Analysis of Evolved Sensory-Motor Controllers

D. Cliff, P. Husbands, I. Harvey CSRP 264, December 1992

Cognitive Science Research Paper

Serial No. CSRP 264

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

This is an expanded version of a paper to be presented at the Second European Conference on Artificial Life, Brussels, May 24–26 1993.

Analysis of Evolved Sensory-Motor Controllers

Dave Cliff ^{1,2} and Philip Husbands¹ and Inman Harvey¹ School of Cognitive and Computing Sciences ²Neuroscience 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 action of the evolved networks is described. We discuss the process of analysing evolved artificial networks, a process which bears many similarities to analysing biological nervous systems in the field of neuroethology.

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.

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. Instead, we analyse our networks using techniques analogous to those used in the study of biological sensory-motor neural systems. 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. (Neuroethology is the study of the neural mechanisms underlying the generation of a creature's behaviour; see e.g. [7].) For further details of the link between neuroethology and artificial neural network research, see [8, 2].

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

¹For 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.

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. [3, 19, 17] 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. The text refers the readers to past papers 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 [16, 15, 11, 10]. 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 'adaptive behaviour', we mean behaviour which is selected to increase the chances that a situated agent can survive in an environment which is noisy, dynamic, hostile, 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. The artificial systems are commonly 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.² 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 [11] for further details.

For reasons given in [15], 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.

²Namely, ray-tracing with antialiasing via sixteen-fold supersampling (see e.g. [13] for details of such techniques).

- 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 [14].
- Almost all adaptive behaviours benefit from distal (i.e. long-range) sensory information. While there is an established body of successful work studying robots with only tactile sensing (i.e. mechanical whiskers), the proximal nature and low dimensionality of the robot's sensors constrain it to relatively primitive "bumping and feeling" behaviours, such as wall-following. For demonstration of our methods working with only proximal sensors, see [16, 15]. A primary means of gathering distal sensory information is by use of visual sensing, so 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 [6]) 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.
- For reasons of parsimony, 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. [12]).

2.2 Details

In accordance with the last item in the above list, 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 in this sense that the visual morphology is concurrently evolved along with the controller network. For full details of the genetic encoding for both the control networks and the visual system, see [15, 11]. 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). The robot has three wheels arranged to give tripod stability. At the front are two independent drive wheels, each capable of rotating at one of five speeds: full on, half on, off, half reverse, full reverse.

The rear wheel is 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 [11]. The simulated robots are accurate models of such a vehicle, with the addition of visual sensors.

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 [11], 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 [10].

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.

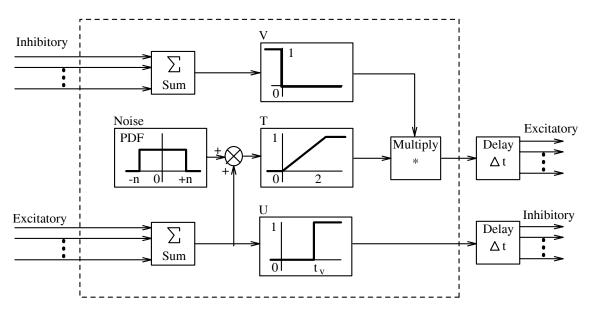


Figure 1: Schematic block diagram showing operations within a single model neuron. See text for further explanation.

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 [15, 11]. 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 [14], 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 (if the best or average \mathcal{E} -value is used, it can be deceptively high).

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.³ 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:

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

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). Put most simply, 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 at the centre, 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. As can be seen, 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

³Typically, 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.

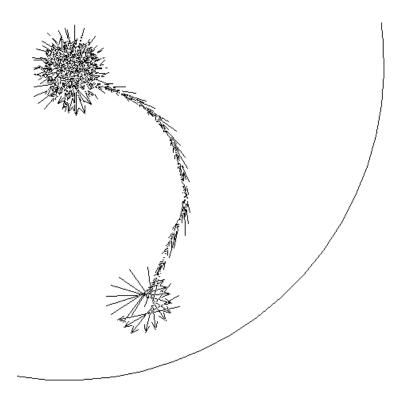


Figure 2: Typical behaviour of the C1 controller. The robot's position at each timestep is shown by an arrow; the midpoint of the arrow 'shaft' is the centre of the robot, and the length of the shaft is the same as the robot's diameter. The robot starts near the edge of the arena, moves to the centre, and then spins on the spot. The 'tip' of the arrow shows the 'front' of the robot, which is not necessarily the direction of travel: although in this case the robot is moving forwards, it can travel in reverse.

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 cf. [18]. 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 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.

Furthermore, while the control network is operating, we can record inputs, outputs, and activity levels for later analysis, along with important measures of the robot's behaviour (such as its 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 (if consistently inactive) can 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 is such that some of the tactile input units have been taken over to act as virtual hidden units; to use the language of neuroscience, they have become higher-order interneurons.

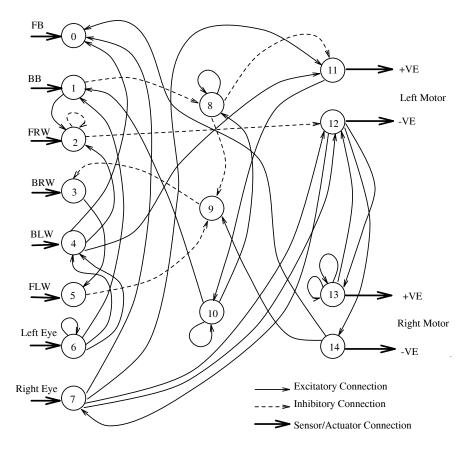


Figure 3: C1 control network. The left-hand column are units originally designated as input units: FB=Front Bumper; BB=Back Bumper; FRW=Front Right Whisker; BRW=Back Right Whisker; BLW=Back Left Whisker; FLW=Front Left Whisker. Right-hand column shows output units, which are paired and differenced to give two motor signals in the range [-1,1] from four 'neuron' outputs in the range [0,1]. Centre column shows 'hidden units'.

To further elucidate what is occurring, we have one more tactic at our disposal: 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. As can be seen, the approach to the centre appears to occur in two phases: an initial low-radius turn followed by a higher-radius turn in the opposite direction, which ends in the spin phase. These phases are marked on the figure as "A1" and "A2" for Approach-1 and Approach-2; and "S" for spin. 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. The redrawn C1 network is shown in Figure 8. As can be seen, unit 2 (initially categorized as an input

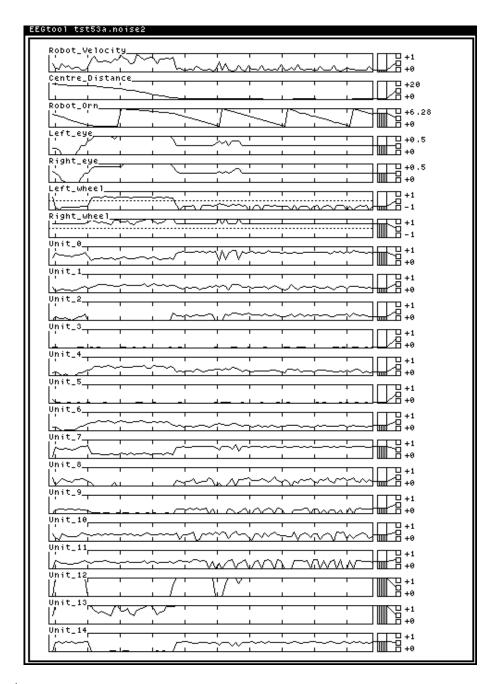


Figure 4: Record of observables and activity levels for the activity illustrated in Figure 2. Horizontal axis is time. From top: robot's velocity; robot's orientation; visual input to left photoreceptor; visual input to right photoreceptor; output of left wheel; output of right wheel; activity levels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

unit) is now acting as a second-order 'interneuron'.

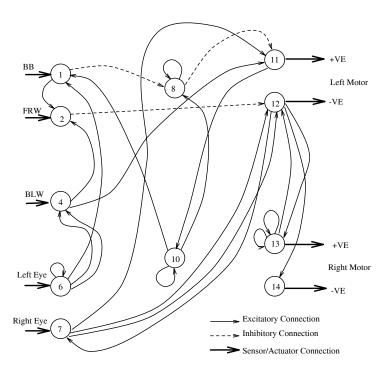


Figure 5: Network with redundant and non-visual units deleted: see text for further details.

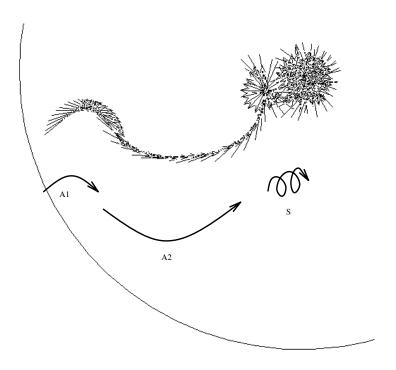


Figure 6: Typical behaviour of the C1 controller in the absence of noise. See text for details.

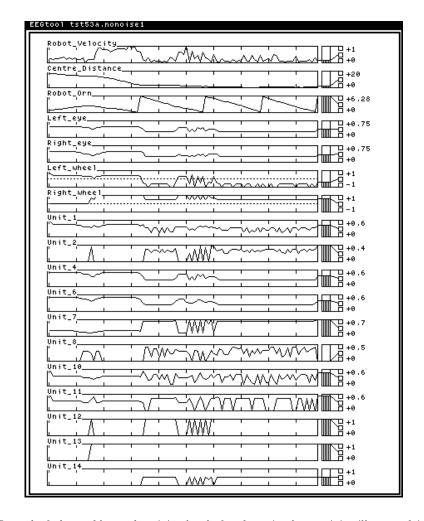


Figure 7: Record of observables and activity levels for the noise-free activity illustrated in Figure 6.

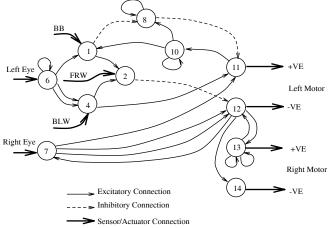


Figure 8: Final C1 network.

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 and the visual input falls, so unit 2 no longer inhibits unit 12.

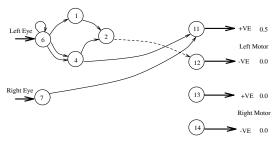


Figure 9: Primary active connections in phase A1. Units and connections not directly involved in producing behaviour in phase A1 have been deleted for clarity (cf. Figure 3). Motor output values are indicated. See text for further details.

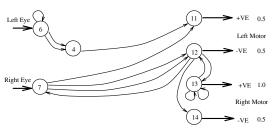


Figure 10: Primary active connections in the momentary transition between phases A1 and A2.

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 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,⁴ 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 (i.e. for practical purposes, the visual input is identical at the centre and also at positions slightly off-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

⁴The 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: in this case it 'sees' mainly the floor and ceiling, with the distant far wall taking up little of the visual field.

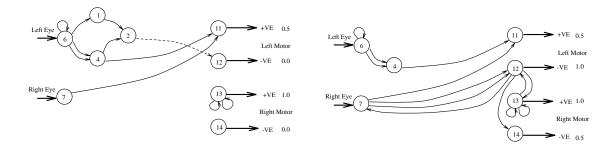


Figure 11: Active connections in phase A2.

Figure 12: Active connections in phase S.

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 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. Typical behaviour of C2 (with noise) is shown in Figure 13.

However, despite these behavioural similarities, the C2 morphology and controller differ significantly from the C1 controller. First, the C2 visual morphology specifies 45° photoreceptors (as with C1), but they are placed 60° either side of the robot's centre-line (cf 6° specified for C1). Furthermore, the control network bears very little resemblance to the C1 network. Figure 14 shows the full network, while Figure 15 shows the final visual-guidance pathways in the network, revealed using the same analysis techniques as for C1: the activity trace for the behaviour of Figure 13 is shown in Figure 16; while a noise-free behaviour sequence and activity trace are shown in Figures 17 and 18 respectively.

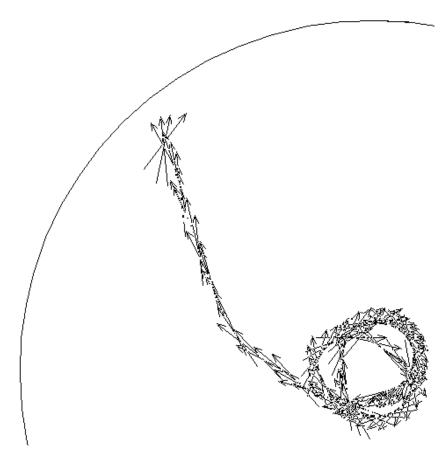


Figure 13: Typical behaviour of the C2 controller, with noise. Display format as for Figure 2. The robot starts near the edge of the arena, moves to the centre, and then spins on the spot. As can be seen, the C2 controller drives the robot in reverse (backwards).

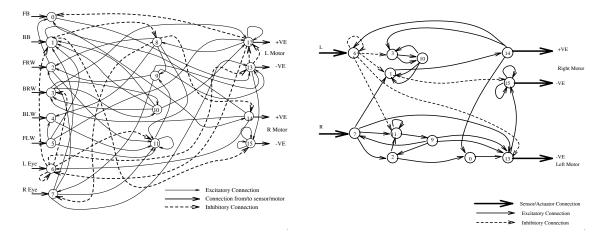


Figure 14: Full C2 control network. Display format as for Figure 3.

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



Figure 16: Record of observables and activity levels for the (with-noise) activity illustrated in Figure 13.

Examination of the activity traces (both with and without noise) allow the analysis of C2 to be taken further. First, unit 6 provides only veto outputs to other units, and it is clear from Figures 16 and 18 that the total input to unit 6 is never sufficiently high to go over the veto-output threshold, so unit 6 is effectively redundant in the context of producing the behaviours illustrated in the figures. For this reason, unit 6 can be eliminated and the C2 network re-drawn accordingly: see Figure 19. This implies that C2 is employing 'monocular' vision, using just the input from the right-hand photoreceptor to perform visual guidance.

To further elucidate how C2 operates, one more analysis technique can be used to improve the legibility of the network diagrams: it can be seen in Figure 19 that some units recieve activation from only one unit, pass that activation through the excitation transfer function, and then provide excitatory input to other units. We refer to such units as distributor units. For example, in Figure 19, unit 9 is a distributor for unit 11.

If, for ease of analysis, we ignore the internal noise in distributor units, then if unit i connects to unit k via a distributor unit j, the only effects of the distributor j are to act as a weight on the connection strength between i and k, and to double the time-delay on activity passing from i to k: the nature of the weight is determined by the excitation transfer function. Given that in the current system all units have the same excitatory transfer function, with a fixed gradient of 0.5 on the linear ramp between the lower and upper thresholds (cf. Figure 1 and [10]), distributor units are acting as doubly-delayed connections with weight 0.5. In Section 2.3 it was stated that all links have a weight

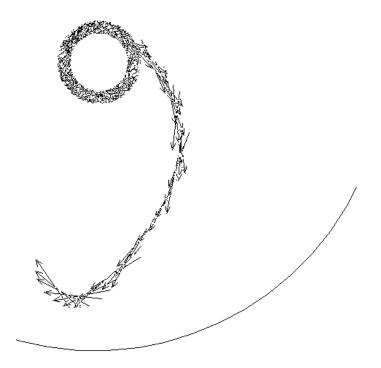


Figure 17: Typical noise-free behaviour of the C2 controller. Display format as for Figure 2.

of one and impose a unit time delay: the use of distributor units allows for "virtual connections" to evolve which have different weights or delays. Similarly, if there are N excitatory links from unit n to unit m, then they can together be considered as a single "virtual connection" with weight N and unity delay: for example, in Figure 19, the two connections from unit 10 to unit 1 form a virtual connection of weight 2.0.

Thus, distributor units and multiple connections between units can be eliminated from the network diagrams, and the network re-drawn with the various weights indicated: the final "weighted" version of C2 is shown in Figure 20. The "weighted" forms of the networks are useful analytic tools: from Figure 20 it is fairly clear that the operation of C2 depends crucially on unit 1: if there is sufficient visual input to the right photoreceptor (through unit 7), unit 1 inhibits unit 15, and the robot enters a low-radius turn: the turn is only sustainable when the robot is within the central isoluminance zone; at other locations the turn will reduce visual input, thereby preventing continued inhibition of unit 15, so the robot's path to the centre is a straight line punctuated by brief bursts of low-radius turns as unit 15 is intermittently inhibited.

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. This is discussed further below.

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

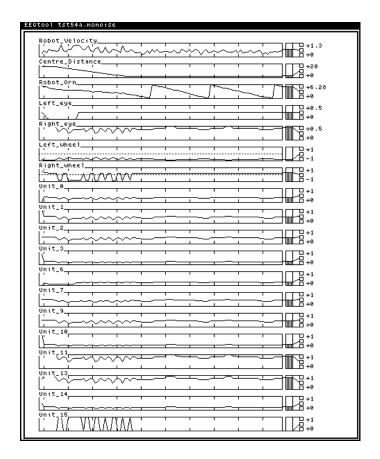


Figure 18: Record of observables and activity levels for the noise-free activity illustrated in Figure 17.

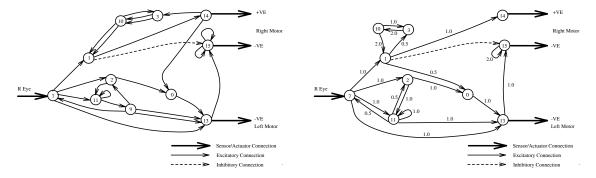


Figure 19: C2 as a 'monocular' network.

Figure 20: C2 as a 'weighted' network: see text for further details.

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 [14].

Such networks exhibited graceful degradation in the presence of increased noise. During evolution, an internal noise distribution of ± 0.1 was used; we found the robots could still approach the centre with noise distributions as high as (in the case of C1) ± 0.8 : see [10].

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 [4, 5] for details of subsumption architectures, and e.g. [12] 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 slightlyless-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 [1, 9].

It is important to note that both the C1 and C2 controllers were evolved in a fixedsize 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

- [1] J. S. Altman and J. Kien. New models for motor control. *Neural Computation*, 1:173–183, 1989.
- [2] R. D. Beer. Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology. Academic Press, 1990.
- [3] R.D. Beer. A dynamical systems perspective on autonomous agents. Technical Report CES-92-11, Case Western Reserve University, Cleveland, Ohio, 1992.
- [4] R. A. Brooks. A robust layered control system for a mobile robot. A.I. Memo 864, M.I.T. A.I. Lab, September 1985.
- [5] R. A. Brooks. Achieving artificial intelligence through building robots. A.I. Memo 899, M.I.T. A.I. Lab, May 1986.
- [6] R. A. Brooks. Artificial life and real robots. In F. J. Varela and P. Bourgine, editors, Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91), pages 3-10, Cambridge MA, 1992. M.I.T. Press — Bradford Books.
- [7] J. M. Camhi. Neuroethology: Nerve Cells and the Natural Behaviour of Animals. Sinauer Associates Inc., Sunderland, Mass., 1984.
- [8] D. T. Cliff. Computational neuroethology: A provisional manifesto. In J.-A. Meyer and S. W. Wilson, editors, From Animals to Animals: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90), pages 29–39, Cambridge MA, 1991. M.I.T. Press Bradford Books. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP162.
- [9] D. T. Cliff. Neural networks for visual tracking in an artificial fly. In F. J. Varela and P. Bourgine, editors, Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91), pages 78–87. MIT Press Bradford Books, Cambridge, MA, 1992.
- [10] D. T. Cliff, I. Harvey, and P. Husbands. Incremental evolution of neural network architectures for adaptive behaviour. Technical Report CSRP 256, University of Sussex School of Cognitive and Computing Sciences, 1992.
- [11] 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, 1993. In Press. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP220.
- [12] N. Franceschini, J.-M. Pichon, and C. Blanes. Real time visuomotor control: from flies to robots. In Proceedings of the 1991 International Conference on Advanced Robotics, Pisa, 1991.

- [13] A. S. Glassner, editor. An Introduction to Ray Tracing. Academic Press, London, 1989
- [14] I. Harvey. Species adaptation genetic algorithms: A basis for a continuing SAGA. In F.J. Varela and P. Bourgine, editors, Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91), pages 346–354. M.I.T. Press Bradford Books, Cambridge MA, 1992. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP221.
- [15] I. Harvey, P. Husbands, and D. 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). M.I.T. Press Bradford Books, Cambridge MA, 1993. In Press. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP219.
- [16] 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.
- [17] T. Smithers. Taking eliminative materialism seriously: A methodology for autonomous systems research. In F. J. Varela and P. Bourgine, editors, Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91), pages 31–40. MIT Press Bradford Books, Cambridge, MA, 1992.
- [18] D. Stork, B. Jackson, and S. Walker. 'Non-Optimality' via pre-adaptation in simple neural systems. In C. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 409–429. Addison Wesley, 1992.
- [19] T. van Gelder. What might cognition be if not computation. Technical Report 75, Indiana University Cognitive Sciences, 1992.