


2 Search spaces and neutrality

This section introduces two of the main concepts used in the paper. The *fitness landscape* (section 2.1), first introduced by Wright (1932), describes the search space as a multi-dimensional landscape defined by the genotype-to-fitness mapping through which evolution moves. The classical idea of searching this landscape for good genotypes focuses on the difficulty of climbing up to the globally optimal fitness solution, and avoiding locally optimal solutions. Here we

adaptiv

e

Figure 1: A two-dimensional model fitness






Figure 2: Three two-dimensional model fitness landscapes showing the possible advantage of neutrality in a simple landscape with one globally-optimal and one locally-optimal peak. (a) shows the two peaks as unconnected; populations evolving to the locally optimal peak will have difficulty moving to the global optimum. (b) shows the two peaks connected by a single neutral network; a population on the sub-optimal peak will eventually find the neutral pathway. (c) shows the two peaks connected by a broad plateau; the population will move easily from the sub-optimal peak to the global optimum.

of the population during these periods of neutral drift from the behaviour of a population stuck in a local optimum. The first key feature is that of $\frac{dD}{dt} > 0$, i.e. is the population moving significantly in genotype space. The second key feature is that of $\frac{dN}{dt} = 0$, i.e. the number of previously unencountered phenotypes seen over time is constant. An immediate corollary of this property is that any phenotype should be accessible from a large enough neutral network; neutral drift can eventually find a higher fitness genotype and jump up to a higher neutral network.

These key features have been analysed in many theoretical landscapes. Barnett (1998) introduces the NKp landscape, a tunably neutral variant on Kauffman's NK systems (Kauffman, 1993),

through exhaustive mutation of genotypes at the start of the plateau that no transitions exist to a higher fitness level; the neutral evolution phase is necessary.

robotic controller architectures. The search spaces defined by these two genetic systems show a high degree of neutrality, and no differences are seen using standard measures of ruggedness and modality (Smith et al., 2001a). Further details of the genotype-to-fitness mapping are given in section 4.

Other biological research in evolvability is also of relevance to evolutionary computation, e.g. the work on adaptation to change in environment through such mechanisms as alleles providing increased mutation rates (Taddei et al., 1997; Sniegowski et al., 1997). However, in this paper we focus on evolvability in terms of the properties of the solutions' local search space. The next section outlines the offspring transmission function, and defines a simple set of evolvability metrics.

3.1 The transmission function


In this paper, we follow the definition of evolvability as the ability of individuals and populations to produce fit variants, specifically the ability to both produce fitter variants, and to not produce less fit variants. This definition is intimately tied in with research on the fitness landscape.

T

3.2 Evolvability metrics

The evolvability of a solution genotype h and fitness k

Thus the evolvability metrics provide more information than measures of correlational structure in the landscapes, giving detail on how easy it will be to find solutions of varying fitness. In particular, they correctly predict that the time required to find good solutions for varying K will depend on how good a solution we need: for low fitnesses a maximally rugged landscape is best, but as our required fitness increases, a smaller degree of epistasis results in faster search. Note, this is of course affected by random solutions already having expected fitnesses of 0.5, but holds if we are starting from low fitness solutions.



applied evolutionary robotics genotype-fitness mapping used in the work presented here.

linearly increased or decreased depending on whether the node is emitting or not. Note $T(t)$ saturates at a maximum of 1 and a minimum of 0. The total concentration at any point in the network is found by summing the concentrations from all emitting nodes.

o on

The transfer parameter value for the i th node at time step n , k_i^n (see equation 8), is changed (or modulated) by the presence of gases at the site of the node. Gas 1 increases the value of k_i^n in a concentration dependent way, while gas 2 decreases its value. This modulation is described by equations 12 to 14 and happens on every time step as the network runs. This provides a form of plasticity very different from that found in most traditional artificial neural networks.

$$k_i^n = [\text{index}_i^n], \quad = \{-4.0, -2.0, -1.0, -0.5, -0.25, 0.0, 0.25, 0.5, 1.0, 2.0, 4.0\} \quad (12)$$

where,

$$\text{index}_i^n = f \left(\text{index}_i^0 + \frac{C_1^n}{C_0 \times K} (N - \text{index}_i^0) - \frac{C_2^n}{C_0 \times K} \text{index}_i^0 \right) \quad (13)$$

$$f(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x < N \\ N & \text{else} \end{cases} \quad (14)$$

where $[i]$ refers to the i th element of set \mathcal{K} , index_i^n is node i 's index into the set of possible discrete values k_i^n can assume, N is the number of elements in \mathcal{K} , index_i^0 is the genetically set default value for index_i , C_1^n is the concentration of gas 1 at node i on time step n , C_2^n is the concentration of gas 2 at node i on time step n , and C_0 and K

triangle by the end of the trial period, and the evaluated fitness was returned as the weighted sum of 16 trials of the controller from different initial conditions:

$$F = \frac{1}{136} \sum_{i=1}^{i=16} i \frac{D_i^F}{D_i^S} \quad (15)$$

where D_i^F is the distance to the triangle at the end of the i th trial, and D_i^S the distance to the triangle at the start of the trial, and the i trials are sorted in descending order of $\frac{D_i^F}{D_i^S}$. Thus good trials, in which the controller moves some way towards the triangle, receive a smaller weighting than bad trials, encouraging robust behaviour on all 16 trials.




Figure 5: Screen shot of the simulated arena and robot. The bottom-right view shows the robot position in the arena with the triangle and square. Fitness is evaluated on how close the robot approaches the

- Initialise population of 100 solutions on 10x10 grid.
- Evaluate each solution fitness.
- Repeat until success criterion met, or MaxGenerations reached:
 - Repeat 100 times for 1 generation:
 - Select solution at random.
 - Create mating pool of solution plus 8 nearest grid neighbours.
 - Pick parent P through rank-based roulette wheel selection on mating pool.

in the mutated offspring networks when compared to the parents: table 1 gives some evidence that even higher mutation rates are useful.

4.6 Previous results – GasNets are faster

Over a large sample of evolutionary runs with GasNet and NoGas conditions, GasNet networks allowed to use the gaseous signalling mechanism reached success significantly faster than the NoGas networks (remember from section 4.3 that success is defined as 100% fitness over 30 consecutive generations). This speed difference was seen in several different evolutionary robotics scenarios (e.g. Husbands, 1998; Smith and Philippides, 2000) and over several different mutation rates, e.g. see table 1.

Condition	$\mu = 0.01$	$\mu = 0.02$	$\mu = 0.04$	$\mu = 0.08$	$\mu = 0.16$
GasNet	7,354	3,436	675	449	671

5 Evolvability analysis of a search space

In previous work (Smith et al., 2001a), we have shown that several standard analyses (correlation structure (Weinberger, 1990), local modality (Weinberger, 1991), and information structure (Vassilev et al., 2000)), derived to predict the difficulty of finding good solutions in a given space, fail to predict the evolutionary speed differences described in section 4.6. There are three primary reasons for these failures. First, the fitness distribution of the space is extremely skewed, with the vast majority of genotypes having extremely low fitness: statistics calculated through random sample techniques fail to explore the interesting parts of the space. Second, the large amount of noise inherent in the fitness evaluation tends to obscure measures of the space ruggedness and local modality. Third, most of the evolutionary runs show extremely long periods during which fitness does not apparently increase. Barnett (1998) has shown the unreliability of correlation structure measures when applied to spaces with neutrality; if these periods of no apparent fitness increase are indeed neutral epochs then we might expect many standard measures to fail when predicting search difficulty.

In spaces with a high degree of neutrality, it may well be that analysis of the evolvability of solutions will predict the difficulty of finding successful solutions. This is backed up by the evidence presented on the evolvability of tunably neutral NK landscapes, shown in section 3.3, which tallies with the time required to evolve good solutions (Newman and Engelhardt, 1998). Thus we frame the hypothesis that evolvability analysis of t







Figure 9: The fitness distribution histogram of a single genotype evaluated 10,000 times. (A) All evaluations plotted, and (B) Only the best fitness every 50 evaluations plotted. Fitness ranges from [0.23, 0.37] and was used to calculate the range of the neutral band in figure 10. Possible fitness [0, 1].

neutral network within this band; the first lower fitness band lies roughly between generations 50 and 100, the second between generations 100 and 477. No further analysis is done on possible multiple neutral networks in the band; this would be an interesting area of further study.





experiment empirically tests this prediction, repeating the evolution from different points along the neutral epoch. Five populations, from generations {100, 200, 300, 400, 477}






Figure 17: Median generations before given fitness level reached, for GasNet and NoGas conditions.

extremely deleterious mutations from the GasNet solutions, rather than an increased number of intermediate mutations. By contrast, the probability of mutations which increase fitness, and the overall expected fitness of mutations were extremely similar for the two search spaces (the differences in the expected fitness of the worst mutations has only a relatively small impact on the overall expected fitness).

Although the differences are small, the larger number of neutral mutations in the GasNet space, especially at higher fitnesses, will result in a larger number and variety of good solutions in the population at any one time. The probability of a deleterious mutation being subsequently reversed is small², thus at any one time, evolution in the GasNet space has a greater probability of finding the good mutations: both the increased number of good solutions, and the increased variety will aid evolution.

Discussion

The promise of artificial evolution lies in producing good solutions to difficult problems. However, if we are to apply such techniques to evolving solutions for problems with time consuming fitness evaluation, we must develop suitably evolvable solution architectures. In recent work we have developed the "GasNet", and shown in a series of evolutionary robotics experiments that a significant decrease in the numbers of evaluations is required to evolve successful GasNet controllers (Husbands, 1998; Husbands et al., 1998; Smith and Philippides, 2000). In this paper we have applied the concept of evolvability in order to analyse the search spaces underlying two different genetic systems, the GasNet and NoGas robot control architectures.

We have shown how a set of evolvability metrics defined in terms of the transmission function, or the distribution of solution offspring fitnesses, can predict the difficulty of finding good solutions in a theoretical class of tunably rugged and neutral landscapes (Kauaman, 1993; Newman and Engelhardt, 1998). In particular, we have shown that the metrics can predict the difficulty of finding solutions of a given fitness, not possible with single measures of landscape "difficulty" such as the correlation length.

² The probability of a deleterious mutation being subsequently reversed is small, thus at any one time, evolution in the GasNet space has a greater probability of finding the good mutations: both the increased number of good solutions, and the increased variety will aid evolution.

funded by a British Telecom sponsored Biotechnology and Biology Science Research Council Case award.

References

- Altenberg, L. (1994). The evolution of evolvability in genetic programming. In Kinnear Jr, K., editor, *Genetic Programming: Proceedings of the First Annual Conference*, chapter 3, pages 47–74. MIT Press.
- Barnett, L. (1998). Ruggedness and neutrality: The NKp family of fitness landscapes. In Adami, C., Belew, R., Kitano, H., and Taylor, C., editors, *Genetic Programming: Proceedings of the Second Annual Conference*, pages 18–27. MIT Press / Bradford Books.
- Burch, C. and Chao, L. (2000). Evolvability of an RNA virus is determined by its mutational neighbourhood. *Journal of Molecular Evolution*, 406:625–628.
- Cavalli-Sforza, L. and Feldman, M. (1976). Evolution of continuous variation: Direct approach through joint distribution of genotypes and phenotypes. *Evolution*, 73:1689–1692.
- Cli, D. T., Harvey, I., and Husbands, P. (1993). Explorations in evolutionary robotics. *Artificial Intelligence Magazine*, 14(2):27–33.

Hordijk, W. (1996). A measure of landscapes. *Artificial Intelligence*, 4(4):335–360.

Husbands, P. (1998). Evolving robot behaviours with distributed gas networks. In Husbands and Meyer (1998), pages 71–86.

Husbands, P. and Meyer, J.-A., editors (1998). *Artificial Intelligence, Robotics and the World*. Springer-Verlag, Berlin.

Husbands, P., Smith, T., Jakobi, N., and O’Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Artificial Intelligence*, 10(3-4):185–210.

- Shackleton, M., Shipman, R., and Ebner, M. (2000). An investigation of redundant genotype-phenotype mappings and their role in evolutionary search. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 493–500. IEEE, San Diego, USA.
- Smith, T., Husbands, P., and O’Shea, M. (2001a). *Measuring evolvability: Initial exploration of an evolutionary robotics search space*. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 9–16. IEEE, Seoul, Korea.
- Smith, T., Husbands, P., and O’Shea, M. (2001b). An evolvability measure of landscapes. In preparation.
- Smith, T. and Philippides, A. (2000). Nitric oxide signalling in real and artificial neural networks. *Neuroinformatics*, 18(4):140–149.
- Sniegowski, P., Gerrish, P., and Lenski, R. (1997). Evolution of high mutation rates in experimental populations of *E. coli*. *Science*, 275:703–705.
- Stadler, P. (1996). Landscapes and their correlation functions. *Journal of Mathematical Biology*, 20:1–45.
- Taddei, F., Radman, M., Maynard-Smith, J., Toupance, B., Gouyon, P., and Godelle, B. (1997). Role of mutator alleles in adaptive evolution. *Science*, 275:700–702.
- Thompson, A. (2001). Neutrality in evolutionary hardware experiments. In preparation.
- Turney, P. (1999). Increasing evolvability considered as a large-scale trend in evolution. In Wu, A., editor, *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 43–46.
- van Nimwegen, E., Crutchfield, J., and Huynen, M. (1999). Neutral evolution of mutational robustness. *Science*, 286:9716–9720.
- Vassilev, V., Fogarty, T., and Miller, J. (2000). Information characteristics and the structure of landscapes. *Evolutionary Computation*, 8(1):31–60.
- Vassilev, V. and Miller, J. (2000). The advantages of landscape neutrality in digital circuit evolution. In Miller, J., Thompson, A., Thomson, P., and T., F., editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 1801 of *LNCS*, pages 252–263. Springer-Verlag.