

Genetic Convergence in a Species
of
Evolved Robot Control Architectures

I. Harvey, P. Husbands, D. Cliff

CSRP 267, January 199

Cognitive Science Research Paper

Serial No. CSRP 267

The University of Sussex
School of Cognitive and Computing Sciences
Falmer
BRIGHTON BN1 9QH
England, U.K.

A version of this paper will be presented at the Fifth International Conference on Genetic Algorithms, University of Illinois at Urbana-Champaign, 17–22 July 1993



any one genetic operation. That is, any change of length should only be slight, the first SAGA principle¹.

Of course small changes in the short term can build up to arbitrarily large changes in the arbitrarily long term. But if in the long term the maximum genotype length of a population increases, say from g to G , then all the members of the later population will be descended from some of the earlier population, despite the g -dimensional earlier 'search-space' being minute in comparison to the later G -dimensional one. It follows that all, bar perhaps the very original, populations

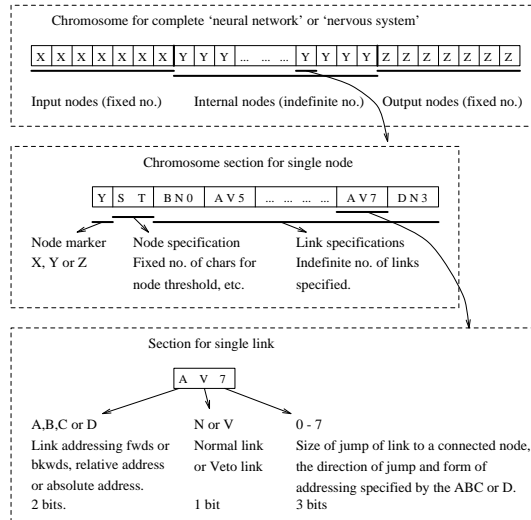


Figure : *The genetic encoding scheme*

mutation-selection cycle do not explore further away in an undirected fashion, but rather seek out any ridges of relatively high fitness in the landscape that may lead to even higher peaks (Fig. 1). To quote from a similar context in [5]:

“In conventional natural selection theory, advantageous mutations drove the evolutionary process. The neutral theory introduced selectively neutral mutants, in addition to the advantageous ones, which contribute to evolution through random drift. The concept of quasi-species shows that much weight is attributed to those slightly deleterious mutants that are situated along high ridges in the value landscape. They guide populations toward the peaks of high selective values.”

The fitness landscape metaphor is potentially misleading, in that high-dimensional spaces have properties very different from our intuitions about 2-D or n -D spaces. Whereas in a normal n -D landscape there can at best be a single ridge between two hills taking the direct shortest route, this is no longer the case in sequence space, which can be thought of as having n dimensions where n is the genotype length. As indicated in Figure 2, between two points Hamming distance d apart in binary genotype sequence space, there are $d!$ shortest paths, and far more slightly longer ones. This is why, in any high-dimensional landscape that is smooth enough for there to be some correlation in height or fitness between neighbouring points, any local optimum (other than the global one) is almost inevitably connected by short paths, without any intermediate points of much lower fitness, to other better regions — hyper-spatial bypasses.

In the n -dimensional sequence space, defined by binary genotypes of length n where Hamming-neighbours are connected, suppose that through mutation points up to Hamming-distance d from the current position can be sampled. There are $\mathcal{M}(d, n) = \sum_{i=1}^d n!/i!(n-i)!$ of these. For the current position to be effectively a local optimum from which escape is impossible, *all* of these points must be less fit. But with increase in n , increase in $\mathcal{M}(d, n)$ is roughly $\mathcal{O}(n^d)$; the higher the dimension, the more hyper-spatial bypasses there are.




Figure 4: *Typical path of a successfully evolved robot, which heads fairly directly for the*

sequential description of the properties of each node, first input nodes, then an arbitrary number of internal nodes, and then outputs; each preceded by a marker. For each node, an initial part specifies properties such as threshold values; then a variable number of groups each describe individual links from that node. Each group specifies whether it is a normal or veto link, and then the address of its target node is specified in either an absolute fashion (by distance from the first or last in genotype order) or in a relative fashion (by distance along the genotype from the current node). Although the genotype is in effect a direct description of the network, the phenotype which is evaluated is the *behaviour* of the robot in which the network is instantiated, and there is no direct relationship between the genotype and the behaviour. At some time in the future, we propose to change to a different form of encoding on the genotype, such that it constrains a developmental process which results in a network; in this way it is hoped that open-ended evolution, instead of being confined to just a lengthening list of units, will be able to build for itself a higher-level, perhaps hierarchical, system, reusing parts of the genotype many times just as a program calls procedures many times. For details of the genetic coding for vision, and of how the input visual signals are calculated in simulation by ray-tracing, see [2].

The task set in these trials is navigating within a simple closed cylindrical room, with black walls, and white floor and ceiling. Apart from the bumpers and whiskers on collision with the walls, the only inputs available to the robot for navigation are the two visual inputs, varying according to position and orientation w.r.t. the wall. On each trial the robot is started at a random orientation, and randomly placed near to the wall. The evaluation function is the sum over the limited time of the trial of a gaussian function \mathcal{G} based on the distance d from the centre of the room at each time-step:


$$\mathcal{G} = \exp(-d^2/c)$$

where the constant c ensures that \mathcal{G} is near-zero towards the walls. Implicitly this sets the goal of heading for the centre of the room as quickly as possible and then staying there. As well as noise in the internal nodes of the network, noise is included in the simulation of the physics of the world, including any collisions with the walls.

To induce robustness in the presence of noise — which it is hoped will be carried over into the real physical implementations, absorbing discrepancies between simulation and reality — each control network was evaluated over a number of trials, and the *worst* score achieved was used as the final score. Typically behaviour interpretable as ‘sensible’ appeared in less than 100 generations, using a population of size 60; a high-scoring trajectory is shown in Fig. 4, and the network that produced this behaviour in Fig. 5. For further analysis of such networks, and how they produce the behaviour, see []. The present paper concentrates on the issues of genetic convergence.

The evolutionary principles on which these experiments are based allow for incremental adding of tasks, requiring additional new behaviours or changes in old ones. The present analysis, however, is restricted to just a first task.

Genetic Convergence






Figure 8: All 60 genotypes in the 76th generation are listed according to their differences from the consensus sequence. Those 224 loci on the genotype with 100% agreement are ignored, the others displayed as '-' where they agree with the consensus, '*' where they differ. They are ordered in terms of Hamming distance from the consensus, distances shown on left.

degree of convergence had already been reached (Fig. 7). The evaluation based on the worst of a number of noisy trials can




Figure 9: A principal components analysis of all 60 members of the population in the 76th generation. The numbering here is in order of distance from the consensus sequence, itself numbered 0, hidden in the cluster at the origin, $(0,0)$.





Figure 10: Focusing on the central group shown in the previous figure — both axes are now in units of 10^{-3} .



gen54 4.9.54oup921
 TR -531875 -80.7430460.448(2R -5312289091 (2f.505(0.00)-61.89144BT / Td (76141 0)-6133047 6904 4242.2302f -5302 4.60.18 Td-87.50(19)1003047 0-.0355)Tj44en)T.00)-616d [(pp)1 5(-20219)10037109 0.97109844 BT45)Tj -13.7387.242(28)Tj




Figure 1 : A principal components analysis of the top scorers in every 5th generation from 36 to 76

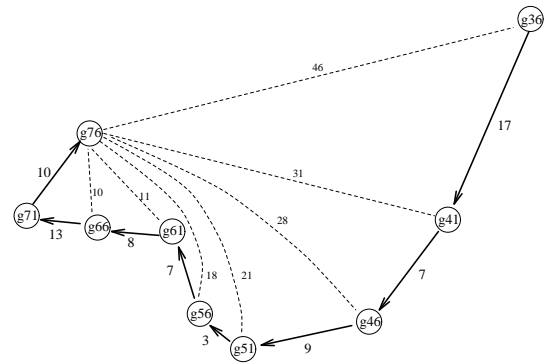


Figure 14: The same top scorers are shown with their Hamming distances both to the next-displayed one (solid lines), and to the one displayed for generation 76 (dashed).

The figure shows what could, if the fitnesses were relatively high, be considered as

practice.

Firstly, although genotypes need to be of arbitrary length, any changes in length which the genetic operators allow should be restricted to very gradual ones. Secondly, the population should be expected to be genetically converged, both in genotype lengths and in the alleles at each locus. Thirdly, in view of this convergence, mutation is promoted from its normal background role to one of greater importance in allowing continued improvement.

Examples have been given from successful runs of simulated robots with vision, using these principles. Principal components analysis has been introduced as a useful visual tool for analysing the movement of populations across sequence space. It has been shown that despite the genetic convergence, mutation is a sufficiently powerful force for genetic movement along 'ridges' to potentially fitter regions.

References

- [1] D. T. Cliff, I. Harvey, and P. Husbands. Incremental evolution of neural network architectures for adaptive behaviour. Technical Report CSRP 256, School of Cognitive and Computing Sciences, Univ. of Sussex, 199 .
- [2] D. T. Cliff, P. Husbands, and I. Harvey. Evolving visually guided robots. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *Proceedings of the Second International Conference on Simulation of Adaptive Behaviour (SAB92)*. MIT Press Bradford Books, Cambridge, MA, 1992. In Press.
- [] D. T. Cliff, P. Husbands, and I. Harvey. Analysis of evolved sensory-motor controllers. Technical Report CSRP 264, School of Cognitive and Computing Sciences, Univ. of Sussex, 199 .
- [4] Kenneth De Jong. Are genetic algorithms optimizers? In R. Männer and B. Manderick,

- [9] D.S. Hirschberg. A linear space algorithm for computing maximal common subsequences. *Communications of the A.C.M.*, 18(6): 41–44, 1975.
- [10] P. Husbands and I. Harvey. Evolution versus design: Controlling autonomous robots. In *Integrating Perception, Planning and Action, Proceedings of 3rd Annual Conference on Artificial Intelligence, Simulation and Planning*, pages 109–146. IEEE Press, 1992.
- [11] Stuart