show significantly faster evolution.

We go on to show that both the uniform and plexus models show increased levels of local neutrality, i.e. that the fraction of equal fitness mutants increases, when compared with the *GasNet* and dispersed models. We further show that both the dispersed and plexus models show decreased coupling, or ruggedness, between the gas diffusion and electrical synapse mechanisms, when compared with the *GasNet* and uniform models.

We argue that the speed of evolution results can be explained by the interaction between the twin effects of in-

creased neutrality and decreased coupling. Neutrality is useful for evolutionary exploration of the fitness landscape, without loss of phenotypic fitness. However, neutrality alone is not enough to produce an evolvable system. Similarly, weak coupling between different components of the system, e.g. between the gas diffusion and electrical synapse mechanisms, allows each to be tuned more easily in isolation. However, this reduction in system ruggedness is not enough to produce increased speed of evolution alone. Thus we argue that it is the increase in neutrality and decrease in ruggedness that leads to increased evolvability.

## 2 Neutrality and ruggedness

To evolve successfully, an organism must satisfy the conflicting pressures of *phenotypic stability* and *genetic instability*, i.e. that the organism be robust to phenotypic change (to not fall off the current adaptive peak), and amenable to genotypic change (to allow movement to a new adaptive peak). Conrad [2] identifies genetic redundancy and multiple weak interaction as possible mechanisms by which these two conflicting pressures can be satisfied.

Such loosely coupled redundant systems contain the potential for genotypic change without phenotypic change; both multiple weak interactions and redundancy allow for gradual transformation of function through genetic variation [2]. In such systems, phenotypic fitness is likely to be highly correlated across the genotype landscape, either (or both) through significant levels of neutrality and low levels of ruggedness. Such systems are also robust to phenotypic change; complex systems picked at random are more likely to be stable if the system is characterised by either multiply connected weakly interacting components, or sparsely connected strongly interacting components [4, 8].

By contrast, strongly coupled non-redundant systems are far less amenable to variation; change in one component is more likely to affect the entire system, leading to phenotypic instability. We see this effect clearly in the theoretical *NK* fitness landscapes, where a higher degree of epistatic connection between the components leads to a less correlated fitness landscape [7]. In other words, even small changes in the genotype in a strongly coupled system lead to large changes in the phenotype. However, in tunably neutral versions of the *NK* landscapes, high degrees of neutrality compensate in some measure for the strong coupling, allowing genetic variation without massive phenotypic variation [1].

Clearly, both ruggedness and neutrality are of importance

when designing evolvable systems. However, it is unclear to what extent the two attributes, are complementary and to what extent they overlap. In this paper we introduce three variants on the original GasNet model (section 3), and investigate both neutrality and ruggedness. In particular, we show that both the model incorporating extra redundancy, and the model incorporating weak coupling show no increase in the speed of evolution. However, the model incorporating both effects shows a significant speed of evolution increase.

### 3 GasNets

The *GasNet* incorporates a mechanism based on the neuron-modulating properties of a diffusing signalling gas into a more standard sigmoid-unit ANN [5]. In previous work the networks have been used in a variety of evolutionary robotics tasks, comparing the speeds of evolution for networks with and without the gas signalling mechanism active, showing that GasNets are consistently faster to evolve than more standard ANNs. [5]. A number of related studies have investigated the nature of the GasNet fitness landscapes [see e.g. 13, 14], in order to elucidate the reasons for the faster evolutionary search. However, in this paper, we take a different approach, introducing three variants on the basic GasNet - the *uniform*, *dispersed* and *plexus* models - and using them to investigate the reasons for faster evolutionary search.

#### 3.1 The GasNet model

The GasNet is an arbitrarily recurrent ANN augmented with a gas concentration model, in which the instantaneous activation of a node is a function of both the inputs from connected nodes and the current concentration of gas(es) at the node. The basic network model consists of connected sigmoid transfer function nodes overlaid with a model of gas concentration; the gas does not alter the electrical activity in the network directly but rather acts by changing the gain of transfer function mapping between node input and output.

In order to incorporate the gas concentration model, the network is placed in a 2D plane, with node positions specified genetically. Node connections are also specified genetically, allowing arbitrarily recurrent network architectures. Any node may receive external sensor input, and the four external output nodes are fixed as the first four nodes in the network; in the experiments described here visual input is received by the network, and the output nodes used to generate motor commands for a robot.

The gas diffusion model is implemented through allowing any node to emit one of two gases, which diffuse out from the node in the network plane. Spatially, the gas concentration varies as an inverse exponential of the distance from the emitting node with a spread governed genetically:

$$C(d, t) = \begin{cases} e^{-2d/r} \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (1)$$

where  $C(d, t)$  is the concentration at a distance  $d$  from the emitting node at time  $t$ ,  $r$  is the genetically specified radius of

influence of the node, see figure 1. The maximum concentration at the emitting node is one, and the concentration builds up and decays from this value through  $T(t)$ , which specifies a linear function dependent on time of emission.

The conditions for gas emission to start are also specified genetically; never; when node activity is high; or when gas concentration at the node is high. The effect of gas concentration at a node is to change the node transfer function in a concentration dependent fashion. In the basic model, two gases are used, the first of which increase the gain of the node transfer functions, while the second similarly decreases gain [for further details of the GasNet model see 5].

#### 3.2 Extensions to the basic GasNet: The uniform, dispersed and plexus models

Philippides [9] introduces a number of variants to the original GasNet model, based on research into the diffusion of the gas Nitric Oxide (NO) in real brains [10]. In this paper we investigate three such variants, the uniform, dispersed and plexus models. All three models are directly inspired by the type of signalling seen in the cerebral cortex, where activity in a neuron is translated via a plexus of nNOS-expressing fibres into a volume signal in a different part of the network.

The three variants of the GasNet we investigate are as follows. The *uniform* model simply replaces the exponentially decaying gas concentration function  $C(d, t)$  (equation 1) with a constant function over the volume of effect, see figure 1:

$$C(d, t) = \begin{cases} 0.5 \times T(t) & d < r \\ 0 & \text{else} \end{cases} \quad (2)$$

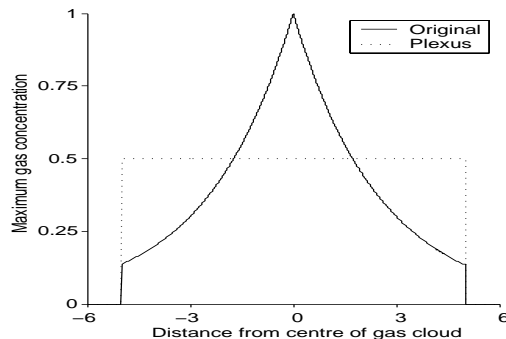


Figure 1: The spatial distributions of gas concentration for the different GasNet models. The solid line denotes the spatial distribution for the GasNet and dispersed models, while the dotted line shows the spatial distribution for the uniform and plexus models.

The *dispersed* model uses the same exponentially decaying gas concentration as the GasNet model, but allows the centre of this gas diffusion cloud to lie anywhere within the space, not just at the emitting node position. Note that this model requires two extra parameters for the gas diffusion centre  $(x, y)$  coordinates. Finally, the full *plexus* model incorporates both the uniform and dispersion models; the gas diffusion produces constant concentration within the area of effect, with the area of effect centred anywhere in the space.

	Original	Uniform	Dispersed	Plexus
Number of runs	40	40	40	40
Mean (S.D.)	3042 (3681)	2783 (3553)	3576 (4097)	1579 (2609)
Median	1201	1016	1128	512
Best	136	90	241	101
Worst	> 10000	> 10000	> 10000	> 10000

Table 1: Number of generations before consistent success is achieved, for the four models described in section 3. Only the plexus model results were significantly faster than the original GasNet results; neither the uniform or dispersed model results showed any such difference. NB runs not achieving consistent success by generation 10000 were terminated.

### 3.3 Visual shape discrimination

The evolutionary task at hand is a visual shape discrimination task; starting from an arbitrary position and orientation in a black-walled arena, the robot must navigate under extremely variable lighting conditions to one shape (a white triangle) while ignoring the second shape (a white square). Both the robot control network, an arbitrarily recurrent ANN, and the robot sensor input morphology, i.e. the position of the input pixels on the visual array, were under evolutionary control. Fitness over a single trial was taken as the fraction of the starting distance moved towards the triangle by the end of the trial period, and the evaluated fitness was returned as the average over 16 trials of the controller from different initial conditions. Success in the task was taken as an evaluated fitness of 1.0 over thirty successive generations of the evolutionary algorithm. In the work reported here, fitness evaluations are carried out in simulation, however evolved controllers have been successfully tested on a real robot [5].

A distributed asynchronous updating evolutionary algorithm was used, with a *PopSize* of 100 arranged on a  $10 \times 10$  grid. Parents were chosen through rank-based roulette-wheel selection on the mating pool consisting of the 8 nearest neighbours to a randomly chosen grid-point. The child solution was a mutated copy of the parent (the mutation operator applied a 4% mutation probability per bit, and the same probability per genome of adding or deleting a network node. No crossover was used.) and placed back in the mating pool using inverse rank-based roulette-wheel selection. One generation was specified as *PopSize* such breeding events.

## 4 Speed of evolution results

Table 1 shows the speed of evolution results for the four models. Forty runs were carried out with each model, with runs being terminated once controllers were evolved that achieved 100% fitness over thirty consecutive generations.

Two main points can be made. First, the GasNet, uniform and dispersed models are fairly similar in their speeds of evolution; in fact the differences are not significant. Second, the plexus model evolves good solutions significantly faster (T-tests were carried out to confirm this) than all three models.

Thus neither the uniform nor the dispersed models show any change in the speed of evolution when compared with

the original GasNet. However the combined effects of the two models does produce faster evolution. In the remainder of the paper we investigate the reasons for this.

## 5 Neutrality

### 5.1 Levels of neutrality

In general, neutrality in the genotype-to-phenotype mapping is simply a consequence of the mapping being a many-to-one function. For each distinct phenotype, there may be a set of distinct genotypes that give rise to that phenotype. By contrast, fitness landscapes showing no such neutrality will have one-to-one genotype-to-phenotype mappings.

In situations where there may be more than one level of phenotype, i.e. where the genotype-to-phenotype mapping has one or more intermediate mapping functions, each intermediate stage may be a source of neutrality. For example, in RNA folding landscapes, genotypes are mapped to a secondary folded structure, which in turn may be allocated a fitness dependent on the distance away from some target secondary structure [see e.g. 6]. Many genotypes will give rise to the same folded structure, and many folded structures will be equidistant from the target structure. Thus two sources, or *levels*, of neutrality will be present.

In the GasNet models described in this paper, we can similarly identify a number of levels of neutrality in the genotype-to-phenotype mapping, and identify the amount of neutrality attributable to each level. We identify the levels as follows:

**Structural neutrality** Many distinct genotypes will code for the same neural network structure.

**Functional neutrality** Many distinct neural network structures will produce the same functional mapping from sensory input to motor output.

**Behavioural neutrality** Many distinct mappings from sensory input to motor output will produce the same final evaluated selective fitness behaviour.

However, the existence of such neutrality is not enough in itself; the neutral genotypes must be distributed in such a manner as to allow neutral movement through the search space. Such neutral networks may be useful for evolutionary search processes; neither small volumes consisting entirely

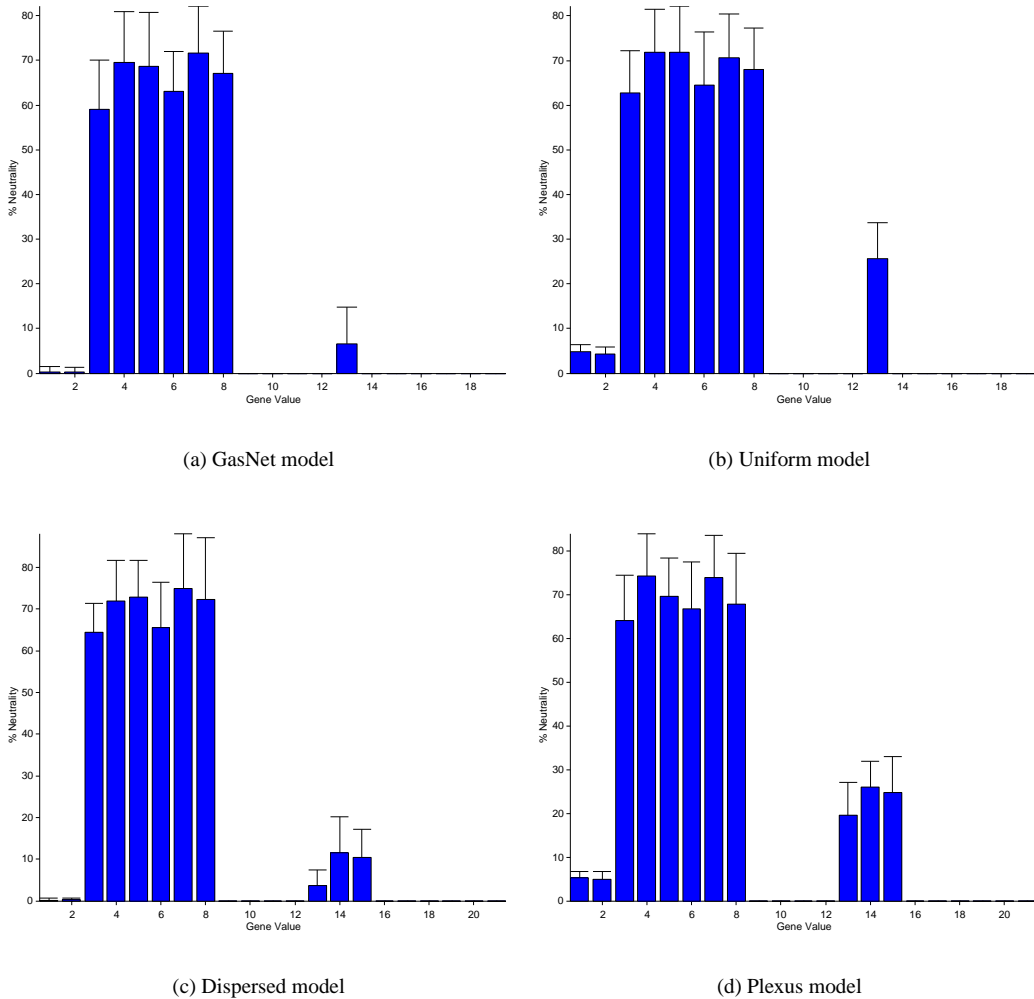


Figure 2: Functional neutrality for the four models. All successfully evolved genotypes were exhaustively sampled using one-point mutation, and the average number of distinct neural network structures producing identical functional mappings from sensory input to motor output plotted against the network node property changed. See text for further details.

	Original	Uniform	Dispersed	Plexus
Number of cases	33	32	30	37
Synaptic connections (S.D.)	1.89 (0.52)	1.70 (0.57)	1.59 (0.47)	1.72 (0.41)
Diffusion connections (S.D.)	2.27 (0.93)	2.05 (0.78)	2.78 (0.75)	2.78 (0.84)
Overlapping connection coupling (S.D.)	40.5% (13.2%)	36.3% (11.1%)	10.4% (7.3%)	10.8% (8.1%)

Table 2: Coupling in the original GasNet and plexus models. For each of the successfully evolved controllers, the number of electrical synaptic connections and number of gas diffusion connections are shown (averaged per neuron). The percentage of connections which overlap, i.e. that connect the same neurons, are also shown. See text for further details.

of neutral solutions, or neutral solutions scattered randomly throughout the space will be useful. In this paper, we focus on such connected neutral sets, calculating the amount of neutral mutations around solutions.

We calculate the different sources of neutrality for a single solution as follows. First, the structural neutrality is derived analytically through consideration of the various node

parameters, e.g. over what range of each parameter is the corresponding network property unchanged? Second, the functional neutrality is empirically derived through exhaustively sampling all one-point mutations of the genotype. The connection matrices, e.g. which nodes are connected and with what strength, are calculated for each mutated solution and compared with the original; these connection matrices are

different only for those networks with different input-output functional mappings. The functional neutrality is then calculated as the fraction of solutions for which the connection matrices are identical to the original, i.e. the functional input-output mapping is identical, averaged over each node property, less the structural neutrality calculated for that node property in the previous step. Behavioural neutrality is calculated as the total fraction of solutions showing identical selective fitness, less both the structural and functional neutrality. It should be noted that the solution neutrality calculated through these methods will not in general be identical to the solution neutrality seen during evolutionary adaptation, as typically the mutation operators used will not be simple one-point mutation. However, we can use the method to compare neutrality between the different models.

In this paper, we focus solely on the second source of neutrality, i.e. functional neutrality. Here we are interested only in the fraction of distinct neural networks that give rise to identical sensor-input to motor-output functional mappings. Although clearly important, behavioural neutrality is crucially dependent on the details of the task at hand so results may not be applicable to a wider range of tasks. Structural neutrality is also not considered, as it is typically caused through weaknesses in the representation of the neural network structure, e.g. encoding binary parameters as  $[0, 1]$  range variables (or in our case,  $[0, 99]$  range integers).

## 5.2 Functional neutrality of the models

In figure 2, we show the functional neutrality against node loci number, derived over all successfully evolved controllers under the four models described in section 3. Each node in the network is coded for by 19 or 21 parameters (the extra two being for the gas diffusion centres in the dispersed and plexus models).

Four main points can be made. First, the parameters 3-8 are highly neutral for all four models; these relate to the way in which electrical synaptic connections are formed. Second, the only difference seen between the GasNet and dispersed model lies in the neutrality in the dispersed model parameters 14 and 15, which relate to the centres of gas diffusion and are not present in the GasNet model. Third, the uniform cloud model shows increased neutrality in parameters 1, 2 (the  $(x, y)$  node coordinates) and 13 (the radius of gas diffusion effect), when compared with the GasNet model. Fourth and finally, the plexus model shows no significant change in neutrality when compared with the uniform model, again apart from the gas cloud centre parameters 14 and 15.

We conclude that the uniform model introduces significant sources of neutrality into existing parameters in the system, allowing change in the node  $(x, y)$  positions and the gas cloud radius. Similarly, the dispersed model introduces some neutrality through the new parameters, although the neutrality in these parameters is much greater in the full plexus model. The extra neutrality in the plexus model is thus seen to be explained through the incorporation of both the dispersed and

uniform cloud model.

## 6 Coupling

As described in section 3.2, the dispersed and plexus models allow network nodes to emit gas from anywhere in the grid (another two parameters are required to specify this  $(x, y)$  gas emission coordinate). This partly separates the gas diffusion and the electrical synaptic activity mechanisms; synaptic connections are formed from the current node position, while gas diffusion connections are formed from the gas emission position. Thus gas connections in the grid can be changed through modifying the gas emission position, while synaptic connections can be altered through moving the node itself. Note that the two mechanisms are not entirely separate: both act on the actual position of the destination nodes. It should also be noted that this separation could be achieved in a number of different ways, however in this paper we focus solely on the plexus model.

For a given genotype in the GasNet model, both the synaptic and diffusion outgoing connections are dependent on the node position; changing these genetic parameters will potentially radically alter the phenotype. For a given genotype in the dispersed and plexus models, the two outgoing connection processes are dependent on two different sets of parameters; changing either of these will affect only the relevant connection mechanism. Thus we argue that the degree of coupling between the parameters of the system will be higher for the GasNet model than for the dispersed and plexus models.

There is no simple way of calculating the degree of coupling in the four systems; in principle one can measure the degree of ruggedness through correlation lengths or similar methods. However, Smith et al. [13] shows that these types of measures do not discriminate between highly heterogeneous problem spaces such as the GasNets. In this section, we introduce a simple description to measure the degree of coupling between the gas diffusion and electrical synapse mechanisms. We calculate the two connectivity matrices for a given solution (section 5.1), and calculate the coupling as the number of overlapping connections, i.e the number of elements which are non-zero in both connectivity matrices.

Table 2 shows this coupling between the electrical and gas diffusion processes for the four models. The number of electrical synaptic connections and number of gas diffusion connections (averaged per neuron), and the percentage of overlapping connections, are shown for each of the successfully evolved controllers, over all models. Three points can be made. First, there are no significant differences between the numbers of electrical synaptic connections across the four models. Second, there is some evidence that the number of gas diffusion connections is higher for the dispersed and plexus models than for the GasNet and uniform models. This is not discussed further in this paper, but is some indication that the plexus and dispersed models are based on more multiple interactions between nodes; further work will

investigate this effect. Third, the percentages of overlapping connections in the GasNet and uniform controllers are significantly higher than those in the dispersed and plexus models; this coupling between the electrical and diffusion processes is far stronger in the GasNet and uniform models than in the dispersed and plexus models.

However, this decoupling of processes in the dispersed and plexus models does not lead inevitably to faster evolutionary search, see table 1. The dispersed model shows no such increase over the GasNet and uniform models. However the plexus model shows significantly faster evolution, when compared with all three other models.

## 7 Discussion

A number of studies have shown that the performance of evolutionary optimisation processes is affected by the presence of fitness landscape neutrality [see e.g. 1, 15, 3]. Similarly, the effect of ruggedness on the performance of optimisation processes is well known [see e.g. 16, 7]. Studies in theoretical landscapes have explored interactions between the two properties [1, 11, 12], however the interaction between the processes in real fitness landscapes is not well understood.

In this paper we have introduced preliminary analysis of the interaction between neutrality and ruggedness in a complex fitness landscape. We have developed a number of evolutionary robotic control architectures designed to incorporate both increased redundancy and decreased coupling. In a series of evolutionary experiments we saw that only the plexus model, incorporating both increased redundancy and reduced coupling between two inter-neuron signalling measures, shows an increase in evolutionary speed. The uniform cloud model, incorporating extra neutrality but no decoupling of the gas and electrical processes, showed no such faster evolutionary speed. Similarly the dispersed model, incorporating decoupled gas and electrical processes and also some extra neutrality in the new genetic parameters, showed no faster evolutionary speed.

The experiments presented here should be seen as a preliminary investigation into the interaction between neutrality and ruggedness in complex landscapes, and are unlikely to tell the full story. In further work we will investigate this in more detail. In particular, how does the extra neutrality affect the evolutionary search process, i.e. are there differences between the four models in the movement of evolutionary populations over time?

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

## Bibliography

- [1] L. Barnett. Ruggedness and neutrality: The NKp family of fitness landscapes. In C. Adami, R.K. Belew, H. Kitano, and C.E. Taylor, editors, *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*, pages 18–27. MIT Press, Cambridge, Massachusetts, 1998.
- [2] M. Conrad. The geometry of evolution. *BioSystems*, 24:61–81, 1990.
- [3] M. Ebner, P. Langguth, J. Albert, M. Shackleton, and R. Shipman. On neutral networks and evolvability. In *Proceedings of the 2001 Congress on Evolutionary Computation (CEC'2001)*, pages 1–8. IEEE Press, Piscataway, New Jersey, 2001.
- [4] M.R. Gardner and W.R. Ashby. Connectance of large dynamic (cybernetic) systems: Critical values for stability. *Nature*, 228:784, 1970.
- [5] P. Husbands, T.M.C. Smith, N. Jakobi, and M. O'Shea. Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 10(3-4):185–210, 1998.
- [6] M.A. Huynen, P.F. Stadler, and W. Fontana. Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Sciences, USA*, 93:394–401, 1996.
- [7] S.A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, Oxford, UK, 1993.
- [8] R.M. May. Will a large complex system be stable? *Nature*, 238:413–414, 1972.
- [9] A. Philippides. *Modelling the Diffusion of Nitric Oxide in Brains*. PhD thesis, School of Biological Sciences, University of Sussex, UK, 2001.
- [10] A.O. Philippides, P. Husbands, and M. O'Shea. Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *Journal of Neuroscience*, 20(3):1199–1207, 2000.
- [11] C.M. Reidys and P.F. Stadler. Neutrality in fitness landscapes. *Applied Mathematics and Computation*, 117(2-3):321–350, 2001.
- [12] T.M.C. Smith, P. Husbands, P. Layzell, and M. O'Shea. Fitness landscapes and evolvability. *Evolutionary Computation*, 10(1):1–34, 2002.
- [13] T.M.C. Smith, P. Husbands, and M. O'Shea. Not measuring evolvability: Initial exploration of an evolutionary robotics search space. In *Proceedings of the 2001 Congress on Evolutionary Computation (CEC'2001)*, pages 9–16. IEEE Press, Piscataway, New Jersey, 2001.
- [14] T.M.C. Smith, P. Husbands, and M. O'Shea. Local evolvability of statistically neutral GasNet robot controllers. *Biosystems*, 2002. In press.
- [15] E. van Nimwegen, J.P. Crutchfield, and M. Huynen. Neutral evolution of mutational robustness. *Proceedings of the National Academy of Sciences, USA*, 96:9716–9720, 1999.
- [16] E. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325–336, 1990.