

Neutral Networks and Evolvability with Complex Genotype-Phenotype Mapping

Tom Smith^{1,2}, Phil Husbands^{1,3}, and Michael O’Shea^{1,2}

¹ Centre for Computational Neuroscience and Robotics (CCNR)

² School of Biological Sciences

³ School of Cognitive and Computing Sciences

University of Sussex, Brighton, UK

toms@cogs.susx.ac.uk

Abstract. In this paper, we investigate a neutral epoch during an optimisation run with complex genotype-to-fitness mapping. The behaviour of the search process during neutral epochs is of importance for evolutionary robotics and other artificial-life approaches that evolve problem solutions; recent work has argued that *evolvability* may change during these epochs. We investigate the distribution of offspring fitnesses from the best individuals of each generation in a population-based genetic algorithm, and see no trends towards higher probabilities of producing higher fitness offspring, and no trends towards higher probabilities of not producing lower fitness offspring. A second experiment in which populations from across the neutral epoch are used as initial populations for the genetic algorithm, shows no difference between the populations in the number of generations required to produce high fitness. We conclude that there is no evidence for change in evolvability during the neutral epoch in this optimisation run; the population is not doing anything “useful” during this period.

1 Introduction

Genetic algorithms are classically regarded as performing hill-climbing. Populations of solutions are progressively improved until some acceptable level of fitness is reached, with problems occurring if the system gets stuck in local optima before the required fitness is reached. However, the idea of non-adaptive *neutral* mutation [11,9,2,21] extends this picture to incorporate the idea of connected sets of equal fitness solutions, or *neutral networks*. In this scenario, the population of solutions randomly drifts along a neutral network, occasionally undergoing transitions to higher fitness networks¹, see figure 1. Local optima may not even exist; long periods where the system does not improve in fitness may indicate neutral change rather than simply being stuck in an optimum. Investigation of the population behaviour during these neutral epochs is thus important for artificial life search techniques that evolve problem solutions.

In such a space, the emphasis on determining the difficulty of finding good solutions through measures of ruggedness and epistasis may be misplaced. Instead,

¹ The population may also drop in fitness; work on *error thresholds* looks at the conditions under which this transition to lower fitnesses may occur [16].




Fig. 1. The fitness of the best individual in a population, over generations. Note the relatively long periods over which fitness does not increase, seen in many optimisation problems. Instead of being stuck in a local optimum, populations may be exploring neutral networks in the space.

we will need to look at other properties of the space, and how these properties change during the neutral period. One such property is the *evolvability* of solutions (or populations of solutions), argued in this paper to be equivalent to the capacity of individuals to produce fit offspring. Recent work has shown that this capacity may change during non-adaptive evolutionary phases [25,24].

In this paper, we investigate the neutral networks and evolvability in a system with an extremely complex genotype-to-fitness mapping. The genotypes code for arbitrarily recurrent neural networks used as robot controllers in a visual shape discrimination. Fitness is evaluated on how close the robot moves to a target shape over an extended period of operation. It is by no means clear that previous work on both neutrality and evolvability [24,5,15] will apply to such complex genotype-fitness mapping spaces. The paper extends the analysis presented in previous work [17,18], here using the idea of the *transmission function* [1] to investigate how evolvability changes during the neutral phase.

The paper is laid out as follows: Section 2 introduces the ideas of evolvability and the transmission function. Section 3 describes the style of neural network used in the work, and the robotics problem on which the evolved network controllers are evaluated. Sections 4 and 5 outline the two experiments carried out, and the results found, and we conclude with discussion.

2 Evolvability and Neutrality

Evolvability is loosely defined as the capacity to evolve, alternatively the ability of an individual or population to generate fit variants [1,13,22]. Attempts to rigorously define the concept tend to reflect the background of the researchers involved, although recently there has been more work linking the biological and computer science approaches [3,20]. Biologists often talk of organisms and structures already pre-adapted to some environment, and their ability to respond to environmental change [12]. Computer scientists tend to talk of evolvability in terms of the ease of finding good solutions in a given search space, closely tied in with work on the properties of search space ruggedness and modality [23,6,10,14,17].

In this paper, we define evolvability as the ability of individuals 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 the population transmission function [1,4]: $T(i \rightarrow j)$, defined as the probability distribution of offspring j obtained through all possible applications of the genetic operators on the parent i (in this work we do not apply recombination, so only a single parent is required). Such a definition encompasses all variation in both the operators and the representation; instead of referring to good and bad genetic operators or good and bad representations, we can talk about the effectiveness of the transmission function. In the remainder of the work, we use the transmission function as short-hand for the distribution of offspring *fitnesses*. Thus the evolvability of an individual or population, i.e. their ability to generate fit variants, is simply a property of the individual or population transmission function.

Researchers have argued that there may be large-scale trends for evolvability to change during evolution [20], and that the capacity can even increase during neutral epochs through the population moving to “flatter” areas of the search space, with fewer deleterious mutations [25,24]. This effect can occur as the centre of mass of the population moves towards solutions producing a higher number of viable offspring. Evolvability may also change as the population diffuses along neutral networks, thus potentially escaping local optima; adding neutrality may increase evolutionary speed and hence evolvability [15,5]. This paper investigates these claims in an evolutionary robotics setting with complex genotype-to-fitness mapping. The next section introduces this mapping.

3 An Evolutionary Robotics Search Space

The *GasNet*, introduced by Husbands *et al.* [7,8], incorporates a mechanism based on the neuron-modulating properties of a diffusing signalling gas into a more standard sigmoid-unit neural network. 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 [7,8,19]. In a variety of robotics tasks, GasNet controllers evolve significantly faster than networks without the gas signalling mechanism. Initial work aimed at identifying the reasons for this faster search has focused on both the underlying search spaces ruggedness and modality [17], and the non-adaptive phases of evolution [18].

3.1 The Task

The evolutionary task is a visual shape discrimination problem; 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 a second shape (a white square). Both the robot control network (an arbitrarily recurrent neural network incorporating artificial diffusing neuro-modulators) and the robot sensor input morphology (the number and position

of input pixels used in 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 weighted sum of 16 trials of the controller from different initial conditions. For further details of the task, fitness function and genetic algorithm used, see [8,19].

Success in the task was taken as an evaluated fitness of 1.0 over thirty successive generations of the genetic algorithm. Previous work has shown that controllers incorporating the diffusion mechanism can evolve successful solutions significantly faster controllers without the mechanism enabled [8,19]. The research presented here is part of an extensive exploration into the reasons for this faster search [17,18].

3.2 The Solution Representation and Mutation Operator


The neural network robot controllers were encoded as variable length strings of integers, with each integer allowed to lie in the range [0, 99]. Each node in the network was coded for by nineteen parameters, controlling such properties as node connections, sensor input, and node bias. In all experiments, the GA population were initially seeded with networks containing ten neurons. For further details see [8,19].

Three mutation operators were applied to solutions during evolution. First, each integer in the string had a 3% probability of mutation in a Gaussian distribution around its current value (20% of these mutations completely randomised the integer). Second, there was a 3% chance per *genotype* of adding one neuron to the network, i.e. increasing the genotype length by 19. Third, there was a 3% chance per genotype of deleting one randomly chosen neuron from the network, i.e. decreasing the genotype length by 19.

4 The Experiments

We concentrate on a single evolutionary run, chosen at random from a set of runs used in previous work [8]; figure 2(a) shows the population best and mean fitnesses over generations. In previous analysis of this evolutionary run [18] we have shown that the evolutionary phase lying between generations 100 and 477 is indeed a neutral epoch [21]. First, the variation in multiply evaluated fitnesses for a single genotype can explain the variation in the population best fitness over this period, see figure 2(b). Second, the population is moving significantly through the search space during this phase, and thus not simply stuck in a local optimum.

The first experiment investigates the transmission function, or the probability distribution of offspring fitness (section 2), for the best genotype found at each generation during the evolutionary run. Previous work analysed the probabilities of both the best individual and the population making the transition to a higher fitness [18]. Here we look at the change in the distribution of



(a) The GA population best and mean fitnesses over generations

(b) Population best fitness, and variation in multiple fitness evaluations of a single genotype (grey band).

Fig. 2. (a) Behaviour of the GA population best and mean fitnesses during the evolutionary run, (b) The variation in the population best fitnesses during the neutral epoch (generations 100-477) can be accounted for by the variation in multiple evaluations of a single genotype (shown as the grey band) [18].

offspring fitnesses over generations, calculating an approximation to the transmission function through recording the fitnesses of 100,000 applications of the genetic mutation operators to the best individuals of each generation (note that typical genotypes have roughly a few hundred loci, so we are exploring a significant fraction of the neighbouring space through 100,000 offspring). In particular, we are interested in whether the offspring fitness distributions highlight changes in the evolvability of the best individual. The high convergence of the population at each generation implies that changes in the evolvability of the best individual will reflect changes in the population evolvability.

The second experiment empirically tests the predictions made from the evolvability results: does the population evolvability predict the speed of evolution? Five populations from the evolutionary run (populations [100, 200, 300, 400, 477]) were used as initial populations for ten runs each of the evolutionary process (fifty runs in all), to see whether there was a difference in the time taken to reach good solutions. The next section describes the results from the two experiments.

5 Results

5.1 The Transmission Functions

For the best individual of each generation, the transmission functions were approximated through recording fitnesses from 100,000 applications of the mutation operator. Figure 3 shows the fitnesses of the best individual, the best mutation found and the mean mutation over generations. The graphs closely




Fig. 3. 100,000 mutations were applied to each of the best-of-generation individuals (the middle line), to approximate the transmission function distribution of offspring fitnesses. The best and mean mutation fitnesses (the top and bottom lines respectively) show no clear trend during the neutral epoch over generations 100-477.

follow the best individual fitness, rising sharply during the initial hill-climbing period, then staying roughly constant once the neutral epoch is reached around generation 100 (although there is a single high mutation fitness found just after generation 60).

Evolvability is equated with the likelihood of obtaining fit variants, and of not obtaining unfit variants, i.e. the upper and lower tails of the transmission function distribution. Whereas figure 3 tracked only the best and mean fitness of variants over generations, figure 4 shows the percentage of mutations above a fixed fitness (the fitness used here is 0.3 - other fitnesses show similar results), and figure 5 shows the percentage of mutations below two fixed fitnesses (we use 0.1 and 0.2, again other fitnesses show similar results).

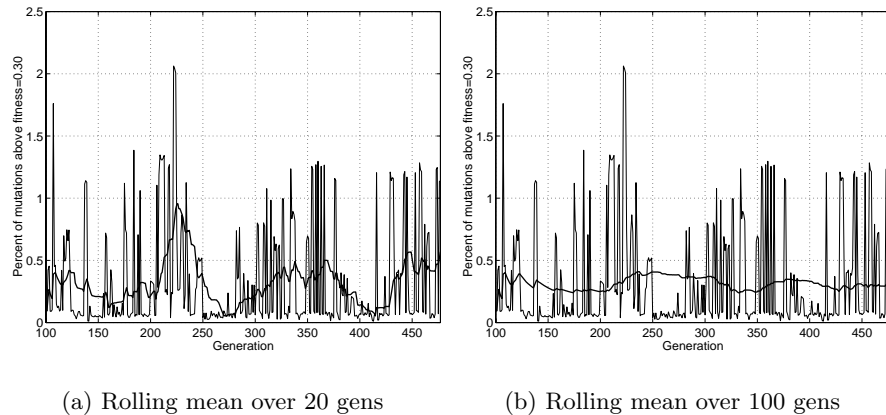


Fig. 4. The percentage of mutations above fitness 0.3 from the best-of-generation individuals during the neutral epoch (generations 100-477). **(a)** shows the short-term trends, while **(b)** shows the longer term trend.

We see that the percentage of mutations above a fitness of 0.3 (figure 4(a)) remains extremely low during the neutral epoch, with a maximum of 2%, and the mean remaining generally under 0.5%. Figure 4(a) shows the local trend, the

rolling mean over the last 20 generations, which does show signs of movement. However, the longer term trend (figure 4(b) shows the rolling mean over the last 100 generations) shows no such movement.




Fig. 5. The percentage of mutations from the best-of-generation individuals below fitness= [0.2,0.1]. **(a)** shows the short-term trends over the entire evolutionary run, while **(b)** shows the longer term trend during the neutral epoch (generations 100-477).

The picture for the percentage of mutations below two fixed fitnesses (figure 5) tells a similar story. The number of deleterious mutations falls quickly during the hill-climbing phase of evolution, then stays roughly constant during the neutral epoch; the short-term trends shown in figure 5(a) show the same movement as for the good mutations results, but the long term trend (figure 5(b)) shows no such movement. There is an interesting result from the end of the whole evolutionary run; as fitness dramatically increases at generation 478, the percentage of mutations below fitness 0.1 actually increases. This is discussed further in section 6.

5.2 Repeated Evolution

The results from section 5.1 show that for the search space at hand, there is no long-term trend for change in evolvability during the neutral epoch between generations 100-477. The second experiment tests whether there is any difference in speed of evolution from populations across the epoch. Five populations, from generations [100, 200, 300, 400, 477], were used as the initial populations for the genetic algorithm, and the evolutionary process repeated ten times for each population. Table 1 shows the number of generations required for 100% success on each of the evolutionary runs, while figure 6 shows the median number of generations taken to reach certain fitnesses. Statistical analysis shows no significant differences between the five sets of runs, supporting the hypothesis that there is

	Pop 100	Pop 200	Pop 300	Pop 400	Pop 477
Mean	2,008	2,096	1,901	1,680	3,024
Median	1,522	1,464	932	1,093	1,597
Maximum	4,713	7,696	>10,000	5,707	>10,000
Minimum	353	365	107	353	290

Table 1. Statistics on the number of generations required before the GA reaches 100% success, starting from the 5 populations saved on generations [100, 200, 300, 400, 477]. Note: Runs not reaching success in 10,000 generations were counted as 10,000 for averaging purposes. No significant differences were seen between the populations (Kruskal-Wallis analysis).




Fig. 6. Median from 10 runs of the number of generations required before the GA reaches a given fitness level, starting from the 5 populations saved on generations [100, 200, 300, 400, 477]. No significant differences were seen between the populations (Kruskal-Wallis analysis).

no change in the evolutionary potential, or evolvability, of the population across the neutral plateau.

6 Discussion

Many complex artificial-life problems such as evolutionary robotics show long periods over which fitness does not markedly increase, classically regarded as points where the population is stuck in local optima. Recent work has shown that the population may be moving along neutral networks in the search space during such periods, and also that the population evolvability may be changing.

In this paper, we have investigated a neutral epoch during an evolutionary process with complex genotype-to-fitness mapping (an evolutionary robotics experiment), and found no evidence for such change in evolvability. The distribution of offspring fitnesses from the best individuals of each generations showed no trend towards a higher probability of producing higher fitness offspring, and no trend towards a higher probability of not producing lower fitness offspring. Further, a second experiment in which populations from across the neutral epoch were used as initial populations for the genetic algorithm, showed no difference between the populations in the number of generations required to produce high fitness. This supports previous work [18], suggesting both that population evolvability does not change across the neutral epoch, and that populations from

across the epoch do equally well when straight hill-climbing is used instead of rerunning the GA.

This has implications for artificial-life techniques when evolving solutions to complex problems. In the evolutionary run we have studied here, there is no sense in which the population is doing something “useful” during the neutral epoch - it is not moving to better, i.e. more evolvable, areas of the fitness landscape. Thus the existence of neutral networks in the search space, which allow the evolutionary process to escape from local optima, does not necessarily provide any advantage; in this problem landscape the population does not evolve any faster due to inherent neutrality.

There is no doubt that the presence of neutrality can and does affect population dynamics during evolution, but it may well be that only in a certain class of search spaces does neutrality aid evolution. The use of genetic operators operating on several loci simultaneously, with the ability to alter the genotype length, may render the presence of neutrality less useful than in the fixed length and single-loci mutation genotype-phenotype mappings typically studied in more theoretical work.

Acknowledgements: The authors would like to thank the anonymous ECAL reviewer, Lionel Barnett, Inman Harvey, Andy Philippides, Anil Seth, Lincoln Smith 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 is funded by a British Telecom sponsored BBSRC Case award.

References

1. L. Altenberg. The evolution of evolvability in genetic programming. In K.E. Kinnear Jr, editor, *Advances in Genetic Programming*, chapter 3, pages 47–74. MIT Press, 1994.
2. 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 / Bradford Books, 1998.
3. C.L. Burch and L. Chao. Evolvability of an RNA virus is determined by its mutational neighbourhood. *Nature*, 406:625–628, August 2000.
4. L.L. Cavalli-Sforza and M.W. Feldman. Evolution of continuous variation: Direct approach through joint distribution of genotypes and phenotypes. *Proceedings of the National Academy of Sciences, USA*, 73:1689–1692, 1976.
5. 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: CEC2001*, pages 1–8. IEEE, Seoul, Korea, May 2001.
6. W. Hordijk. A measure of landscapes. *Evolutionary Computation*, 4(4):335–360, 1996.
7. P. Husbands. Evolving robot behaviours with diffusing gas networks. In P. Husbands and J.-A. Meyer, editors, *Evolutionary Robotics: First European Workshop, EvoRobot98*, pages 71–86. Springer-Verlag, Berlin, April 1998.

8. 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, December 1998.
9. 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.
10. T. Jones and S. Forrest. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In L.J. Eshelmann, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms (ICGA95)*, pages 184–192. Morgan Kaufmann, CA, 1995.
11. M. Kimura. *The Neutral Theory of Molecular Evolution*. Cambridge University Press, 1983.
12. M. Kirschner and J. Gerhart. Evolvability. *Proceedings of the National Academy of Sciences, USA*, 95:8420–8427, July 1998.
13. P. Marrow. Evolvability: Evolution, computation, biology. In A.S. Wu, editor, *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program (GECCO-99 Workshop on Evolvability)*, pages 30–33. 1999.
14. B. Naudts and L. Kallel. A comparison of predictive measures of problem difficulty in evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 4(1):1–15, April 2000.
15. M.E.J. Newman and R. Engelhardt. Effects of selective neutrality on the evolution of molecular species. *Proceedings of the Royal Society of London, B*, 265:1333–1338, 1998.
16. M. Nowak and P. Schuster. Error thresholds of replication in finite populations - mutation frequencies and the onset of Muller's ratchet. *Journal of Theoretical Biology*, 137:375–395, 1989.
17. 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: CEC2001*, pages 9–16. IEEE, Seoul, Korea, May 2001.
18. T.M.C. Smith, P. Husbands, and M. O'Shea. Neutral networks in an evolutionary robotics search space. In *Proceedings of the 2001 Congress on Evolutionary Computation: CEC2001*, pages 136–145. IEEE, Seoul, Korea, May 2001.
19. T.M.C. Smith and A. Philippides. Nitric oxide signalling in real and artificial neural networks. *BT Technology Journal*, 18(4):140–149, October 2000.
20. P. Turney. Increasing evolvability considered as a large-scale trend in evolution. In A.S. Wu, editor, *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program (GECCO-99 Workshop on Evolvability)*, pages 43–46. 1999.
21. 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.
22. G.P. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.
23. E. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325–336, 1990.
24. C.O. Wilke. Adaptive evolution on neutral networks. *Bulletin of Mathematical Biology*, 2001. Submitted.
25. C.O. Wilke, J.L. Wang, C. Ofria, R.E. Lenski, and C. Adami. Evolution of digital organisms lead to survival of the flattest. *Nature*, 2001. Submitted.