

# Analysis of Evolved Sensory-Motor Controllers

Dave Cliff<sup>1,2</sup> and Philip Husbands<sup>1</sup> and Inman Harvey<sup>1</sup>

<sup>1</sup>School of Cognitive and Computing Sciences

<sup>2</sup>Neuroscience IRC, School of Biological Sciences

University of Sussex, BRIGHTON BN1 9QH, U.K.

davec or philh or inmanh, all @cogs.susx.ac.uk

## Abstract

We present results from the concurrent evolution of visual sensing morphologies and sensory-motor controller-networks for visually guided robots. In this paper we analyse two (of many) networks which result from using incremental evolution with variable-length genotypes. The two networks come from separate populations, evolved using a common fitness function. The observable behaviours of the two robots are very similar, and close to the optimal behaviour. However, the underlying sensing morphologies and sensory-motor controllers are strikingly different. This is a case of convergent evolution at the behavioural level, coupled with divergent evolution at the morphological level.

The mathematics of our particular style of network are such that it would be difficult or impossible to derive closed-form equations describing the action of the networks.<sup>1</sup> Instead, our analysis uses techniques analogous to those used in the study of biological sensory-motor neural systems. However, the qualitative analyses presented in this paper can be augmented with quantitative analyses, as explained in (Husbands *et al.* 1993). In trying to understand how our artificially evolved networks generate behaviours in the robot, we are performing a task directly analogous to the task faced by biological scientists in the field of *neuroethology* — the study of the neural mechanisms underlying the generation of a creature's behaviour; see e.g. (Camhi 1984). For further details of the link between neuroethology and artificial neural network research, see (Cliff 1991, Beer 1990).

---

<sup>1</sup>For example, the transfer functions used in our model neurons are all nonlinear with discontinuities in the first derivative, and non-Gaussian noise is introduced at a number of points in the sensory-motor system.

# 1 Introduction

As part of our ongoing work in using genetic algorithms to develop 'neural' networks which act as controllers for visually guided robots, we have analysed the final evolved networks in order to identify how they work. This is an essential step in moving away from the treatment of artificially evolved neural networks as magical black boxes.

We view the networks we evolve as continuous dynamical systems, rather than as computational devices transforming between representations: inputs to the system might *perturb* the trajectory of the network in state space, so it enters a different state which might be interpreted by an external observer as a new behaviour. We find this perspective less encumbering than the traditional computational perspective, and also less amenable to the use of potentially misleading intentional language; see e.g. (Beer 1992, van Gelder 1992, Smithers 1992) for further discussion of the benefits of adopting a dynamical systems perspective.

Most of this paper deals with analysing two networks from separate populations, each evolved to perform the same task. We demonstrate that although the final observed behaviour from the two networks is very similar, the underlying mechanisms are remarkably distinct: the two populations converged at the behavioural level, while maintaining distinct sensory-motor morphologies.

The primary focus of this paper is on analysing networks resulting from the evolutionary processes; other papers are cited for further details of the genetic encoding, the genetic algorithm employed, and description of the vision system. Nevertheless, Section 2 offers a brief overview of most of the important details. Following that, Section 3 describes our experimental regime, and provides analysis the two networks. Finally, Section 3.3 discusses the implications of our work.

## 2 Background

### 2.1 Rationale

The rationale for our work, and some early results, have been discussed elsewhere (Husbands and Harvey 1992, Harvey *et al.* 1993, Cliff *et al.* 1993, Cliff *et al.* 1992a). The notes below present a brief summary of the important concepts.

In common with a growing number of other researchers, we believe that the generation of *adaptive behaviour* should form the primary focus for research into cognitive systems. By this, we mean behaviour which is selected to increase the chances that a situated agent can survive in an environment which is noisy, dynamic, and uncertain. Almost all animals in the natural world exhibit some form of adaptive behaviour, and there is increasing interest in the creation of artificial systems which are capable of acting in an adaptive manner — either simulated 'virtual agents', or real robots.

Our work to date has involved using artificial evolution on populations of simulated robots. The simulations involve a model of a real robot built at Sussex, and the simulated vision employs advanced computer graphics techniques.<sup>2</sup> Work is currently underway on the construction of specialised robotic equipment which eliminates the need for simulating perception and action, while still allowing the use of artificial evolution; see (Cliff *et al.* 1993) for further details.

<sup>2</sup>Ray-tracing with antialiasing via sixteen-fold supersampling; see e.g. (Glassner 1989).

For reasons given in (Harvey *et al.* 1993), we are approaching the task of creating artificial agents that exhibit adaptive behaviour in accordance with the following set of beliefs:

- 'Neural'-network processors are likely to be most useful in building controllers for agents that exhibit adaptive behaviour.
- Manual design of such networks is likely to become prohibitively difficult as increasingly complex or sophisticated behaviours are required. Rather than design-by-hand, we are employing artificial evolution techniques, based on Harvey's SAGA variable-length genotype methods (Harvey 1992).
- Almost all adaptive behaviours benefit from distal information. We believe visually-guided agents should be studied from as early a stage as possible.
- While we could impose on our robot some visual sensors with fixed properties, we advocate, in common with Brooks (Brooks 1992), the concurrent evolution of visual sensor morphology and the control networks; separating morphology from control is a measure which is difficult to justify from an evolutionary perspective, and potentially misleading.
- Studies of visually guided agents should commence by examining minimal systems. The work reported on here involves robots using very simple low-resolution devices coupled to small networks. It is our intention to work towards more complex (i.e. higher resolution) systems. Furthermore, because we intend to transfer our results from simulated robots to the real robot on which the simulation is based, we constrain evolution such that the evolved designs could realistically be built from discrete components and operate in real time. In effect, our intention is to evolve a specification for a robot with electronic compound eyes (c.f. (Franceschini *et al.* 1991)).

## 2.2 Details

In accordance with the last item above, our current studies have addressed evolving visually guided robots with just two photoreceptors (i.e. two 'pixels' in the input images). The direction of view of the photoreceptors, and their acceptance angles, are under evolutionary control: it is in this sense that the visual morphology is concurrently evolved along with the controller network. For full details of the genetic encoding for control networks and visual system, see (Harvey *et al.* 1993, Cliff *et al.* 1993). Because there are only two photoreceptors, we can only expect to evolve robots which exhibit relatively simple behaviours. Nevertheless, we have concentrated on evolving robots which perform tasks that would be difficult or impossible using only tactile information.

Physically, the Sussex robot is cylindrical: it has a circular bottom-plate on which the motors and wheels are mounted, and a circular top-plate where a notebook computer is situated (the computer simulates the control networks). There are two independent drive wheels, each capable of five speeds: full on, half on, off, half reverse, full reverse. There is also a large ball-bearing freewheel castor. The robot is equipped with tactile sensors giving a six-bit input vector: it has four radially oriented binary 'whiskers', and binary 'bumper-bars' at front and rear. For illustration, see (Cliff *et al.* 1993). The simulated robots are accurate models of such a vehicle, with the addition of visual sensors.




Figure 1: Schematic block diagram showing operations within a single model neuron.

While our early tactile-only work involved the robot roving around cluttered office-like environments, all the visually-guided tasks have been set in a closed circular arena. The arena has black walls, while the floor and ceiling are white. There are no obstacles: the arena contains only the robot.

The visual input from each of the robot's photoreceptors at any particular moment in time depends on the robot's visual morphology, and the position and orientation of the robot in the arena. Essentially, the population of robots has to evolve to correlate the visual input with its position in the world, so as to satisfy whatever fitness evaluation we impose on the robot's behaviours. As was demonstrated in (Cliff *et al.* 1993), visual guidance emerges without explicit reference to vision in the evaluation process. In the early stages of evolution, the tactile sensors can be useful in helping correlate visual input with the robot's position. However, as will be demonstrated below, later generations typically tend to rely *only* on visual information.

### 2.3 Networks and the 'Neuron' Model

The controller networks are continuous dynamical systems, built from model 'neurons' (i.e. processing units), which can have asymmetric and recurrent connectivities. Activation values (all real numbers in the range  $[0, 1]$ ) are transmitted between units along the connections, all of which have a weight of one, and impose a unit time delay in transmission. Fully asynchronous processing is simulated by fine-time-slice approximation techniques with random variation in time-cycling on each unit to counter periodic effects.

The neuron model has separate channels for excitation and inhibition. A schematic of the operations for one unit is shown in Figure 1. The inhibition channels operate as a 'veto' or 'grounding' mechanism: if a unit receives *any* inhibitory input, its excitatory output is reduced to zero (but it can still inhibit other units). Excitatory input from sensors or other units is summed: if this sum exceeds a specified inhibitory output threshold, the unit produces an inhibitory output. Independently, the sum of excitatory inputs has uniform noise (distribution:  $\pm n$ , where  $n$  is a real number) added, and is then passed through an excitation transfer function, the result of which forms the excitatory output for that unit, so long as the unit has not been inhibited. For further details of the excitation transfer function, see (Cliff *et al.* 1992a).

We have found that this neuron model is sufficiently sophisticated that there has been no need to introduce variable connection weights or variable delays for controllers based on the minimal visual systems studied so far. Nevertheless, we are actively investigating the use of placing such parameters within evolutionary control, as well as exploring other neuron models.

### 3 Evolving Network Controllers

The evolutionary process starts with a population of genotypes; in the work reported here, we have used populations of size 60. Each genotype consists of two chromosomes: one is an encoding of the control network, the other encodes parameters governing the visual morphology (Harvey *et al.* 1993, Cliff *et al.* 1993). Initially, all the genotypes in the population are random. On every generation, each genotype is evaluated, and assigned a fitness score. The genotypes are then 'interbred', with mutation and crossover according to SAGA principles (Harvey 1992), thereby creating a new population. This process continues for a specified number of generations (in the work discussed here, genotypes were evolved over 100 generations).

The evaluation of each genotype involves decoding the chromosomes to create a simulated robot, then testing the robot a number of times (we use eight tests per genotype). On each test, the robot is positioned at a random orientation and position in the arena (with a bias towards positions near to the wall), and then it is allowed a fixed amount of simulated time, during which its behaviour is rated according to an evaluation function  $\mathcal{E}$ .  $\mathcal{E}$  varies according to the behaviour we want the population of robots to exhibit. At the end of the eight tests, the *lowest* value of  $\mathcal{E}$  scored on the tests is used as the robot's fitness value in the reproductive phase: this ensures robust solutions.

It was our intention to impose as little structure as possible on the control networks, but it is necessary to designate some units as 'input' units (receiving activity from the robot's tactile or visual sensors), and some as 'output' units (the activity level of which determines the output of the two drive motors). Units which are neither 'input' or 'output' are referred to as 'hidden'. As will be seen later, the evolutionary process can blur these distinctions.

The initial random genotypes are created to encode for networks with all the necessary input and output units, and either one or two hidden units. Because we use Harvey's SAGA genetic algorithm, the genotypes can vary in length: longer genotypes can arise, where the increase in length corresponds to more connections or extra hidden units; but such increases in the size of the network will only be carried forward to subsequent generations if they achieve higher fitness ratings in the evaluation process. In this sense, more complex networks will develop in an *incremental* fashion.

For each  $\mathcal{E}$  we have studied, we set up eight separate random populations, and allowed them each to evolve for 100 generations.<sup>3</sup> When this was complete, we took the genotype with the highest fitness from each population, and analysed its performance. Typically, in each batch of eight populations, 3–5 of them had only improved moderately on the performance of the initial random genotypes, while the remainder were scoring close to maximum fitness. In sections 3.1 and 3.2 we illustrate the analysis process on two genotypes taken from separate populations. Both genotypes were the most-fit in their population after 100 generations, and they come from the two highest-scoring populations evolved according to the evaluation function:

<sup>3</sup>Typically, it takes approximately 24 hours on a Sun SPARC workstation to evolve one population; we evolved the eight populations in parallel, on eight separate workstations.




Figure 2: Typical behaviour of C1 controller, the robot's position at each timestep shown by an arrow. It starts near the edge of the arena, moves to the centre, and then spins on the spot.




Figure 3: C1 control network. The left-hand column are units designated as inputs: FB=Front Bumper, BLW=Back Left Whisker, etc. Right-hand column shows four output units, which are paired and differenced to give two motor signals in the range [-1,1]. Centre column shows 'hidden units'.

$$\mathcal{E} = \sum_{\forall t} \exp(-s|\mathbf{r}(t)|^2)$$

where  $\mathbf{r}(t)$  is the 2-D vector from the robot's position to the centre of the floor of the circular arena at time  $t$ , and  $\forall t$  denotes the duration of the evaluation test (the sum is essentially a discrete approximation to a temporal integral). The more time the robots spend at or near the centre of the arena, the higher they are rated. The value  $s$  is a scale factor which ensures that the robots collect no score if they are near the walls of the arena.

Under this evaluation function, the optimal behaviour is, from a random initial starting position, to move towards the centre of the arena as fast as possible, and then stay there. As will be seen, such behaviours were exhibited by both the controllers examined below. Controller 1 produced the best behaviour; Controller 2 is the second-best. For brevity, they are referred to as C1 and C2 respectively.

### 3.1 Controller 1

Typical behaviour for C1 is shown in Figure 2. The robot starts at the edge, moves to the centre, and then stays there. It holds its position at the centre by spinning on the spot; this is acceptable behaviour insofar as  $\mathcal{E}$  does not impose any penalties for energy expenditure. The genotype for C1 specifies that the two photoreceptors should have 45° acceptance angles, and be placed 6° either side of the robot's centre-line. The network for C1 is shown in Figure 3.

As is clear from Figure 3, the C1 network is unlike networks designed by humans: the way in which it works is not at all clear from examination of the diagram. However, we can identify redundant units and connections (e.g. unit 0 has no outputs, so it – and any connections to it – can be eliminated from consideration). Many of the redundant units or connections are likely to be “evolutionary scaffolding”: i.e. vestigial parts of the network which served a purpose in earlier generations but are now no longer useful. Furthermore, we can attempt to identify different sensory-motor *pathways*. For example, some of the units and connections may be involved purely in dealing with efficiently turning away from the wall if a whisker or bumper is triggered by a collision, while other parts of the network




Figure 4: Record of observables and activity levels for the activity illustrated in Figure 2. Horizontal axis is time. From top: robot's velocity; robot's orientation; visual input to left photoreceptor; visual input to right photoreceptor; output of left wheel; output of right wheel; activity levels in the control network units 0 to 14.




Figure 5: Network with redundant and non-visual units deleted: see text for further details.




Figure 6: Typical behaviour of the C1 controller in the absence of noise. See text for details.

may be dedicated to generating the visually-guided behaviour of moving to the centre and staying there. For this reason, the rest of the analysis concerns the identification of only those sensory-motor pathways involved in visual guidance.

While the control network is operating, we can record inputs, outputs, and activity levels for later analysis, along with measures of the robot's behaviour such as velocity, orientation, or distance from the centre. Figure 4 shows such a record for the behaviour sequence illustrated in Figure 2. As can be seen, some of the units are largely inactive for the duration of the sequence, and may be eliminated from consideration in the visual pathway. The results of eliminating redundant and tactile-only units are shown in Figure 5.

From Figure 5, it becomes clear that the initial categorisation of units into 'input', 'hidden', and 'output' is no longer sensible; the opportunistic nature of evolution allowing some of the tactile input units to become virtual hidden units; to use the language of neuroscience, they have become higher-order interneurons.

Correlations in activity levels are not particularly clear in Figure 4, because of the disruptive effect of internal noise; but we can switch off the internal noise and observe the controller functioning as a 'perfect' system (external noise, e.g. in the kinematics model, is *not* disabled). This is a great advantage in analysing

simulated systems, and one which is not available to neuroethologists. The performance of the robot does not degrade significantly when the noise is eliminated, although there are notable differences: Figure 6 shows typical behaviour in the absence of noise. The approach to the centre appears to occur in two phases: an initial low-radius turn (marked A1 in the figure) followed by a higher-radius turn in the opposite direction (A2), which ends in the spin phase (S). The corresponding activity-trace is shown in Figure 7.

Analysis of noise-free results such as those illustrated finally allows us to explain the activity of the network. The explanation is made easier by redrawing the network, abandoning our prior categorization of unit-types where appropriate, as in Figure 8. Unit 2 (initially categorized as an input unit) is now acting as a second-order 'interneuron'.

There follows a short explanation of the action of the network, with reference to Figures 6 to 8. All units initially have zero activity. The units active in each phase are illustrated in Figures 9 to 12.

**A1** Initially, relatively high visual input to unit 6 excites unit 2, which inhibits unit 12, so units 12 and 13 stay inactive. Meanwhile, the effects of visual input arriving at unit 11 gives a low-radius turn. Eventually, the robot turns towards the dark wall, visual input falls, so unit 2 no longer inhibits unit 12.

**Transition: A1 to A2** Momentarily, unit 12 becomes active and excites units 13 and 14. This initiates a low-radius turn in the *opposite* direction, turning the robot away from the wall. Therefore the visual input rises again, re-activating unit 2, which re-inhibits unit 12. As a consequence, unit 14 goes




Figure 7: Record of observables and activity levels for the noise-free activity illustrated in Figure 6.

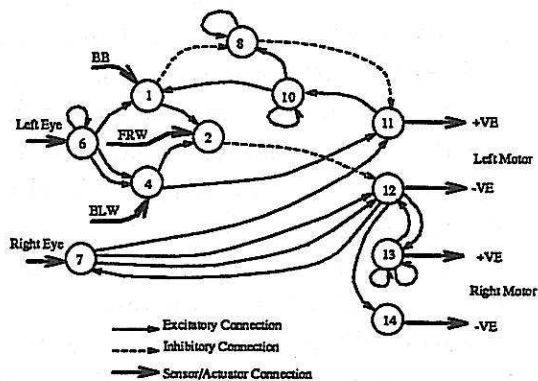


Figure 8: Final C1 network.






Figure 9: Primary active connections in phase A1. Units and connections not directly involved in producing behaviour in phase A1 have been deleted for clarity (cf. Figure 3). Motor output values are indicated.




Figure 10: Primary active connections in the momentary transition between phases A1 and A2.




Figure 11: Active connections in phase A2.




Figure 12: Active connections in phase S.

inactive, but unit 13 stays active by self-excitation.

**A2** Combined activity in units 11 and 13 give a high-radius turn, which takes the robot toward the centre of the arena. Once at the centre, the visual input drops,<sup>4</sup> and unit 2 no longer inhibits unit 12.

**S** Unit 12 becomes active, and excites unit 14. Units 11 and 13 are still active from the A2 phase. The combined activity in units 11 to 14 makes the robot spin on the spot, in the same direction as the A2 phase. During the spin, the interactions between units 1, 8, 10, and 11 can intermittently cause unit 11 to go briefly inactive, which has the effect of making the spin-position drift slightly. This is useful, in that there is a fairly large isoluminance zone near the centre. The slight drift while spinning increases the chances of the robot moving over the exact centre of the arena, where  $\mathcal{E}$  is highest, which is a better policy than fixing the spin just inside the border of the isoluminance zone. If the robot spins outside the centre-zone, it will revert to phase A2 (this can be seen in Figures 6 and 7).

The above explanation appears to account for all of the observed behaviour of the C1 controller in the absence of noise. It is clear that unit 2 is very important, acting as a switch between approach and spin phases (the A1-A2 transition, initiated by unit 2 going inactive, may be viewed as a very brief 'spin'). The same behaviour phases can also be witnessed in the with-noise behaviour, although when noise is present it is possible for unit 13 to become either active or inactive

<sup>4</sup>The visual input drops because, at the centre, the two photoreceptors (as specified by the C1 vision chromosome) 'see' more of the (dark) walls than the (light) floor or ceiling - visual input is maximal for C1 when the robot is against a wall, oriented towards the centre.




Figure 13: Typical behaviour of the C2 controller, with noise. Display format as for Figure 2. The robot starts near the edge of the arena, moves to the centre, and then spins on the spot. As can be seen, the C2 controller drives the robot in reverse.

via internal noise and its self-excitatory connections: the random noise induces a "drunkard's walk"-style drift in the excitation of unit 13, which means that in the approach to the centre the C1 controller may switch between A1 and A2 approach modes a number of times. Nevertheless, the central role of unit 2 in switching between 'approach' and 'spin' is maintained.

### 3.2 Controller 2

As with C1, the behaviour of the robots controlled by C2 is close to the best behaviour: the C2 robots make a smooth approach towards the centre, and then stay there. At the behavioural level, the performance of C2 differs from C1 in the final phase: instead of spinning on the spot, the robot makes low-radius cycling movements which hold its position near the centre; see Figure 13.

However, despite these behavioural similarities, the C2 morphology and controller differ significantly from the C1 controller. Here the differences will be briefly sketched; for full details see (Cliff *et al.* 1992b).

First, the C2 visual morphology specifies 45° photoreceptors (as with C1), but they are placed 60° either side of the robot's centre-line. Figure 14 shows the full C2 control network, while Figure 15 shows the final visual-guidance pathways in the network.

The discussion in (Cliff *et al.* 1992b) shows that unit 6 is effectively redundant, which implies that C2 is employing 'monocular' vision, using just the input from

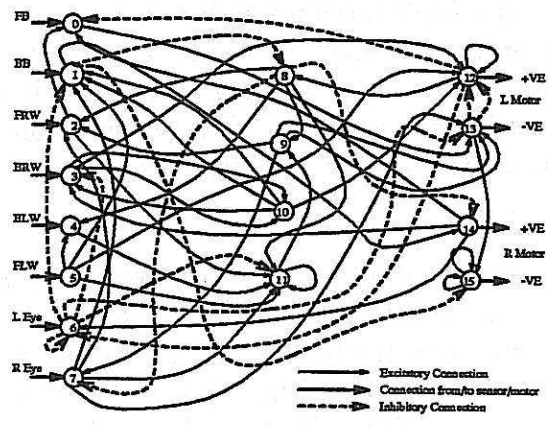


Figure 14: Full C2 control network. Display format as for Figure 3.




Figure 15: C2 visual guidance pathways. For clarity, the positions of the left and right motor outputs have been interchanged.




Figure 16: C2 as a 'monocular' network.




Figure 17: C2 as a 'weighted' network.

the right-hand photoreceptor to perform visual guidance. Further analysis shows that some units in C2 are operating as *distributor* units, which can be considered as allowing for "virtual connections" with non-unit weights or delays; for example, in Figure 16, unit 9 is a distributor for unit 11.

From the above analysis, it is clear that while C2 produces similar observable behaviour to C1, the internal mechanisms responsible for generating these behaviours operate on markedly different principles.

### 3.3 Discussion

The primary factor of note in comparing controllers C1 and C2 is that, although they were evolved separately, they had indistinguishable initial populations (i.e. both populations were composed entirely of random genotypes). After 100 generations, both populations show a high degree of convergence, in that the genomes for all individuals in the population are fairly similar. Also, both populations perform approximately similar behaviours. Yet, as was made clear above, there are significant differences between C1 and C2 in both visual morphology and control networks. The two populations therefore show a form of speciation, in that the two populations can be considered as different *species*, performing the same task. This is an accordance with the principles underlying the SAGA genetic algorithm we used (Harvey 1992).

Such networks exhibited graceful degradation in the presence of increased noise. During evolution, an internal noise distribution of  $\pm 0.1$  was used; we found the robots could still approach the centre with noise distributions as high as (in the case of C1)  $\pm 0.8$ : see (Cliff *et al.* 1992a).

In almost all of the networks we have analysed, there has been no clearly identifiable structure. C2 is a clear example. Nevertheless, we find the structure of C1 intriguing: the role of unit 2, which can disable unit 12 (and, in doing so, also disables unit 14) seems vaguely reminiscent of a two-layer subsumption architecture, in that units 12 and 14 are responsible for generating the 'spin' behaviour; a behaviour 'subsumed' by the approaching behaviour. See (Brooks 1985, Brooks 1986) for details of subsumption architectures, and e.g. (Franceschini *et al.* 1991) for an example of a two-layer subsumption visually guided robot. Clearly, it is too early to make strong claims, but we suspect that it is not infeasible that subsumption-style architectures could evolve within our scheme: because we use *truly* incremental evolution, it is possible that mechanisms generating elementary low-level behaviours evolve first, with structures responsible for generating higher-level behaviours coming later. Such an evolutionary trajectory would make sense,

given the need for satisfying intermediate viability (i.e. good controllers have to be built from minor changes to earlier slightly-less-good controllers – there is no opportunity for a total re-design from scratch). This may go some way toward explaining why subsumption-style controllers (i.e. behavioural decomposition) have been identified in biological creatures (Altman and Kien 1989, Cliff 1992).

It is important to note that both the C1 and C2 controllers were evolved in a fixed-size arena, and hence are dependent on the ratio of the height of the arena's walls to the diameter of the floor. It is this ratio, combined with the controller's particular visual morphology, that determines the brightness values in the central isoluminance zone discussed in Section 3.1. Work is currently underway on varying the arena dimensions on each evaluation, in order to evolve truly general-purpose controllers which should operate in *any* circular arena.

### 3.4 Conclusion

We have examined two controller networks evolved using incremental genetic algorithms, and found a form of speciation, in that two controllers evolved in separate populations produce convergent behaviours while employing divergent mechanisms for generating those behaviours. Nevertheless, both controllers perform in a close-to-optimal manner, and are robust in the presence of noise. While both the robot's world and behaviours are relatively trivial, we can see no reason why our methods, suitable extended beyond the specifics described here, should not prove successful in increasingly complex domains.

The important achievement in this paper is not that we got a simulated robot to perform a particular visually guided behaviour, nor that the behaviours were generated by evolved neural networks. What matters is that we haven't treated the evolved networks as magic black boxes. We specified *what* the robots should do, but not *how* the controllers work. Nevertheless, analysis lets us know what's going on inside the box. And, for the record, we don't think that it's computation (at least, not in the conventional sense).

### References

- [Altman and Kien 1989] J. S. Altman and J. Kien. New models for motor control. *Neural Computation*, 1:173–183, 1989.
- [Beer 1990] R. D. Beer. *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology*. Academic Press, 1990.
- [Beer 1992] R.D. Beer. A dynamical systems perspective on autonomous agents. Tech. Report CES-92-11, Case Western Reserve University, Cleveland, Ohio, 1992.
- [Brooks 1985] R. A. Brooks. A robust layered control system for a mobile robot. A.I. Memo 864, M.I.T. A.I. Lab, September 1985.
- [Brooks 1986] R. A. Brooks. Achieving artificial intelligence through building robots. A.I. Memo 899, M.I.T. A.I. Lab, May 1986.
- [Brooks 1992] R. A. Brooks. Artificial life and real robots. In F. Varela and P. Bourguine, eds., *Toward a Practice of Autonomous Systems: Proc. First Eur. Conf. on Artificial Life (ECAL91)*, pages 3–10, Cambridge MA, 1992. M.I.T. Press Bradford Books.
- [Camhi 1984] J. M. Camhi. *Neuroethology: Nerve Cells and the Natural Behaviour of Animals*. Sinauer Associates Inc., Sunderland, Mass., 1984.

- [Cliff *et al.* 1992a] D. T. Cliff, I. Harvey, and P. Husbands. Incremental evolution of neural network architectures for adaptive behaviour. Tech. Report CSRP 256, Univ. of Sussex, School of Cognitive and Computing Sciences, 1992.
- [Cliff *et al.* 1992b] D. T. Cliff, P. Husbands, and I. Harvey. Analysis of evolved sensory motor controllers. Technical Report CSRP 264, Univ. of Sussex, School of Cognitive and Computing Sciences, 1992.
- [Cliff *et al.* 1993] D. T. Cliff, P. Husbands, and I. Harvey. Evolving visually guided robots. In J.-A. Meyer, H. Roitblat, and S. Wilson, eds., *Proc. Second Intl. Conf. on Simulation of Adaptive Behaviour (SAB92)*. MIT Press Bradford Books, Cambridge, MA. In Press. Also available as Univ. of Sussex, COGS Tech. Report CSRP220.
- [Cliff 1991] D. T. Cliff. Computational neuroethology: A provisional manifesto. In J.-A. Meyer and S. W. Wilson, eds., *From Animals to Animats: Proc. First Intl. Conf. on Simulation of Adaptive Behavior (SAB90)*, pages 29–39, Cambridge MA, 1991. M.I.T. Press — Bradford Books. Also available as Univ. of Sussex, COGS Tech. Report CSRP162.
- [Cliff 1992] D. T. Cliff. Neural networks for visual tracking in an artificial fly. In F. J. Varela and P. Bourguine, eds., *Toward a Practice of Autonomous Systems: Proc. First Eur. Conf. on Artificial Life (ECAL91)*, pages 78–87. MIT Press Bradford Books, Cambridge, MA.
- [Franceschini *et al.* 1991] N. Franceschini, J.-M. Pichon, and C. Blanes. Real time visuo-motor control: from flies to robots. In *Proc. 1991 Intl. Conf. on Advanced Robotics, Pisa*, 1991.
- [Glassner 1989] A. S. Glassner, editor. *An Introduction to Ray Tracing*. Academic Press, London, 1989.
- [Harvey *et al.* 1993] I. Harvey, P. Husbands, and D. Cliff. Issues in evolutionary robotics. In J.-A. Meyer, H. Roitblat, and S. Wilson, eds., *Proc. Second Intl. Conf. on Simulation of Adaptive Behaviour (SAB92)*. M.I.T. Press Bradford Books, Cambridge MA. Also available as Univ. of Sussex, COGS Tech. Report CSRP219.
- [Harvey 1992] I. Harvey. Species adaptation genetic algorithms: A basis for a continuing SAGA. In F.J. Varela and P. Bourguine, eds., *Toward a Practice of Autonomous Systems: Proc. First Eur. Conf. on Artificial Life (ECAL91)*, pages 346–354. M.I.T. Press — Bradford Books, Cambridge MA. Also available as Univ. of Sussex, COGS Tech. Report CSRP221.
- [Husbands and Harvey 1992] P. Husbands and I. Harvey. Evolution versus design: Controlling autonomous robots. In *Integrating Perception, Planning and Action, Proc. 3rd Annual Conference on Artificial Intelligence, Simulation and Planning*, pages 139–146. IEEE Press, 1992.
- [Husbands *et al.* 1993] P. Husbands, I. Harvey, and D. Cliff. Analysing Recurrent Dynamical Networks Evolved for Robot Control. In *Proc. 3rd IEE Intl. Conf. on ANNs.* IEE Press, 1993.
- [Smithers 1992] T. Smithers. Taking eliminative materialism seriously: A methodology for autonomous systems research. In F. Varela and P. Bourguine, eds., *Toward a Practice of Autonomous Systems: Proc. First Eur. Conf. on Artificial Life (ECAL91)*, pages 31–40. MIT Press Bradford Books, Cambridge, MA.
- [van Gelder 1992] T. van Gelder. What might cognition be if not computation. Technical Report 75, Indiana University Cognitive Sciences, 1992.