

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 independently from sensor input to motor output. Later ‘wired in’ behaviours interact with earlier ones either through the environment or by suppression or inhibition mechanisms.

As the number of layers or modules within such an architecture increases, the number of potential interactions increases much faster. The foresight needed to design by hand will soon be outstripped by the exponentially increasing complexity (Husbands and Harvey 1992). As with so many AI problems, progress beyond relatively toy domains becomes infeasible. So the possibility of automatic evolution of the cognitive architecture without explicit design becomes very attractive.

Natural evolution is the existence proof for the viability of this approach, given appropriate resources.

Genetic Algorithms (GAs) are a form of search technique, modelled on Darwinian evolution, primarily used for function optimisation (Goldberg 1989). An evolutionary approach to robotics necessarily means an incremental approach, and yet this is something that standard GAs cannot handle; in fact standard GAs, though borrowing ideas from natural evolution, are themselves of no use for *applied* evolution. Below, I shall introduce SAGA, which is a framework extending GAs for applied evolution. This will demonstrate that necessarily, in incremental evolution of, e.g., a robotic cognitive system, the population will be genetically converged; in other words the cognitive structures of all the robots will be fairly similar, and the genotypes will be positioned around some hill in a fitness/genotype landscape.

From this it will follow that the evolutionary search process will involve hill-crawling as much as hill-climbing in the fitness landscape; this has implications for mutation rates and for the selection mechanism for the evolutionary algorithm. This in turn leads to theoretical reasons why tournament selection is appropriate; there are in any case *practical* reasons why it might be appropriate for evolutionary robotics.

2 Related Work

Evolutionary approaches, often using variants of GAs, to Artificial Life in simulations have been widespread (Langton 1989, Langton *et al.* 1991). Evolutionary robotics was proposed for philosophical reasons in (Cariani 1989). It is only recently that serious proposals have been made to use evolutionary approaches to real-world robots (PRANCE 1991, Brooks 1992, Husbands and Harvey 1992). Earlier, a student of Brooks discussed some of the issues involved, with reference to subsumption architectures, in (Viola 1988). De Garis (de Garis 1992) proposed using GAs for building be-

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 *et al.* 1993b). The control systems evolved are noisy dynamic recurrent networks (Cliff *et al.* 1993a). We agree with Beer (Beer 1992) in his advocacy of a dynamical systems perspective on autonomous agents, and are generally sympathetic to the enactive approach to cognition (Varela *et al.* 1991); in this paper these concerns will not be elaborated on.¹

3 Simulation versus Reality

Any evolutionary technique is going to need large numbers of trials of robots, and practical constraints mean that these should be done on simulated robots if this is viable. Traditionally, and for good reason, those who have built real robots have tended to scorn simulations as implicitly assuming that all the really hard real-world problems have been solved. To quote from (Brooks 1991)

First, there is no notion of the uncertainty that the real world presents . . . Second, there is a tendency to not only postulate sensors which return perfect information (e.g., the cell ahead *contains food* — no real perception system can do such a thing) but there is a real danger of confusing the global world view and the robot's view of the world. . . .

It is standard practice for a commercial pilot to convert to flying a new model of a plane by training in

¹Papers from the Sussex Evolutionary Robotics Group may be obtained from the present author, or by anonymous ftp from 192.33.16.70; files pub/reports/csrp/csrp???.ps.Z, for ??? = 219,220,221,222,223,256,264,265,267,278.

a flight simulator, so that the first real flight in the new plane is carrying passengers. Two things should be noted: the conversion is from one plane type to another similar one, and commercial flight simulators are so complex that, although cheaper than flying the real thing, the cost is not of a totally different order.

If robot simulations are being used to save money, then they are likely to be the equivalent of a flight simulator on a workstation rather than a commercial airline simulator. The usefulness of robot simulations may be compared with the usefulness of practising on a personal flight simulator program for a pilot learning for the first time to fly a real light aircraft. There can be no substitute for experience of the real thing, though the simulations can be of benefit.

The simulator is limited by the knowledge of the programmer of relevant factors to be included, and if, for instance, no account of the effects of wind shear is put into a flight simulator, the first encounter with this in the real world will be hazardous. So any benefits brought to an evolutionary approach by using simulations will inevitably have to be paid for by major effort being put into the realism of the simulator. Trials of evolved architectures on real robots will have to be carried out at frequent intervals for the dual purpose of validating the fitnesses *and* providing feedback for improvement of the simulator. At Sussex we have found that doing simulations of vision, using ray-tracing, has been so computationally expensive that we have built hardware to allow automation of multiple evaluations with real vision (Cliff *et al.* 1993b); it is faster.

4 SAGA and gradual increase in complexity

Some hints from natural evolution have been used by the GA community to produce effective search techniques for complex multi-dimensional search spaces. But this use of GAs for function optimisation is problem-solving in what is, although enormous, a pre-defined space of possibilities of known size — this size being a maximum of a^l when genotypes are of length l with a possible alleles at each position. But the most impressive feature of natural evolution is how over aeons organisms have evolved from simple organisms to ever more complex ones, with associated increase in genotype lengths. This aspect of evolution has been completely ignored in the standard GA literature. GAs have been adapted to problem-solving, and the problem-solving metaphor or frame of mind is, I believe, much of the time inappropriate for considering both natural evolution and potential robot evolution.

The theoretical underpinning for GAs, Holland's Schema Theorem (Holland 1975, Goldberg 1989) is no longer valid when the genotypes within a population vary in length. Some GA systems have used variable

lengths, e.g. Smith’s LS-1 classifiers (Smith 1980) and Koza’s genetic programming (Koza 1990); but the analyses offered in these two examples do not satisfactorily extend the notion of a schema such that schemata are preserved by the genetic operators.

The conceptual framework of SAGA was introduced in 1991 in order to try to understand the dynamics of a GA when genotype lengths are allowed to increase (Harvey 1992b). Working with a finite population, a standard GA often starts with a random distribution that spans the whole search space; the genetic operators, particularly recombination, shift the population over successive generations until hopefully it converges around some optimum (see figure 1). If genotype lengths are going to be allowed to increase indefinitely, then there is no finite search space of pre-determined size, and this picture can no longer be valid. In (Harvey 1992b) it is shown, using concepts of epistasis and fitness landscapes drawn from theoretical biology (Kauffman 1989), that progress through such a genotype space will only be feasible through relatively gradual increases in genotype length. A general trend towards increase in length is associated with the evolution of a *species* rather than global search. The word *species* I use to refer to a fit population of relative genotypic homogeneity.²

In contrast to the goal-seeking metaphor of figure 1, a journey through SAGA space can be characterised in the form of figure 2. The conclusion of (Harvey 1992b), that only gradual increases in genotype length are likely to be viable, means that the finite resources of the population in searching around its current focus should be concentrated on just such gradual increases. The analysis given was supplemented by experimentation using an NK model (Kauffman 1989) which gave confirmatory results shown in figure 3. This conclusion accords well with, for instance, Brooks’ approach of ‘wiring in’ new behaviours one at a time, and waiting until current behaviours are thoroughly debugged before ‘wiring in’ the next.

In general, the ‘problem-solving’, or ‘goal-seeking’ metaphor for evolution (figure 1) is misleading. Within a SAGA space, however, it can still be useful to use this metaphor in the restricted sense of searching around the current focus of a species for neighbouring regions which are fitter, or in the case of neutral drift, not less fit (figure 4). Such a search takes place through application of genetic operators such as mutation or change-length.

A change-length genetic operator, when subject to the restrictions of only allowing small changes in length in any single application, can be translated into a form equivalent to mutation by sleight-of-hand. In the case

²This is only indirectly related to a biological definition of the word. However it follows from my definition that crosses between members of the same species have a good chance of being another fit member of the same species; whereas crosses between different species will almost certainly be unfit.

of an increase by one gene in a binary genotype with possible alleles 0 and 1, this gene can be considered instead to have 3 possible values 0, 1, and $-$, where the latter value is equivalent to ‘absent’; the new appearance of this gene can be considered as a mutation from $-$ to either 0 or 1. If more than one gene appears in one reproduction event, then this would be equivalent to a simultaneous set of mutations, which for low mutation rates is highly unlikely.

For this reason in the next sections of this paper we will consider mutation only, though bearing in mind that through this sleight-of-hand it will be possible to extend some of the conclusions to an increase-length operator which works at very low rates. In addition we will start by assuming that there is no recombination; i.e. reproduction is asexual. Having sketched out the important factors using this assumption, the impact of recombination on this sketch will be discussed.

5 Species hill-crawling

Usually in GAs premature convergence of a population is something to be avoided. In SAGA, we are continually working with a converged population, and are interested in encouraging search around the local focus while being careful not to lose the gains that were made in achieving the current *status quo*. In the absence of any mutation (or change-length) genetic operator selection will concentrate the population at the current best. The smallest amount of mutation will hill-climb this current best to a local optimum. As mutation rates increase, the population will spread out around this local optimum, searching the neighbourhood, but if mutation rates become too high then the population will disperse completely, and the search will become random with the previous hill-top lost.

The problem is that of *Muller’s ratchet* (Maynard Smith 1978). Call the genotype that represents the very peak of the hill the master sequence (or the ‘wild-type’). As a converged population, the other members will be quite close in Hamming distance to this master sequence, and hence far more mutations will increase this Hamming distance than will decrease it. The only force opposed to this pressure is that of selection preferentially reproducing the master sequence and its nearest neighbours in sufficiently large numbers to allow an occasional copy of the master with fortuitously no mutation; the other possibility, of fortuitous back-mutation from a near neighbour to the master sequence, is so small as to be usually negligible.

Figure 5 sketches the effects of mutation on a population centred around a local optimum. The vertical axis represents fitness or selective values. The horizontal axis indicates distance in sequence space from the master sequence at the top of the hill. The shape of the hill indicates the assumption that within this local neighbour-

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 the master sequence.

If any such outliers reach a second hill that climbs away from the ridge, then parts of the population can climb this hill. Depending on the difference in fitness and the spread of the population, it will either move *en masse* to the new hill as a better local optimum, or share itself across both of them.

So in a SAGA setup of evolution of a converged species, we want to encourage through the genetic operators such hill-crawling down towards ridges to new hills, subject to the constraint that we do not want to lose track of the current hill. To quote from (Eigen *et al.* 1988)

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.

6 SAGA and mutation rates

Although progress of a species through a fitness landscape is not discussed in the standard GA literature, in theoretical biology there is relevant work in the related field of molecular quasi-species (Eigen and Schuster 1979, Eigen *et al.* 1988). In particular, analysis of ‘the error catastrophe’ shows that, subject to certain conditions, there is a maximum rate of mutation which allows a quasi-species of molecules to stay localised around its current optimum. This critical maximum rate balances selective forces tending to increase numbers of the fittest members of the population against the forces of mutation which tend, more often than not, to drag offspring down in fitness away from any local optimum. But a zero mutation rate allows for no further local search beyond the current species, and other things being equal increased mutation rates will increase the rate of evolution. Hence if mutation rates can be ad-

justed, it would be a good idea to use a rate close to but less than any critical rate which causes the species to fall apart. A further possibility, in the spirit of simulated annealing, is to temporarily allow the rate to go *slightly above* the critical rate — to allow exploration — and then cut it back again to consolidate any gains thus made.

For an infinite population, in the particular context of molecular evolution, Eigen and Schuster show (Eigen and Schuster 1979) that these forces just balance for a mutation rate

$$m = \frac{\ln \sigma}{l}$$

where l is the genotype length and σ is the *superiority* parameter of the master sequence — the factor by which selection of the master sequence exceeds the average selection of the rest of the population. The diagrams they show for the very sharp cutoff at the critical rate refer to a fitness landscape with a single ‘needle’ peak for the master sequence, all the rest of the population taken to be equally (un-)fit; where the hill slopes more gently from the master sequence, the cutoff is less abrupt (figure 14). For typical values of σ between 2 and 20, the upper limit of mutation before a quasi-species ‘loses its grip’ on the current hill would be between $0.7/l$ and $3/l$. When the population is of finite large size, yet small enough for stochastic effects of genetic drift to start having an effect, the same overall picture holds except for a reduction in this critical mutation rate (the ‘error threshold’)(Nowak and Schuster 1989). Expressed in terms of the single-digit accuracy of replication $q = 1 - m$, then the critical value of q for a population of size N is related to that for an infinite population by

$$q_N = q_\infty \left(1 + \frac{2\sqrt{\sigma - 1}}{l\sqrt{N}} + \dots \right)$$

Nowak and Schuster suggest that the approximations made in deriving this equation mean that it should only be relied on for large populations of significantly more than 100. Nevertheless, the presence of l , genotype length, in the denominator suggests that for genotypes of length order 100, and populations of size order 100, the error threshold will be extremely close to that for an infinite population.

Since the natural logarithm of the *superiority* parameter σ which enters into the equation for m , variations in this of an order of magnitude do not affect the error threshold as significantly as variations in genotype length. In conventional GAs, choice of mutation rates tends to be a low figure, typically 0.01 or 0.001 per bit as a background operator, decided upon without regard to the genotype length. This despite suggestions from experimentation in (Schaffer *et al.* 1989) that optimal rates $m_{opt} = \alpha / (N^{0.9318} l^{0.4535})$, for some

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.

In figure 13, an artificial fitness space was created with two fitness hills. The fitnesses of any point in the space was given by the (negative of) the minimum of the Hamming distances to each hill. The population was started entirely concentrated at the peak of one hill, and then allowed to drift under various mutation rates, with tournament selection, until the other hill was first arrived at by a single member. Obviously with zero mutation it would never be found; with excessively large mutation, random search would result, also taking excessively long times. The typical U-shaped curve is shown in figure 14. In this case the difference between with and without recombination was only marginally in favour of the former.

Recently in the GA community there has been some discussion of the surprising success (in some circumstances) of what has come to be called *Naive Evolution*; i.e., mutation only, contrary to normal GA folk-lore which emphasises the significance of crossover. It would be interesting to check on those circumstances where it has been found useful, and see whether the population is in fact converged, with hill-crawling being the motive force for progress. The optimal mutation rates that are appropriate when hill-crawling is feasible have been obscured in the GA literature by the usual practice of quoting mutation rates per bit or per symbol, rather than per genotype. It is the optimal mutation rates per genotype that can be found within a band that is nearly invariant over all genotype lengths.

7 Tournament Selection for practical reasons

Realistic simulations take time to run, and it will be necessary to do a large number in parallel. As each simulation is complex, parallel machines with SIMD are no use, and for instance an individual workstation per simulation would be appropriate. In almost all networks of workstations there is a vast unused computational capacity which can be used effectively by running background processes. It then becomes attractive to use an evolutionary algorithm which allows asynchronous processing of the individuals.

Standard GAs tend to evaluate the whole current population, select from these and apply genetic operators to produce the next generation. A steady-state algorithm such as GENITOR (Whitley 1989) replaces an individual at a time rather than a generation at a time. But since it always replaces the currently worst member of the population, it requires global communication of statistics about the whole population before carrying out such a replacement. In a network of processes running asynchronously, with the possibility of individual machines being down for periods of time, this negates some of the benefits of parallelism.

Tournament selection operates by taking two (or sometimes more) members of the population chosen at random, and choosing the best of this tournament to contribute genetic material to a new individual. There are a variety of ways to choose which old individual should sacrifice its place for the new. In a tournament of size two a copy of the winner, after application of genetic operators such as mutation or crossover (recombination), could replace the loser, or replace a randomly chosen member of the population. Sexual recombination can take place between the winners of two different tournaments.

The practical advantage of this procedure when using a network of workstations is that it can be truly asynchronous and decentralised. Communication between machines is largely limited to occasional passing of genotype strings. If one, or a whole group of machines, slows due to loading or even is down for a time, the algorithm can carry on regardless on the remaining machines. It even becomes feasible to use several different networks only occasionally communicating with each other by electronic mail. Most tournaments would be 'local' within one network, as it is only necessary for there to be one transmission of genetic material per 'generation' (i.e. number of tournaments equal to population size) between two otherwise isolated genetic pools for them to stay together. It thus becomes possible to use enormous amounts of computing power currently little used on present facilities.

For robotics applications it should be noted that the use of tournament selection reduces the evaluation of the robots to a very simple question: of two given robots which is (probably) the 'best'? If the robots are tested (in reality or in simulation) on a series of tasks of increasing complexity in a noisy environment, then the evaluation will become something like: which of the two got further before stopping? It is the power of evolution that complexity can be built up through a succession of such trivial 'questions and answers' each containing at the very maximum 1 bit of information. As with any selection mechanism which is equivalent to ranking, it is not necessary to have an evaluation function that returns a scalar value, which may simplify matters greatly.

However, the extent to which noise in tournament selection — occasionally selecting the less fit by mistake — affects the stability of a species will be considered below.

In the game of *Twenty Questions* each reply can provide a maximum of 1 bit of information, hence at the end potentially discriminating between 2^{20} choices; but answers to inappropriate questions will provide much less than 1 bit, and frequently zero information. In an evolutionary approach to robotics it is the task of the genetic operators such as mutation, change-length, and recombination to generate new test cases that ‘ask new questions’ which tend to be appropriate. The following sections look at the problem from a theoretical perspective, and demonstrate that tournament selection can aid the genetic operators in this task.

8 Tournament Selection for theoretical reasons

Tournament selection relies on selecting randomly from the population a small number of contestants for a tournament, and taking the winner for further genetic processing to contribute to the future population. We will only consider a tournament of size 2, but within this constraint there is still a wide range of possibilities.

Tournament selection in GAs are generally used on a generational basis; i.e. the current generation is held fixed, and an appropriate number of tournaments are held to build up a genetic pool for mating or mutation, after which this becomes the new generation. It can be shown (Goldberg and Deb 1990) that this is equivalent to a ranking scheme³, in which the highest rank on average contributes 2 members to the genetic pool, the middle rank 1 member, and the bottom rank none. In Eigen and Schuster’s terms, the superiority $\sigma = 2$. Such a method loses the advantage of asynchronous parallel computing mentioned earlier.

A steady-state method, which can be done asynchronously, involves replacing the offspring derived from the winner after each tournament. In the case of a draw, the winner is chosen at random. The interesting question is, who is to die so as to make way for the offspring, and some possibilities are (see figure 8);

1. A randomly chosen member of the population.
2. A randomly chosen member of the tournament.
3. The loser of the tournament.
4. A subtle variation — do *not* reproduce from the winner, but remove the loser of the tournament and replace it by an offspring of a randomly chosen member of the population.

³The advantages of a ranking scheme are discussed in (Whitley 1989)

The effects of these methods will be described in terms of a notional generation, when a number of tournaments equal to the population size have been run. The first one is similar in effect to the original generational basis, except that the superiority of the first rank is approximately e rather than 2 — based on the fact that $(1 + 1/N)^N$ tends towards e for large N . This superiority only holds true while each member of the population is in general separately ranked, and ceases being valid as soon as a significant number of the tournaments are draws — in other words, Eigen and Schuster’s analysis no longer becomes directly applicable as their selection mechanism is radically different from tournament selection.

The second method has the same effect as the first — a randomly chosen member of the tournament will, in the long run, be just as ‘averagely fit’ as a random member of the whole population. Since in half of these tournaments the winner will be replacing itself, then these N such tournaments are equivalent, as far as σ is concerned, to $N/2$ tournaments using method 3. Hence a full notional generation of N tournaments with method 3 yields a σ value of e^2 . Surprisingly, method 4 gives the same value of σ as methods 1 and 2.

Method 3 will be the one discussed here, and the effect of mutation rates on a converged species (with a binary genotype) will be assessed for three different scenarios (see figure 9):

1. A mutation rate of μ bits flipped per genotype, where μ is a small integer and the genotype length is so long that the possibility of back-mutations towards the current master-sequence can be ignored.
2. A similar mutation rate of μ bits per genotype on a genotype of length l , with the possibility of back-mutations.
3. An average mutation rate of μ bits per genotype, calculated independently at the rate of μ/l at each locus.

8.1 Long genotypes

In the first scenario, we can classify each member of the population by the Hamming-distance from the master sequence. This will increase by μ at each replication, giving possible distances of $0, \mu, 2\mu, \dots$, so without loss of generality we need only consider $\mu = 1$. In the context of hill-crawling, our interest is in how the population distributes itself within different mutant-classes of size r_i , whose members have Hamming-distance i from the master-sequence. Given a tournament between members of distance i and j , for $i < j$ the winner i will remain in the population, and j will be replaced by a mutant of distance $i + 1$. A tournament-draw between

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 further to allow for the stochastic effects of genetic drift. The analysis of this case is not attempted here.

8.2 Shorter genotypes

Experiments run in simulation with populations of various sizes with genotypes of length 100, and with a mutation rate of exactly one bit flipped at replication, demonstrate what happens when the possibility of back-mutation is no longer negligible. The histograms in fig-

ure 11 demonstrates a similar shape to those in the previous case, except that much more of the population stays close to the wild-type. It can be seen that in the case of unreliable choice discussed above, this effect will supplement that of any given value of q in countering the loss of the wild-type.

8.3 Mutations assessed independently at each locus

When the same experiment is run on genotypes of length 100, but in this case instead of exactly one bit flipped per genotype, there is a 1/100 chance of flipping at each locus on the genotype, the results shown in figure 12 are startlingly different. The reason for this is that although the expected number of mutations per genotype is 1, this is made up from a probability of about $1/e$ of no mutations, about $1/0.99e$ of one mutation, about $1/1.98e$ of 2 mutations, and so on. The significant probability of there being no mutation has a similar effect to that given by the deliberate introduction of a probability $1 - q$ of no mutation, discussed above in the long genotype scenario and in Appendix B. Appendix C shows that the proportion of the population expected to be the wild-type is $2/(1 + e^\mu)$, which when $\mu = 1$ gives a value of about 0.538. When there is a probability p of making a mistake in a tournament, it is also shown in this appendix that in an infinite population the wild-type will not be lost if $p > e^\mu/(1 + e^\mu)$, which for $\mu = 1$ gives a minimum value for p of 0.731. To use this equation in the other direction, if it is known that $p > 0.9$ then the maximum value of μ would be $\ln 9 \approx 2.19$. These figures would be worse in a finite population due to stochastic effects.

9 SAGA and recombination

It has been suggested above that the application of a change-length genetic operator at the very low rates required in SAGA can be treated in a similar fashion to low mutation rates, although the sleight-of-hand used is equivalent to increasing the number of possible alleles at the relevant loci. In the context of adjusting genetic operators so as to be able to influence hill-crawling without losing the current wild-type, the introduction of recombination makes a major impact.

One virtue of recombination within a species is that when two different favourable but improbable mutations take place within two different members of a population, then sexual replication can at a stroke produce an individual combining both; whereas asexual replication would require both improbable events to occur within the same single line of descent. It is this virtue which is stressed in standard GAs, yet here we will concentrate on another virtue — the other side of the same coin — which is that recombination functions as a form

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 practical computational problems in dealing with recombination with the variable-length genotypes that are necessary in evolutionary robotics; given a crossover point in one parent genotype, where should the crossover point in the other parent genotype be? A discussion of this problem, and an algorithm which provides an efficient technique, are presented in (Harvey 1992a).

10 Elitism in noisy finite populations

We have been seeking ways of avoiding loss of the wild-type, while promoting appropriate exploration. In GAs, the policy of retaining the best, unchanged, for inclusion in the next generation is known as *elitism*, and we have seen that tournament selection gives you elitism for free, when the tournaments are 100% reliable. In the presence of noise, we can counter the operation of Muller's ratchet by some of the measures mentioned above which guarantee preservation of the wild-type in non-zero proportions in infinite populations. For finite populations of a practical size, stochastic effects are significant, and the counter-measures are less effective.

11 Conclusions

An evolutionary approach to robot design, working from simple towards more complex cognitive architectures, implies species evolution within the SAGA conceptual framework. This requires a very different analysis from standard GAs, and abandoning the goal-seeking metaphor associated with them. A new metaphor of hill-crawling of a converged species has been introduced, and this needs an analysis of the conflicting forces of exploitation and exploration — which here means efficiently searching down the current hillside and along high-value ridges in the fitness landscape while being careful not to lose track of the current hilltop.

Whereas theoretical biologists are trying to analyse the selection mechanisms they believe exist in the natural world, in simulated evolution we can choose our own selection mechanism. Arguments have been presented that tournament selection can be used for hill-crawling,

with significantly higher mutation rates than are used in conventional GAs; higher mutation rates enable a faster rate of evolution. It has also been argued that with the complex simulations that would be needed in evolutionary robotics, requiring serious computing power for each individual being evaluated, tournament selection allows a practical evolutionary setup to be highly distributed over an asynchronous network or networks of machines with minimal intercommunication. In addition, tournaments reduce the selection process to a succession of binary decisions as to which of two individuals is the better, avoiding scaling problems with any evaluation function.

Analytical results have been shown for the effects of tournament selection in the case of infinite populations, with and without reliable tournament decisions. Stochastic effects of genetic drift in small populations alter these results. Results from simulations with finite populations under different conditions have been shown, and practical ways to overcome Muller's ratchet have been suggested.

References

- [Barhen *et al.* 1987] J. Barhen, W.B. Dress, and C.C. Jorgensen. Applications of concurrent neuromorphic algorithms for autonomous robots. In R. Eckmiller and C.v.d. Malsburg, editors, *Neural Computers*, pages 321–333. Springer-Verlag, 1989.
- [Beer and Gallagher 1991] R.D. Beer and J.C. Gallagher. Evolving dynamic neural networks for adaptive behavior. Technical Report CES-91-17, Case Western Reserve University, Cleveland, Ohio, 1991.
- [Beer 1992] R.D. Beer. A dynamical systems perspective on autonomous agents. Technical Report CES-92-11, Case Western Reserve University, Cleveland, Ohio, 1992.
- [Brooks 1986] R. A. Brooks. Achieving artificial intelligence through building robots. A.I. Memo 899, M.I.T. A.I. Lab, May 1986.
- [Brooks 1991] R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
- [Brooks 1992] Rodney A. Brooks. Artificial life and real robots. In *Proceedings of the First European Conference on Artificial Life*. MIT Press/Bradford Books, Cambridge, MA, 1992.
- [Cariani 1989] P. Cariani. *On the Design of Devices with Emergent Semantic Functions*. PhD thesis, State University of New York at Binghamton, 1989.
- [Cliff *et al.* 1993a] D. T. Cliff, I. Harvey, and P. Husbands. Incremental evolution of neural network architectures for adaptive behaviour. In M. Verseylen,

- editor, *Proceedings of the First European Symposium on Artificial Neural Networks, ESANN'93*, pages 39–44. D facto 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 Bradford Books, Cambridge, MA, 1993.
- [de Garis 1992] Hugo de Garis. The genetic programming of steerable behaviors in GenNets. In F. J. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 272–281. MIT Press/Bradford Books, Cambridge, MA, 1992.
- [De Jong 1975] K. De Jong. *Analysis of the Behavior of a Class of Genetic Adaptive Systems*. PhD thesis, Univ. of Michigan, 1975.
- [Eigen and Schuster 1979] M. Eigen and P. Schuster. *The Hypercycle: A Principle of Natural Self-Organization*. Springer-Verlag, 1979.
- [Eigen *et al.* 1988] M. Eigen, J. McCaskill, and P. Schuster. Molecular quasi-species. *Journal of Physical Chemistry*, 92:6881–6891, 1988.
- [Goldberg and Deb 1990] D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. Technical Report TCGA-90007, TCGA, The University of Alabama, 1990.
- [Goldberg 1989] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, Massachusetts, USA, 1989.
- [Harvey *et al.* 1993] I. Harvey, P. Husbands, and D. T. Cliff. Issues in evolutionary robotics. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour (SAB92)*, pages 364–373. MIT Press Bradford Books, Cambridge, MA, 1993.
- [Harvey 1992a] Inman Harvey. The SAGA cross: the mechanics of crossover for variable-length genetic algorithms. In R. Männer and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 269–278. North-Holland, 1992.
- [Harvey 1992b] Inman Harvey. Species adaptation genetic algorithms: The basis for a continuing SAGA. In F. J. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 346–354. MIT Press/Bradford Books, Cambridge, MA, 1992.
- [Hesser and Männer 1991] J. Hesser and R. Männer. Towards an optimal mutation probability for genetic algorithms. In H.-P. Schwefel and R. Männer, editors, *Parallel Problem Solving from Nature*. Springer-Verlag, Lecture Notes in Computer Science Vol. 496, 1991.
- [Holland 1975] John Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, USA, 1975.
- [Husbands and Harvey 1992] 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 139–146. IEEE Press, 1992.
- [Kauffman 1989] Stuart Kauffman. Adaptation on rugged fitness landscapes. In Daniel L. Stein, editor, *Lectures in the Sciences of Complexity*, pages 527–618. Addison Wesley: Santa Fe Institute Studies in the Sciences of Complexity, 1989.
- [Koza 1990] John R. Koza. Genetic programming: A paradigm for genetically breeding populations of computer programs to solve problems. Technical Report STAN-CS-90-1314, Department of Computer Science, Stanford University, 1990.
- [Langton *et al.* 1991] C. G. Langton, J. D. Farmer, S. Rasmussen, and C. Taylor, editors. *Artificial Life II: Proceedings Volume of Santa Fe Conference Feb. 1990*. Addison Wesley: volume XI in the series of the Santa Fe Institute Studies in the Sciences of Complexity, 1991.
- [Langton 1989] C. G. Langton, editor. *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, NM, Sept, 1987*. Addison-Wesley: volume VI in the series of the Santa Fe Institute Studies in the Sciences of Complexity, 1989.
- [Maynard Smith 1978] John Maynard Smith. *The Evolution of Sex*. Cambridge University Press, 1978.
- [Meyer *et al.* 1993] J.-A. Meyer, H. Roitblat, and S. Wilson, editors. *Proceedings of the Second International Conference on Simulation of Adaptive Behaviour (SAB92)*, Cambridge, MA, 1993. MIT Press Bradford Books.
- [Nowak and Schuster 1989] Martin Nowak and Peter 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.

[PRANCE 1991] PRANCE. Perceptive robots: Autonomous navigation and cooperation through evolution. Unpublished research proposal, PRANCE consortium: Univ. of Sussex, Cap Gemini Innovation, École Normale Supérieure (Paris) and Université Libre de Bruxelles, 1991.

[Schaffer *et al.* 1989] J.D. Schaffer, R.A. Caruana, L.J. Eshelman, and R. Das. A study of control parameters affecting online performance of genetic algorithms for function optimization. In J. D. Schaffer, editor, *Proceedings of the 3rd ICGA*, San Mateo, California, 1989. Morgan Kaufmann.

[Smith 1980] Stephen F. Smith. *A Learning System based on Genetic Adaptive Algorithms*. PhD thesis, Department of Computer Science, University of Pittsburgh, USA, 1980.

[Varela *et al.* 1991] F. Varela, E. Thompson, and E. Rosch. *The Embodied Mind*. MIT Press, 1991.

[Viola 1988] P. Viola. Mobile robot evolution. Bachelors thesis, M.I.T., 1988.

[Whitley 1989] D. Whitley. The genitor algorithm and selection pressure: Why rank-based allocation of reproductive trials is best. In J. D. Schaffer, editor, *Proceedings of the 3rd ICGA*, San Mateo, California, 1989. Morgan Kaufmann.

Appendices

A Rank sizes with long genotypes

Consider a population of size N , with binary genotypes, and classify each individual by Hamming-distance from the current master sequence or wild-type, which without loss of generality can be taken to be a genotype of all 0s. It is assumed that fitness monotonically decreases with this distance, and without loss of generality the negative fitness can be the number of 1s in the genotype. At each replication there is a single mutation, and the genotypes are so long that the chance of a back-mutation is negligible; i.e. all mutations are deleterious.

The number of individuals with i 1s is defined as r_i . The loss from and gain to this class should balance at equilibrium.

Something of rank i is lost in a tournament if there is a tournament between a rank i and a rank $< i$, with probability

$$2 \frac{r_i}{N} \sum_{j=0}^{i-1} \frac{r_j}{(N-1)},$$

and also for one between two different members of rank i , probability

$$\frac{r_i (r_i - 1)}{N(N-1)}.$$

Something of rank i is gained in a tournament between rank $(i-1)$ and a rank $> i$, with probability

$$2 \frac{r_{i-1}}{N} \sum_{j=i+1}^N \frac{r_j}{(N-1)},$$

and also for one between two different members of rank $(i-1)$, probability

$$\frac{r_{i-1} (r_{i-1} - 1)}{N(N-1)}.$$

Setting gains equal to losses, and eliminating $N(N-1)$ we have

$$2r_i \sum_{j=0}^{i-1} r_j + r_i(r_i - 1) = 2r_{i-1} \sum_{j=i+1}^N r_j + r_{i-1}(r_{i-1} - 1).$$

This gives a value for r_i based on values for lower i s. Since we know that $r_0 = 1$, we have

$$2r_1 + r_1(r_1 - 1) = 2(N - 1 - r_1).$$

This quadratic equation yields a positive solution of

$$r_1 = \frac{1}{2}(-3 + \sqrt{8N + 1}).$$

For $N = 100$ this gives $r_1 \approx 12.65$. Successive values for higher i can be found by iteratively solving successive quadratics.

B Noisy decisions in an infinite population

Consider a similar situation to that in appendix A, but where there is a probability p of correctly deciding a tournament, and hence $(1-p)$ of making a mistake. We will be considering the long term possibility of losing all members of rank 0, and will assume a population of infinite size so as to ignore stochastic effects of genetic drift. Let the proportion of the population in rank 0, the master sequence, be a . To keep up the value of a , we will impose a probability q of mutating on replication, and hence $(1-q)$ of there being no mutation.

A wild-type of rank 0 will be lost from the population when rank 0 meets rank 0 and there is a mutation, probability a^2q ; and also when rank 0 meets rank ≥ 1 and the wrong one wins, probability $2a(1-a)(1-p)$.

A rank 0 will be gained when a rank 0 meets a rank ≥ 1 , the right one wins, and there is *no* mutation, probability $2a(1-a)p(1-q)$.

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 s is positive and $(2 - q)$ is positive, so the condition for a to be positive is that the r.h.s. of the above equation is also.

Hence $2(2p - pq - 1) > 0$, which gives $q < (2p - 1)/p$ as the condition for the proportion a of rank 0 to remain positive.

C Mutations assessed independently at each locus

Consider an population of N binary genotypes each of length l . N is assumed to be large enough to avoid stochastic genetic drift, and l is assumed large enough for approximations to be made below.

If there is on average μ mutations per genotype, this is a probability of μ/l at each locus. So the chance of there being no mutation at all l loci on a genotype is $(1 - \mu/l)^l$. For small μ and large l this is close to $e^{-\mu}$.

Let the proportion of rank 0 in the population be a . Then a rank 0 will be lost to the population when rank 0 meets rank 0 and there *is* a mutation, giving a probability of

$$a^2 (1 - e^{-\mu}),$$

a rank 0 will be gained when a rank 0 meets a rank ≥ 1 and there is *no* mutation, a probability of

$$2a(1-a)e^{-\mu}.$$

Setting gains equal to losses, and multiplying through, we have

$$a(e^\mu - 1) = 2(1-a).$$

Hence the proportion of the population expected to be the wild-type at equilibrium is:

$$a = \frac{2}{1 + e^\mu}.$$

Let us now reconsider this scenario with noise added, when the winner of a tournament is selected with probability p . A rank 0 will still be lost to the population when rank 0 meets rank 0 and there *is* a mutation, giving a probability of

$$a^2 (1 - e^{-\mu}),$$

but a rank 0 will *also* be lost when rank 0 meets rank ≥ 1 , and noise makes it lose; probability is

$$2(1-p)a(1-a).$$

A rank 0 will be gained when a rank 0 meets a rank ≥ 1 , wins, and there is *no* mutation, a probability of

$$2pa(1-a)e^{-\mu}.$$

Setting gains equal to losses, and multiplying through, we have

$$a(e^\mu - 1 + 2e^\mu(1-p)(1-a)) = 2p(1-a).$$

Hence

$$a(e^\mu(2p-1) - 1) = 2p - 2(1-p)e^\mu.$$

The contents of the bracket on the l.h.s. are always positive, so that the condition for $a > 0$ is that the r.h.s. is positive,

$$p > (1-p)e^\mu.$$

Hence for preservation of the wild-type we need

$$p > \frac{e^\mu}{1 + e^\mu}.$$

The proportion of the population expected to be the wild-type at equilibrium is:

$$a = \frac{2}{1 + e^\mu}.$$

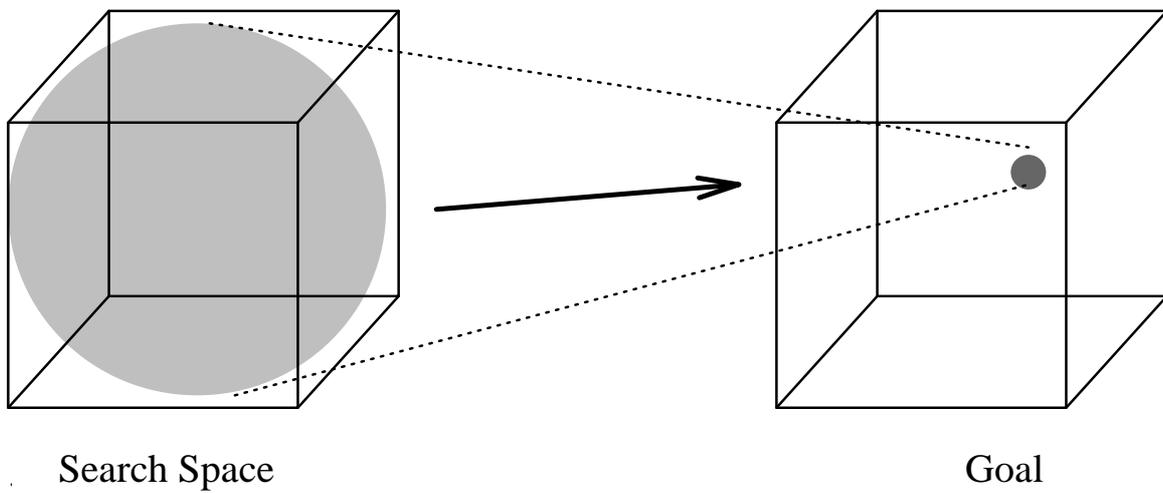


Figure 1: *The evolution of a standard GA in a fixed-dimensional search space; the population initially spans the whole space, and in the end focuses on the optimum.*

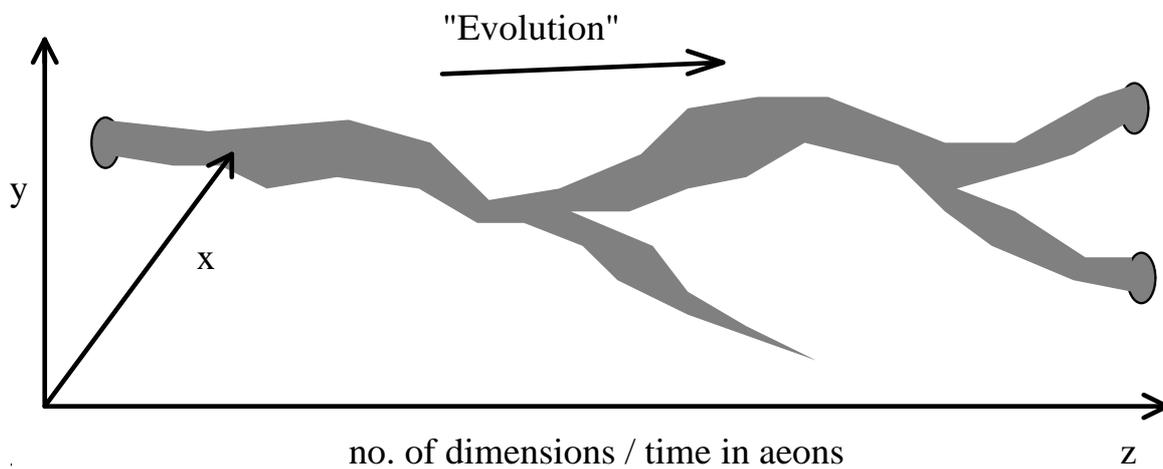


Figure 2: *The progress of the always compact course of a species; the horizontal axis indicates both time and the (loosely correlated) number of dimensions of the current search space. The other axes represent the current dimensions in genotype space. The possibility of splitting into separate species, and of extinction, are indicated in the sketch, although not here discussed.*

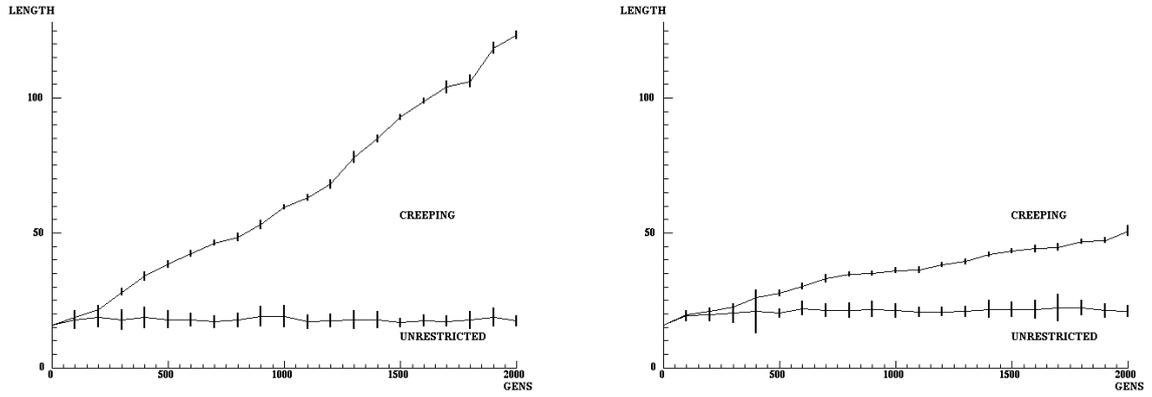


Figure 3: Average genotype lengths against generations; vertical bars show standard deviations. Effects of 'creeping' and 'unrestricted' increase-length genetic operators on a population with the same fitness conditions, epistasis $K = 2$. Left graph, linkage with neighbouring genes. Right graph, random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

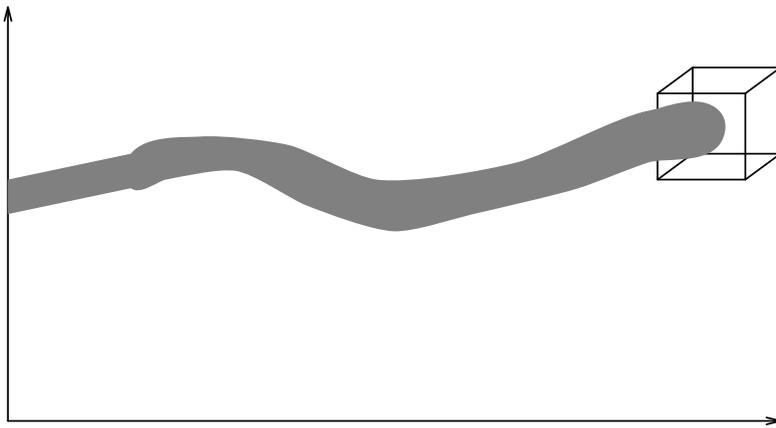


Figure 4: As a species evolves through SAGA space, the search for higher fitness only takes place in a very local search space around the current focus of the species.

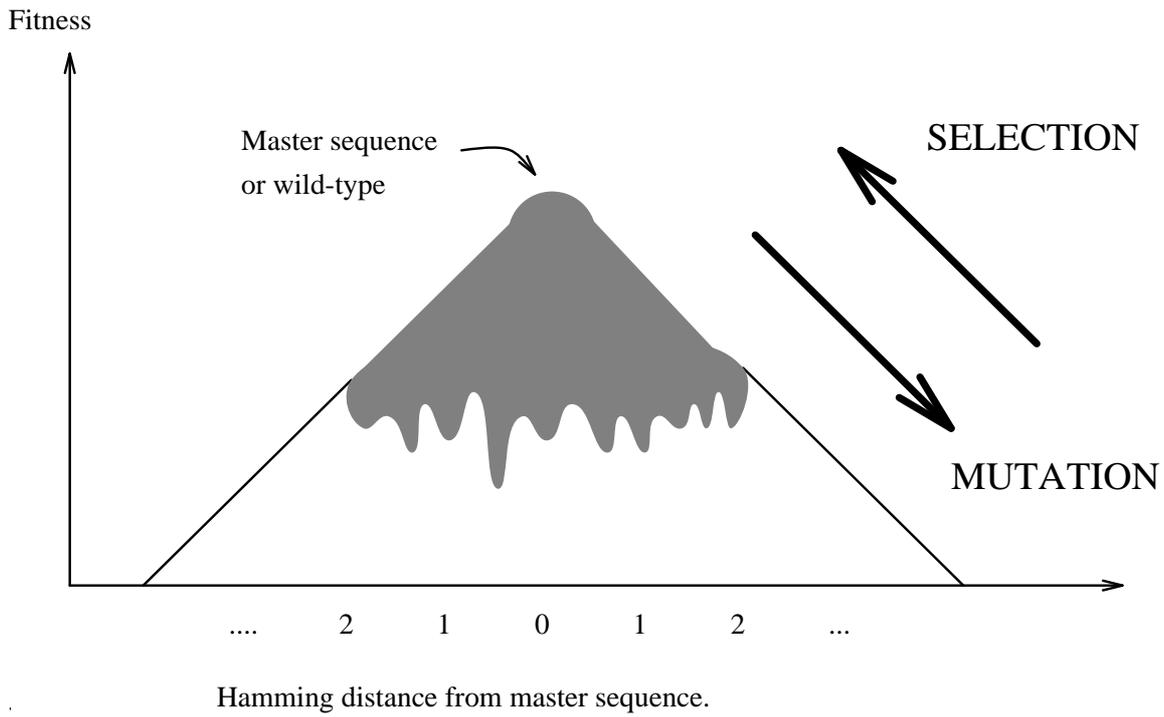


Figure 5: *The opposing forces of mutation and selection on a population centred around a local optimum, where Hamming distance from master sequence is directly related to fitness ranking.*

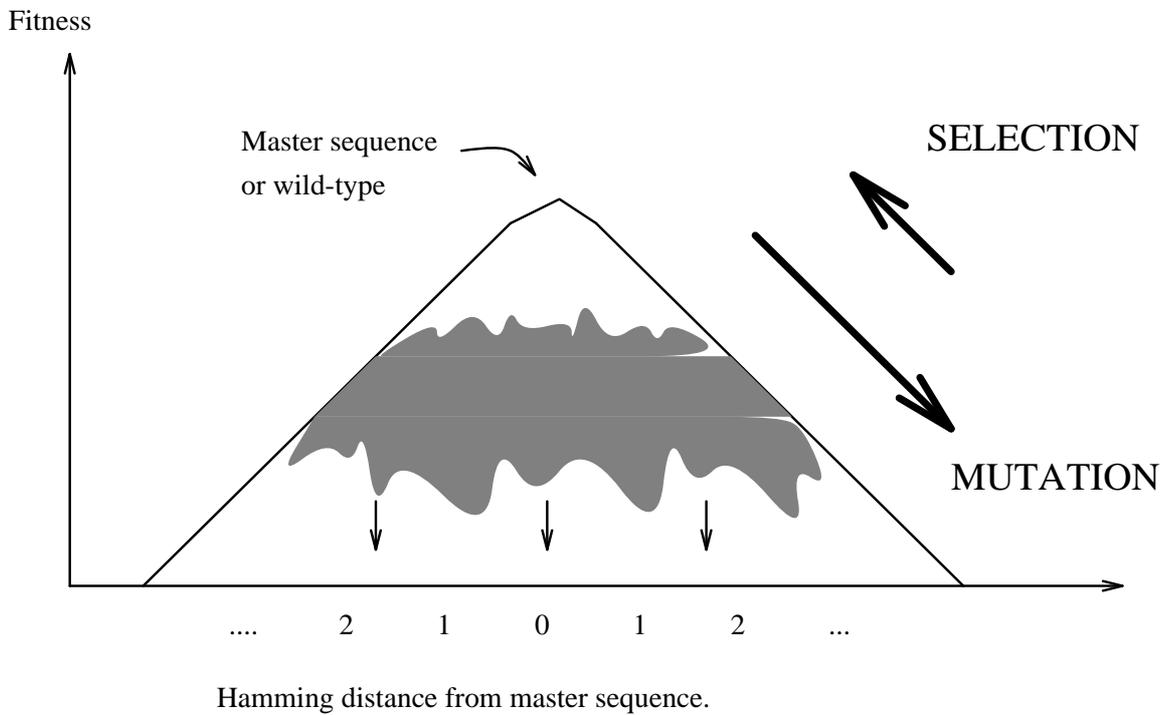


Figure 6: *When mutation outweighs selection so that the fittest rank can be lost, Muller's ratchet inexorably drives the population down the hill.*

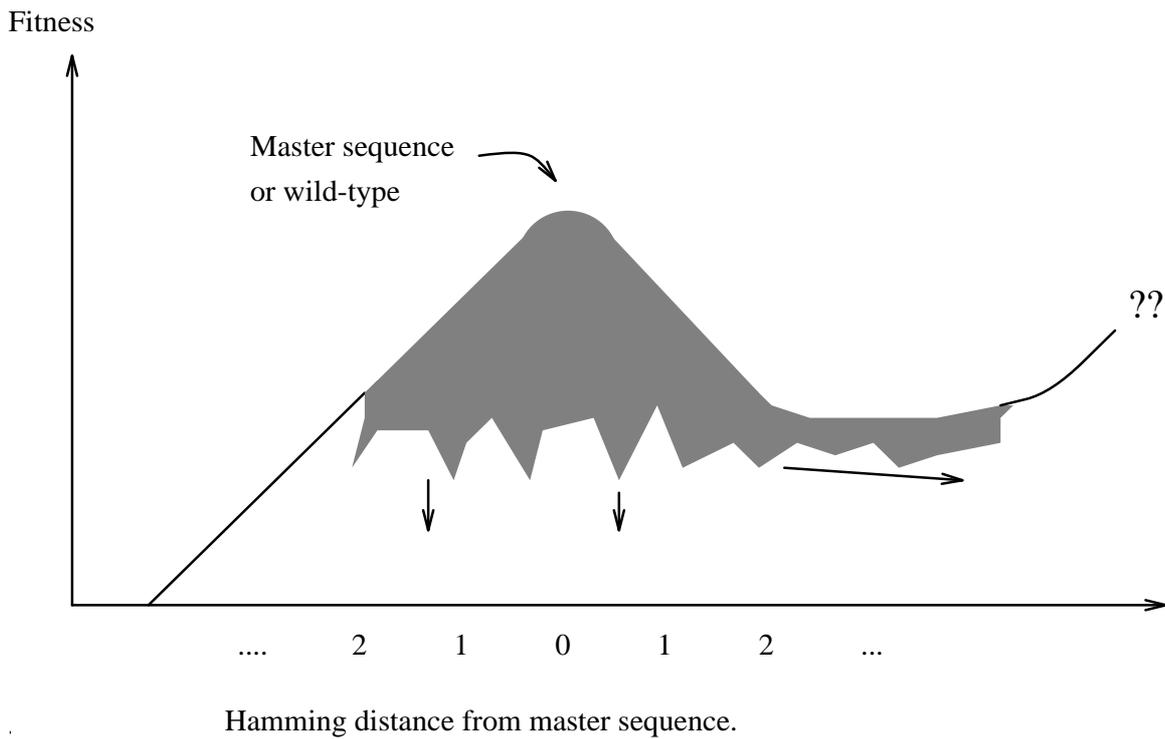


Figure 7: *If the population can crawl down the hill far enough to reach a ridge of relatively high fitness, it will spread along it, potentially reaching new hills.*

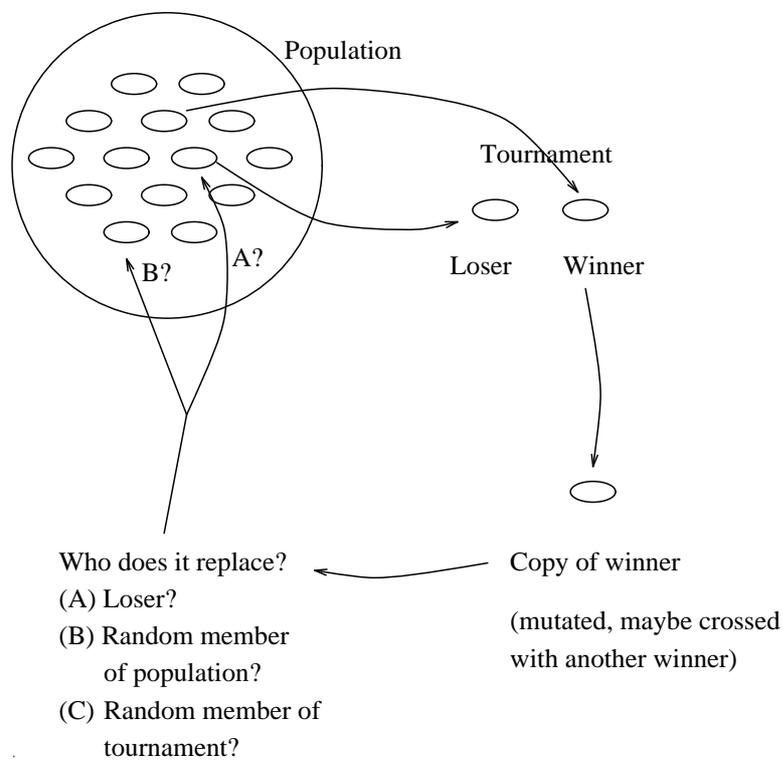


Figure 8: *Tournament Selection. Possible choices for who is to die to make way for the copy of the winner.*



Scenario 1: Exactly m bits flipped, chosen at random. Genotype is so long that chance of favourable back-mutation towards master-sequence is negligible.



Scenario 2: Exactly m bits flipped, chosen at random. Finite length genotype with significant possibility of back-mutation.



Scenario 3: Each bit is flipped with probability of $m/(\text{genotype length})$. I.e. on average m mutations on whole genotype.

Figure 9: Possible ways to apply a given mutation rate.

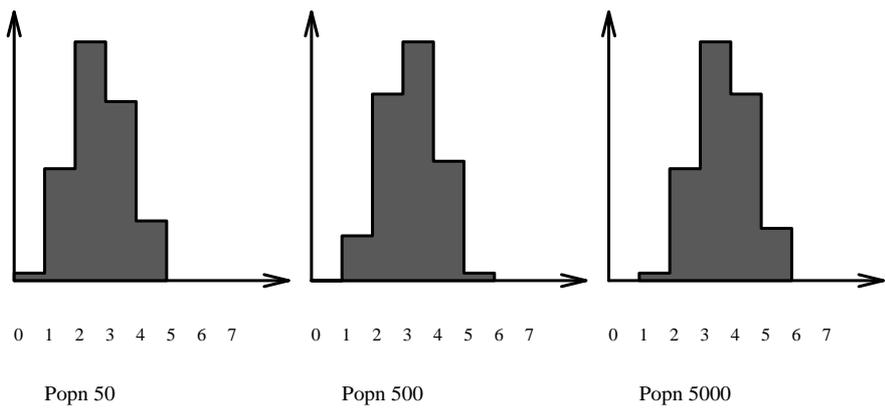


Figure 10: Numbers in each class, ranked by the Hamming-distance from the master-sequence. Effectively infinitely long genotype, in that there are no back-mutations, but exactly one mutation in genotype at replication of winner of tournaments. Results shown are after simulations of number of tournaments equal to 4000 times population size.

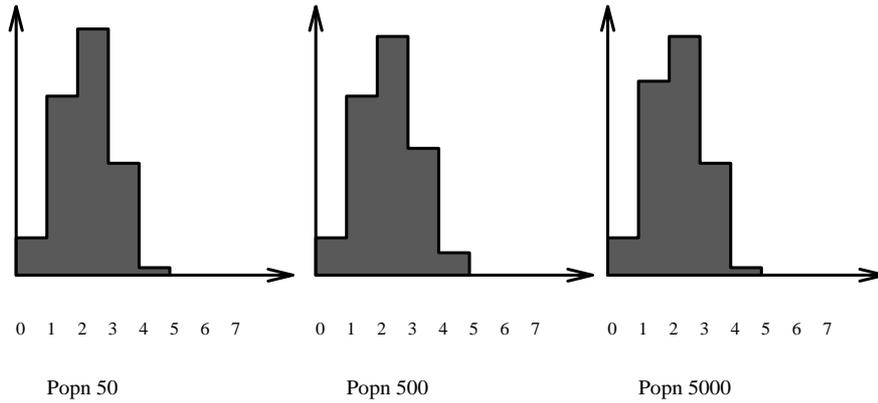


Figure 11: Similar simulations to those shown in previous figure, except that genotypes are of length 100. Exactly one mutation per genotype at replication, which means a significant number of back-mutations towards the master-sequence.

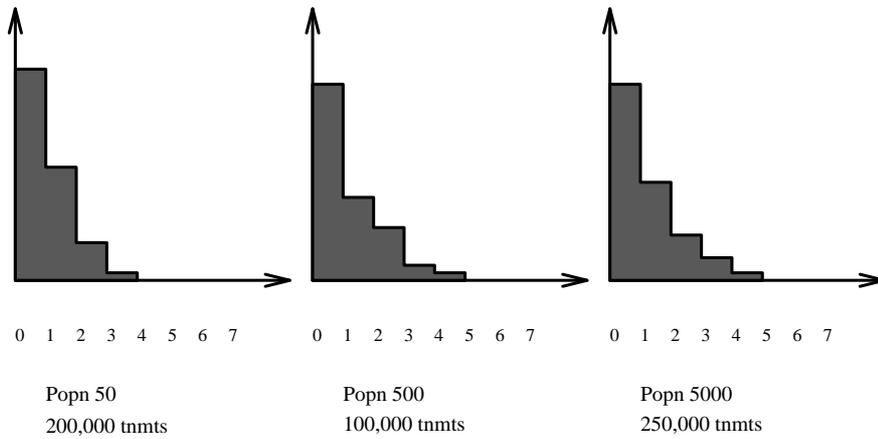


Figure 12: Similar conditions to previous figure, except that there is a 1/100 chance of a mutation at each bit of the genotype, which is of length 100.

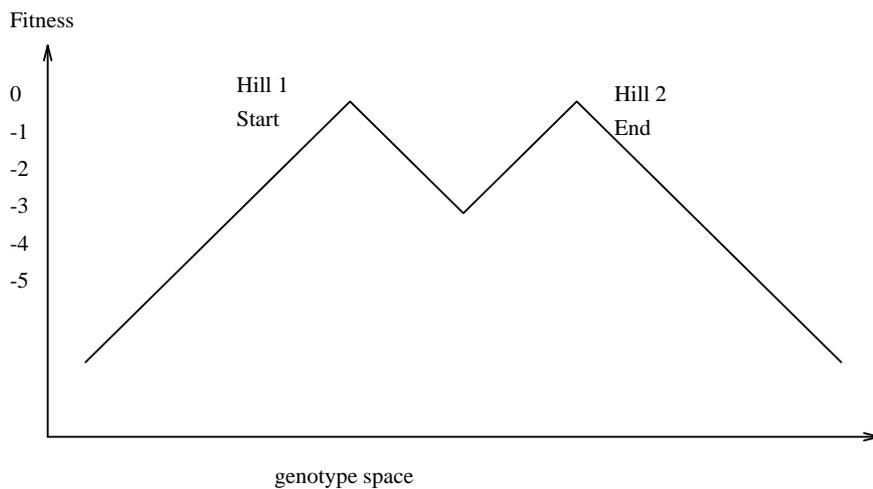


Figure 13: Two genotypes are designated as start and goal, each of them with maximum fitness (of 0). The fitness of any other genotype is defined as (the negative of) the minimum of the Hamming distances to start and to goal. This gives a fitness landscape of two hills as shown. The population is all started off at the peak of the 'start hill'.

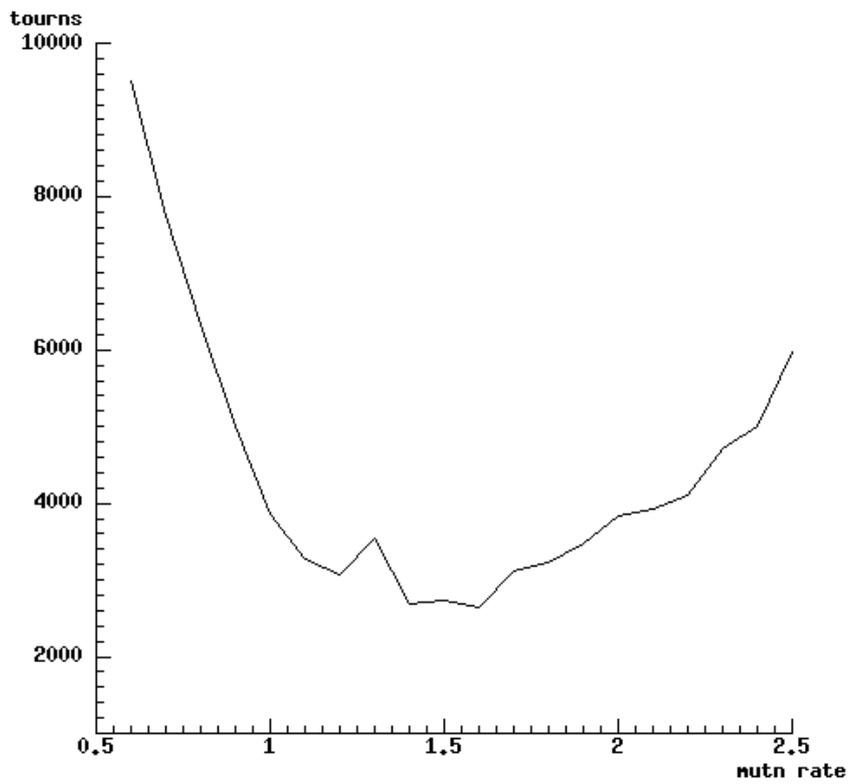


Figure 14: *Number of tournaments for a population centred at one 'hill-top' to have a first member reach a nearby hill-top. Rate specified is the average number of mutations per genotype. Recombination has been used.*