

Genetic Convergence in a Species
of
Evolved Robot Control Architectures

I. Harvey, P. Husbands, D. Cliff

CSRP 267, January 1993

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

Genetic Convergence in a Species of Evolved Robot Control Architectures

Inman Harvey¹ and Philip Husbands¹ and Dave Cliff^{1,2}

¹School of Cognitive and Computing Sciences

²Neuroscience IRC, School of Biological Sciences

University of Sussex

Brighton BN1 9QH, England

email: inmanh or philh or davec@cogs.susx.ac.uk

Abstract

We analyse how the project of evolving ‘neural’ network controllers for autonomous visually guided robots is significantly different from the usual function optimisation problems standard genetic algorithms are asked to tackle. The need to have open-ended increase in complexity of the controllers, to allow for an indefinite number of new tasks to be incrementally added to the robot’s capabilities in the long term, means that genotypes of arbitrary length need to be allowed.

This results in populations being genetically converged as new tasks are added, and needs a change to usual genetic algorithm practices. Results of successful runs are shown, and the population is analysed in terms of genetic convergence and movement in time across sequence space.

1 Introduction

In the context of our ongoing project to evolve ‘neural’ networks which act as controllers for visually guided autonomous robots, some basic questions as to the nature of evolution have to be faced. Can such a project be treated as a function optimisation problem, for which standard genetic algorithms (GAs) have been designed? If, as will be argued below, there is a significant difference between evolution and optimisation, then what changes in GAs are needed to deal with this?

In answering these questions, particular attention will be paid in this paper to issues of genetic convergence, which in standard GAs is usually taken to signal the end of the road. Brief details of the networks which act as controllers, and of the robots themselves which are required to perform simple navigational tasks using vision and touch-sensors, will be given here; for fuller information see the papers cited below. A particular run which results in successful behaviour will be analysed in terms of the movement of the population across sequence space.

2 Evolution versus Optimisation

Genetic Algorithms have been focused in such a concentrated fashion on function optimisation problems that, for instance, when De Jong presented a paper [4] bringing to peoples' attention the fact that Holland's motivation for his initial GA work was the design and implementation of robust adaptive systems, a much broader context, this was greeted with some surprise and scepticism by the audience. De Jong stressed that it was a fallacy to *equate* GAs with function optimisation.

The generic GA was not designed to solve any particular problem, but was rather a high level simulation of a biological adaptive system, Darwinian evolution. One (of many) ways to think of evolution is as a strategy for exploration and traversal of complex, time-varying, fitness landscapes. In natural evolution, making the rather large assumption that 'fitness' can be unproblematically defined, there remains the question of what are to be the 'horizontal' dimensions of such a landscape. They may be phenotypic characteristics, or genotypic ones; whichever they are, they can only be treated as well-specified and meaningful dimensions in the short term, for just so long as changes in the population treated as moving across this landscape are not too radical. There is no sensible single fitness landscape which can simultaneously cater for apes, jellyfish, bacteria and self-reproducing RNA molecules at the very origin of life, although individual landscapes may be usefully posited for each one of these.

Function optimisation can, of course, use fitness-landscape language. But the landscape is always fully specified by the specific function being optimised, and the problem is usually to find the global optimum, or some near-optima, of the whole landscape. In contrast, a fitness landscape in evolution can only be specified with reference to some current population, and their genetic or phenotypic characteristics, and such a population will inevitably be already situated in some confined region of this landscape. Hence, in so far as any question is being asked in evolution, it is not 'what is best?', but rather 'where shall we go to from here?'

In the context of robotics, it has been suggested [10] that the design by hand of the control systems for autonomous robots is reaching the limits of feasibility; and that the only hope of future progress is through some evolutionary process. Although current practice in GAs is an obvious starting place for establishing some such evolutionary process, it is indeed evolution and not optimisation that is required. It is proposed that a (or several) 'species' of robots — or robot control architectures — should be evolved in an incremental fashion. As each new task is added to the specification, the starting place should be the current converged population, rather than a fresh initial random spread.

3 SAGA principles

Initial tasks for an autonomous robot, whose architecture is genetically determined, may be rigorously specified; but it is not possible to specify in advance what future tasks may be indefinitely added. Since there must be some at least loose correlation between the complexity of such an architecture, and the length of a genotype which determines it, then an evolutionary algorithm must be able to deal with genotypes of arbitrary lengths. This lead to the development of SAGA (Species Adaptation Genetic Algorithms) principles in [8],




Figure 1: *Mutation allows a population to explore along ridges towards potentially higher hills in the fitness landscape. This picture is potentially misleading in high-dimensional landscapes ...*




Figure 2: *...as there is not just a single shortest distance between two points Hamming distance d apart in binary genotype sequence space, there are $d!$ shortest routes, and far more than this that are nearly as short.*

briefly introduced and summarised below.

The immediate questions which raise themselves are:

- How, if at all, can Holland's Schema Theorem be accommodated?
- How should genetic operators which allow change in genotype length be handled?
- And how should recombination be done between genotypes of different lengths?

If a coding from genotype to phenotype is chosen which allows indefinite increase in the length of the former, associated with indefinite increase in complexity of the latter, then the notions of schemata needed for Holland's Schema Theorem do not work. Since the class of all genotypes with specified fixed values for particular alleles is now infinite in size — whereas normally with fixed-length genotypes it is of large but finite size — the concept of an average fitness for the members of this class becomes highly questionable. The following route to partially reconciling the Schema Theorem with arbitrary-length genotypes is rather devious, and starts with a detour.

As discussed more fully in [8], for all practical purposes evolution requires the fitness landscape to be 'not too rugged'; for an explanation of the ultimately dead-end nature of adaptation on fully rugged landscapes see [12]. For a landscape to be fairly smooth, this implies that points close together in 'horizontal' distance should in general be reasonably correlated in 'vertical' distance, i.e. fitness. If the 'horizontal' dimensions refer to genotype space, then immediate neighbours are those that can be reached in a single genetic operation such as mutation of a single bit. If one also includes as neighbours those that can be reached by application of a genetic operator that changes genotype length, then one will need to restrict those changes to ones that do not (very often) change fitness by an arbitrarily large amount; the landscape should not contain too many 'precipices'. As explained at the end of the next section, this virtually eliminates the chance of being trapped on some local optimum. If one can characterise the fitness of the whole as composed of the sum of fitness contributions determined by separate parts of the genotype ('genes', if you like), with a reasonable degree of epistatic interaction between these parts, this smoothness requirement translates into one that the genotype length should not change by an enormous amount in

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 over evolutionary time-scales must be genetically converged. Not only must any changes in length be gradual, but all the lengths within a population will be very similar, and will be genetically converged at corresponding loci — a species.

This has consequences for allowable genetic operators for recombination, primarily that such an operator must produce offspring with genotypes of similar (not necessarily identical) lengths to their parents; and homologous segments must be swapped [7].

Returning from this detour to see how the Schema Theorem is rescued, it turns out that in these particular circumstances of only gradual change, small finite bounds can in practice be put on genotype lengths in the short term, and hence the class of all members of any given schema is now finite in size.

4 Convergence and Mutation

The most visible difference between genotypes under evolution, and those in a standard GA for function optimisation, is that at all times (bar perhaps the very start) the population is virtually genetically converged. In a standard GA this is usually considered the end of the story. The received folklore is that recombination is the driving force for genetic search, and mutation is only a background operator. To quote from [13]:

“Clearly the $\mathcal{O}(n^3)$ estimate [for implicit parallelism] is based on a *diverse* population, where many schemata are represented. However, as exponential allocation of observed-best schemata accrues, one can expect that the number of building blocks processed will decrease. This is an inevitable consequence of convergence in the [vanilla-flavour] GA outlined above. After convergence, the GA population will be composed primarily of copies of one individual. The only diversity maintained in the population after convergence is a result of mutation. Note that mutation is a completely random operator that is unguided by the algorithm’s observations of fitness values over time.”

Now whereas it is undisputed that mutation within an individual is completely random, it does not follow that random mutation on the individuals within a converged population under selection results in random undirected movements of the population across the fitness landscape. Treating the converged population as being currently centred around some local hilltop, then mutations can be thought of as explorations away from the peak; with long genotypes the chance of a back-mutation is insignificant. But successive rounds of the

¹How big a change in length can be while remaining ‘slight’ depends on how rugged the fitness landscape is; it must be related to the *correlation length* of the space [11].

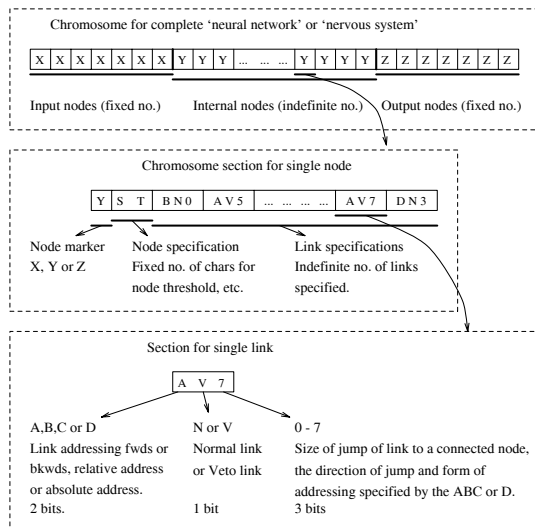


Figure 3: *The genetic encoding scheme*

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 3-D spaces. Whereas in a normal 3-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 centre of the room and circles there, using input from 2 photoreceptors. The direction the robot is facing is indicated by arrows for each time step, largely superimposed.*

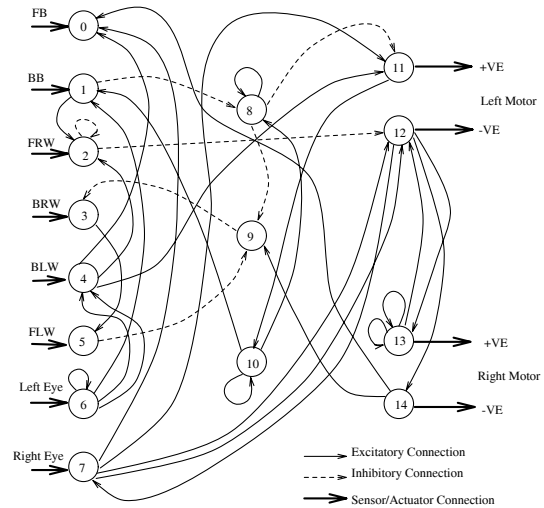


Figure 5: *The network diagram for the robot that produces the behaviour shown in the previous figure.*

5 Evolving Robots

Our aim is to evolve the control systems for autonomous mobile robots, initially in simulation but soon to be transferred to physical robots nearing completion. The robot, on which the simulation is based, is circular, with two independent drive wheels and a trailing rear castor. It has front and back bumpers, and four whiskers each at 45° to the midline of the robot. In the experiments considered here, it has also two photoreceptors, symmetrically placed each side of the direction the robot faces, with the angle of acceptance of each photoreceptor, and the angle of eccentricity away from the forward direction, genetically determined by a ‘vision chromosome’. The control system is a recurrent dynamical neural network, genetically specified by the main ‘chromosome’. Eight of the nodes in the network are designated as inputs, activated by the 8 sensors — bumpers, whiskers, photoreceptors. Four are designated as outputs, controlling the left and right motors, which each can go full-speed or half-speed, forwards or backwards, or stop. The genotype can also specify an arbitrary number of internal, or ‘hidden’ nodes in the network.

For reasons given in [6] we advocate continuous real-valued networks with unrestricted connectivity and time delays between units (i.e., nothing like back-propagation!). To date, all weights in the networks have been fixed, at a value of 1. The nodes themselves act as noisy linear threshold devices. Two types of links between nodes are allowed: normal and veto. The latter is an infinitely inhibitory connection subject to its own threshold; if this is exceeded, then all the normal output of the vetoed node is turned off.

The genetic encoding used is illustrated in Figure 3. The genotype is interpreted as a

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 [3]. 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.

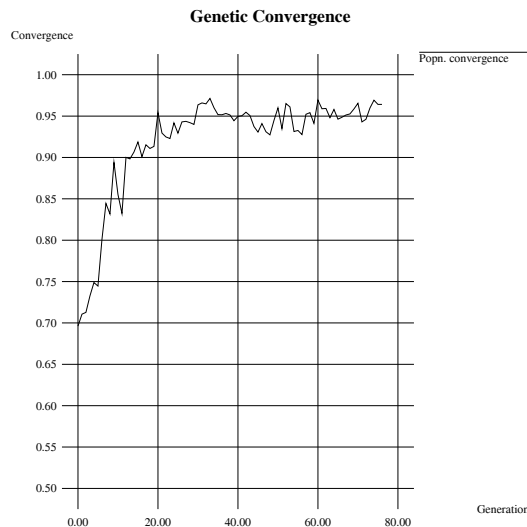


Figure 6: *Genetic convergence against generations, calculated as the percentage agreement between pairs of genotypes taken from the population. For details see text.*

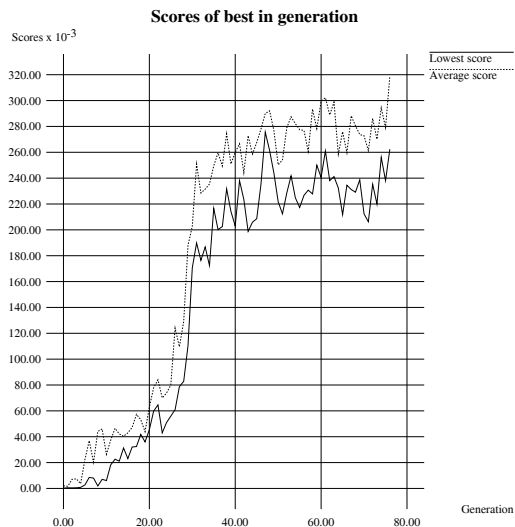


Figure 7: *The scores of the best member of each generation. Shown are the average score that the best member achieved over 8 (noisy) trials; and the lowest score of its 8 trials — it is this figure that is used as the evaluation.*

6 Genetic Analysis

A particular run of some 77 generations with a population of 60 is analysed here. The starting population was randomly initialised, and only the single task was used for evaluation; so there has not yet been any attempt to increase the complexity of the task over time. Selection was rank-based, with a quadratic used to convert ranking into expected contribution to the next generation. The first i members of a population size n have between them a quota proportional to \sqrt{i} . When $n = 60$, this means that the the first 16 all have above average quotas, and the very first contributes an expected $\sqrt{60} \simeq 7.75$ to the pool for the next generation. This is, compared to standard GAs, abnormally high selection which the rank-based method maintains indefinitely. The mutation rate was set at an expected 0.9 bits flipped per genotype. There was 100% recombination, with the single crossover point arranged so that, despite the possible variations in lengths between recombining genotypes, there was minimal change in length in the offspring.

Measuring genetic convergence in a population with varying lengths is non-trivial, even though the lengths in general remain nearly equal. A number (here 8) of pairs of genotypes were selected at random, and for each pair the longest common subsequence (LCSS) was calculated [9]. Convergence for this pair was taken to be the length of the LCSS divided by the average length of the pair; population convergence was taken to be the average convergence of those pairs sampled.

The population converged to around 95% after only some 20 generations (Fig. 6), driven by the strong selection, even though absolute scores were low. The sharp rise in fitness around 30 generations, followed by a prolonged gradual improvement, occurred *after* this




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 be seen to be tracking the average scores closely, indicating the robustness of the networks. Other studies have shown that when the noise is increased to levels higher than those used during evolution, degradation in performance is ‘graceful’ [1].

The run was stopped arbitrarily after 77 generations. At this stage, all the genotypes were the same length except for one minimally shorter and one minimally longer. These two were ‘edited’ to conform in length, and the whole population converted into binary format, which was then 330 bits long. In 224 of these places there was 100% convergence. These identical values were discounted, the consensus sequence calculated (a sequence with the most popular value for each position), and the population of genotypes displayed in terms of their difference from the consensus sequence (Fig. 8). In this figure they are listed in order of Hamming distance from the consensus sequence, these distances ranging from 0 to 19. In fact the consensus sequence itself was present in this particular population, though in the general case this need not be so. The first 21 so listed are at a maximum Hamming distance of 2 from the consensus, indicating a tight cluster.

Some correlation is distinguishable by eye amongst those further from the consensus. A principal components analysis (PCA), of the first and second components, is shown in Fig. 9, confirming the presence of a strong correlation. The central cluster is shown in more detail in Fig. 10, and includes those nearest the consensus. The PCA is here used as a convenient tool to give a 2-dimensional snapshot of what here is a 106-dimensional space. For the first component, a vector is chosen though the consensus sequence in the direction such as to maximise the variation in the projections of all the points onto this vector; the




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} .




Figure 11: This time the same final population is numbered in order of scores, 0 being the best. This best is contained within the ‘ridge’ towards the north-north-east, far from the central cluster. Another high-scorer, 4, is isolated towards the east.




Figure 12: Focusing on the north-eastern ‘ridge’ reveals the best scorer 0, amongst some relatively indifferent ones.

second component is an orthogonal vector chosen so as to maximise the remaining variation.




Figure 13: 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 a ‘ridge’ to the north-north-east of the centre of this graph.

By numbering the individuals within the population according to their ranking, it can be seen that the current winner is in fact within this ‘ridge’, and hence will be contributing to a larger cluster around there when it is preferentially copied in large numbers in the next generation.

Looking at the history from once the population had converged up until the 76th generation, the Hamming distances of the best in every 5th generation, from 36 to 76, from the consensus sequence of this last generation, is shown in Fig. 14. Again, a PCA of these best-of-generations, Fig. 13, shows a continuous trajectory. By looking at the distances moved by the best-of-generation every 5 generations (Fig. 14), and comparing with distances to the consensus of the last generation, it can be seen that relatively large distances can be traversed across the sequence space, despite the high degree of convergence. It should be borne in mind that at this stage, with genotypes of around 330 bits, no two points in the sequence space are further apart than 330. Thus it can be seen that the early genetic convergence is no barrier to movement across sequence space, and possible continued improvement.

7 Conclusions

A distinction has been drawn between evolution, which does not have its ultimate goals fully specified from the start, and function optimisation which necessarily does. For the very practical job of producing controllers for autonomous robots, it is the former that we must use. The SAGA principles presented cover some of the necessary changes in standard GA

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, 1993.
- [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.
- [3] 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, 1993.
- [4] Kenneth De Jong. Are genetic algorithms optimizers? In R. Männer and B. Manderick, editors, *Parallel Problem Solving from Nature, 2*, pages 3–13. North-Holland, 1992.
- [5] M. Eigen, J. McCaskill, and P. Schuster. Molecular quasi-species. *Journal of Physical Chemistry*, 92:6881–6891, 1988.
- [6] I. Harvey, P. Husbands, and D. T. Cliff. Issues in evolutionary robotics. 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.
- [7] 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.
- [8] 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.

- [9] D.S. Hirschberg. A linear space algorithm for computing maximal common subsequences. *Communications of the A.C.M.*, 18(6):341–343, 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 139–146. IEEE Press, 1992.
- [11] 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.
- [12] Stuart Kauffman and Simon Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128:11–45, 1987.
- [13] R.E. Smith and D. E. Goldberg. Diploidy and dominance in artificial genetic search. *Complex Systems*, 6:251–285, 1992.