

Evolutionary Robotics and SAGA:
the case for Hill Crawling
and Tournament Selection

Inman Harvey

CSRP 222, 1992

Cognitive Science Research Paper

Serial No. CSRP 222

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

This paper appears in:
C. Langton, editor, *Artificial Life 3 Proceedings,*
Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol XVI.

Evolutionary Robotics and SAGA: the case for Hill Crawling and Tournament Selection

Inman Harvey

School of Cognitive and Computing Sciences

University of Sussex

Brighton BN1 9QH, England

email: inmanh@cogs.susx.ac.uk

Abstract

This paper will look at an evolutionary approach to robotics; partly at pragmatic issues, but primarily at theoretical issues associated with the evolutionary algorithms which are appropriate. Genetic Algorithms are not suitable in their usual form for the evolution of cognitive structures, which must be in an incremental fashion. SAGA — Species Adaptation Genetic Algorithms — is a conceptual framework for extending GAs to variable length genotypes, where evolution allows a species of individuals to evolve from simple to more complex.

In the context of species evolution the metaphor of hill-crawling as opposed to hill-climbing is introduced, and appropriate mutation rates discussed. On both pragmatic and theoretical grounds, it will be suggested that there are good reasons for using Tournament Selection in evolutionary robotics.

1 Why Evolutionary Robotics?

Subsumption-style cognitive architecture for robots (Brooks 1986, Brooks 1991) in theory analyses independent behaviours of a robot, and ‘wires them in’ largely

havioural modules for artificial nervous systems, or artificial embryology. Beer (Beer and Gallagher 1991) used GAs to synthesize a walking behaviour for a six-legged agent. In a more traditional robotics context, mention is made of an evolutionary approach in (Barhen *et al.* 1987).

Recently the Japanese government research laboratories, ATR in Kyoto, have set up a well-funded research group for Evolutionary Robotics in their Evolutionary Systems department. Similar work is pursued at ETL in Japan, and there is interest from Japanese industry; Mitsubishi sponsored a symposium on Evolutionary Robotics in March 1993. At the Simulation of Adaptive Behavior 1992 conference (Meyer *et al.* 1993) in Hawaii, a group of papers were closely related to this field. The Evolutionary Robotics Group at Sussex has been since 1992 artificially evolving control systems for mobile robots — co-evolved with sensor attributes — for visual navigation tasks. This work started with simulations of a real physical robot, and is now using a specialised piece of hardware allowing real vision to be used in a robot that can have a succession of control systems rapidly and automatically evaluated in sequence (Harvey *et al.* 1993, Cliff

lengths, e.g.

hood increase in numbers of mutations is monotonically related to decrease in fitness. Figure 6 demonstrates the effect of Muller's ratchet when mutation is high enough to cause loss of information. Figure 7 sketches the effects when mutation is high enough (without bringing Muller's ratchet into play) for some elements of the population to crawl down the hill far enough to reach a ridge of high selective values. As discussed in (Eigen *et al.* 1988), this results under selection in a significant proportion of the population working their way along this ridge, and making possible the reaching of outliers further in Hamming-distance in that particular direction from

constant α ; in (Hesser and Männer 1991) that after earlier higher values should decrease exponentially towards $m_{opt} = \alpha'/(N\sqrt{l})$; and in (De Jong 1975) quoted in (Hesser and Männer 1991) as recommending $m_{opt} = 1/l$.

But these rates, and also the error threshold given by Eigen and Schuster, are based on particular assumptions about the selective forces on the population. The use of tournament selection provides a significantly different range of selective forces to a population, which means the above analyses cannot be relied on. However, in all simulations of hill-crawling where different mutation rates are tried, a typical U-shaped curve is found, giving the shortest time needed to reach another hill at around the mutation rates suggested by theory.

members of the same distance i results in a winner i and the loser replaced by $i + 1$.

For maximum hill-crawling without losing the master-sequence (of distance 0) from the population, the long-term fate of this master-sequence should be considered. It can be seen that all tournaments between a 0 and a 0 result in the loss of one 0 to the population, and there is no other way in which 0s can be gained. If all tournaments are constrained to be between two different individuals, then r_0 will soon reduce to one member which will thereafter survive alone for ever. This member, the 'wild-type', will win all its tournaments and continually replenish the flow of mutants down the hill away from it. Histograms of results from populations of various sizes run in a computer simulation is shown in figure 10. An equation which allows one to iteratively derive the expected size of each class is derived in Appendix A.

If the same individual can be chosen twice for a tournament, resulting in the replacement of itself by a mutated copy, then the wild-type will eventually be lost through just such an incident, and Muller's ratchet will start to operate. However, if tournaments are between different individuals, then the wild-type will never be lost, whatever the size of μ . We thus have a selection mechanism which can move the bulk of the population crawling down the hill as much as is desired, without ever encountering the error threshold of Eigen and Schuster. The banding into multiples of μ can be broken up by alternating between two different integers for μ .

But careful ...

There is a dangerous potential flaw in this. We are relying on the choice of winner of a tournament being 100% reliable, and in the context of evolutionary robotics, as discussed earlier, this may very easily not be the case. If the reliability of choice is $p < 1$, then sooner or later the wild-type will be lost and Muller's ratchet will start. A possible counter to this will be to only mutate the replica with probability $q < 1$, and otherwise leave it unchanged. In Appendix B it is shown that this will save the situation in an infinite population for values of $q < (2p - 1)/p$, this being independent of the value of μ . For example, if $p = 0.9$, we should have $q < 0.888 \dots$. In the case of a finite population, q should be reduced

of repair mechanism protecting against Muller's ratchet (Maynard Smith 1978).

With tournament selection, candidates for recombination would be the winners of two separate tournaments, and the two offspring, after crossover and mutation, can replace the two losers. In general, the crossover will produce one offspring closer to the wild-type than the average of the two parents' distances, and another offspring further away than this average; after which mutation adds its toll. This constitutes a force producing a restorative flow towards the wild-type, allowing larger mutation rates without loss of the current local optimum. Simulations confirm this.

There are

editor, *Proceedings of the First European Symposium on Artificial Neural Networks, ESANN'93*, pages 39–44. D factio Publishing, Brussels, 1993.

[Cliff *et al.* 1993b] D. T. Cliff, P. Husbands, and I. Harvey. Evolving visually guided robots. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 3, Proceedings of the Second International Conference on Simulation of Adaptive Behaviour (SAB92)*, pages 374–383. MIT Press Brad-

populations, mutation frequencies and the onset of Muller's ratchet. *Journal of Theoretical Biology*, 137:375–395, 1989.

[PRANCE 1991] PRANCE. Perceptive robots: Autonomous navigation and cooperation through evolution. Unpublished research proposal, PRANCE con-

Setting gains equal to losses, and dividing by a , we have

$$aq + 2(1 - a)(1 - p) = 2(1 - a)p(1 - q)$$

$$a(4p + q - 2pq - 2) = 2(2p - pq - 1)$$

We can assume that $p > 0.5$, say $p = 0.5 + s$ for positive s . The factor on the l.h.s. of the equation, multiplying a , then becomes $(2 + 4s + q - q - 2sq - 2)$, which reduces to $2s(2 - q)$. We know that

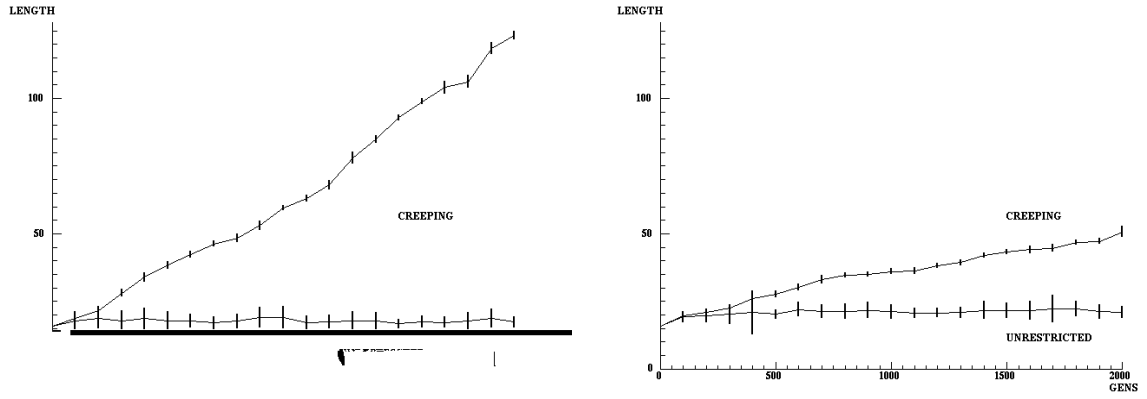


Figure 3: Average genotype lengths against generations; vertical bars show standard deviations. Effects of 'creeping' and 'unrestricted' increase-length genetic operators on a population with the same fitness conditions, epistasis $K = 2$. Left graph, linkage with neighbouring genes. Right graph, random linkage.




Figure 4: As a species evolves through SAGA space, the search for higher fitness only takes place in a very local search space around the current focus of the species.





Figure 5: *The opposing forces of mutation and selection on a population centred around a local optimum, where Hamming distance from master sequence is directly related to fitness ranking.*



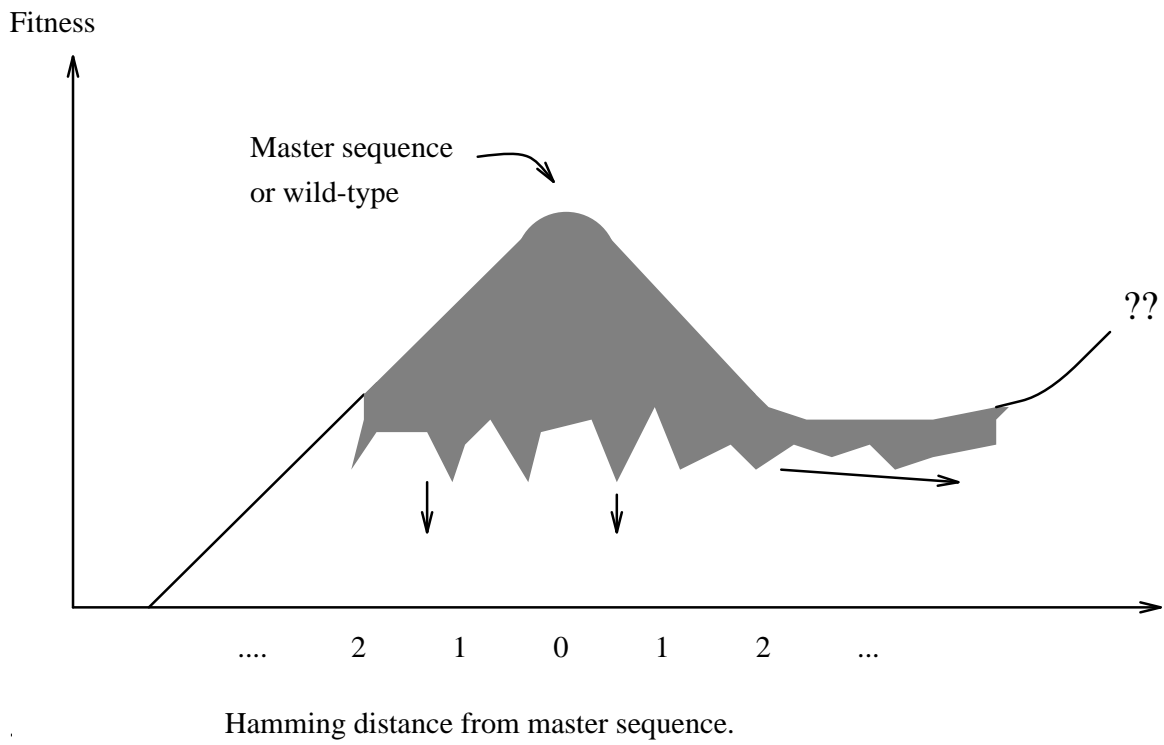


Figure 7: *If the population can crawl down the hill far enough to reach a ridge of relatively high fitness, it will spread along it, potentially reaching new hills.*



_____ - - - .
* * * m = 3 mutns

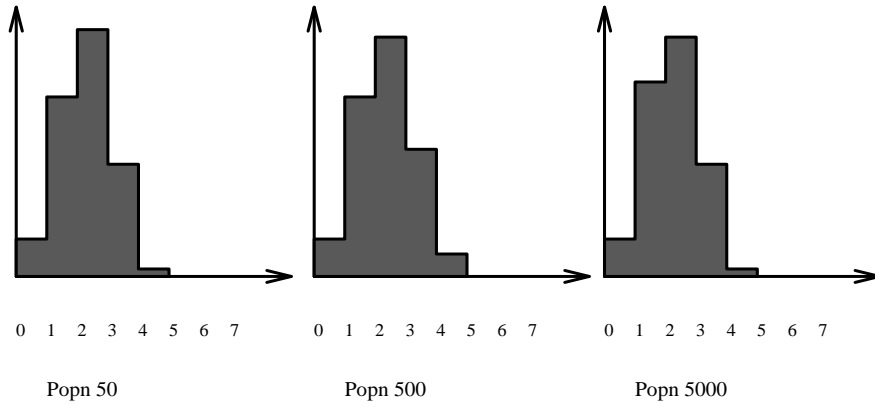


Figure 11: *Similar simulations to those shown in previous figure, except that genotypes are of length 100.*

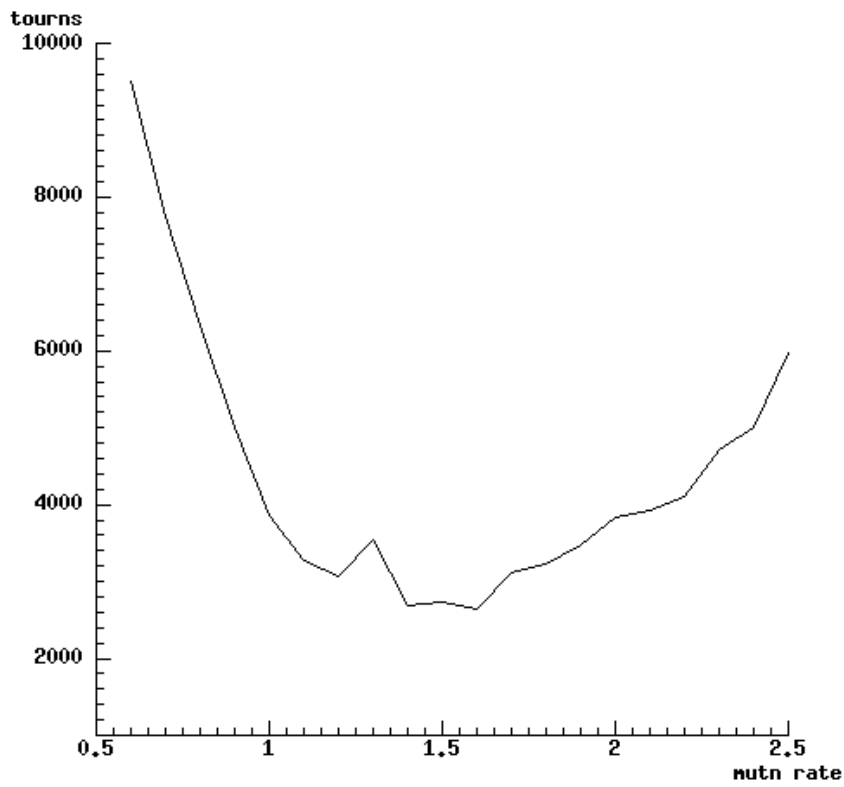


Figure 14: *Number of tournaments for a population centred at one 'hill-top' to have a first member reach a nearby hill-top. Rate specified is the average number of mutations per genotype. Recombination has been used.*