

Artificial Evolution of Visual Control Systems for Robots*

Dave Cliff, Inman Harvey, and Phil Husbands†

Abstract

Many arthropods (particularly insects) exhibit sophisticated visually guided behaviours. Yet in most cases the behaviours are guided by input from a few hundreds or thousands of “pixels” (i.e. ommatidia in the compound eye). Inspired by this observation, we have for several years been exploring the possibilities of visually guided robots with low-bandwidth vision. Rather than design the robot controllers by hand, we use artificial evolution (in the form of an extended genetic algorithm) to automatically generate the architectures for artificial neural networks which generate effective sensory-motor coordination when controlling mobile robots. Analytic techniques drawn from neuroethology and dynamical systems theory allow us to understand how the evolved robot controllers function, and to predict their behaviour in environments other than those used during the evolutionary process. Initial experiments were performed in simulation, but the techniques have now been successfully transferred to work with a variety of real physical robot platforms. This chapter reviews our past work, concentrating on the analysis of evolved controllers, and gives an overview of our current research. We conclude with a discussion of the application of our evolutionary techniques to problems in biological vision.

*To appear in M. Srinivisan and S. Venkatesh (eds) *From Living Eyes to Seeing Machines*. Oxford University Press, in press 1996.

†The authors are all from the Evolutionary and Adaptive Systems Research Group, School of Cognitive and Computing Sciences, University of Sussex, BRIGHTON BN1 9QH, U.K. E-mail `davec` or `philh` or `inmanh`, all `@cogs.susx.ac.uk`

1 Introduction

Typically, machine vision systems employ high-bandwidth images with hundreds of thousands of pixels. Yet many visual animals (particularly arthropods) exhibit sophisticated visually guided behaviours with orders of magnitude fewer “pixels” (i.e. ommatidia in the compound eye). Inspired by this observation, we have for several years been exploring the possibilities of visually guided robots with ultra-low-bandwidth vision. Rather than design the robot controllers by hand, we use artificial evolution (in the form of an extended genetic algorithm: see e.g. [23, 18]) to automatically generate the architectures for artificial neural networks which generate effective sensory-motor coordination when controlling mobile robots.

Initial experiments were performed in simulation, but the techniques have since been successfully transferred to work with a variety of real physical robot platforms.

The mathematics of our particular style of neural networks are such that it would be difficult or impossible to derive closed-form equations describing the action of the networks. 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. 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. [10].) For further details of the link between neuroethology and artificial neural network research, see [11, 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 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, 4, 33, 30] for further discussion of the benefits of adopting a dynamical systems perspective).

In this chapter we show the analysis of 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. Our results do not rely on the commonly held belief that internal representations are central to visual processing.

The primary focus here is on analysing networks resulting from the evolutionary processes. We refer the reader 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 of the two networks. Finally, Section 4 discusses our current research directions.

2 Background

2.1 Rationale

The rationale for our work, and some early results, have been discussed elsewhere [24, 21, 15, 13]. 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, unforgiving, 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.

The work discussed here uses artificial evolution on populations of simulated robots. The simulations involve a model of a real robot built at Sussex, and advanced computer graphics techniques (namely, ray-tracing with antialiasing via sixteen-fold supersampling: see e.g. [17]) are used to faithfully simulate optical image-formation.

For reasons given in [14], 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 [19, 20].
- 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 such robots’ sensors constrain them to relatively primitive “bumping and feeling” behaviours, such as wall-following. For demonstration of our methods working with only proximal sensors, see [24, 21]. Vision is a distal sensory modality that has been extensively studied in both natural and artificial systems, and 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 [9]) 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-bandwidth, low-resolution devices coupled to small networks. It is our intention to work towards more complex (i.e. higher bandwidth) systems. Furthermore, because we are interested in results which could realistically be transferred from simulated robots to the real robot on which the simulation is based, we constrain evolution such that the evolved designs could 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. [16]).

3 Analysing Evolved Sensory-Motor Controllers

In accordance with the last item in the above list, our early work 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 both the control networks and the visual system, see [14]. 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.

Section 3.1 describes the methods used in our early explorations in evolving visual sensory-motor controllers for simulated robots. Because we deliberately chose to start with attempts at evolving simple behaviours in uncomplicated environments with ultra-low-bandwidth vision, the task of analysing the resultant evolved controllers was eased. Section 3.2 discusses insights gained from qualitative analysis of two evolved controllers. The qualitative analysis techniques are similar

to those used in neuroethology. Section 3.3 then gives a brief account of quantitative analysis of the second of the two controllers. The quantitative analysis, which employs techniques from dynamical systems theory, makes it possible to predict how the robot will behave in environments different from the one it was evolved in.

3.1 Methods

3.1.1 The Robot and its Environment

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 robot steers by differential drive: if there is a difference in the drive settings of the two wheels, the robot will move in an arc (or spin on the spot if the wheel settings are equal in magnitude but opposite in sign); if the settings are identical, the robot moves in a (roughly) straight line. 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 [15]. 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, the visually-guided tasks analysed here were 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 visual input (and internal

state) with motor outputs in order to score highly on whatever fitness metric we impose on the robot's behaviours. As was demonstrated in [15], 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 situation. However, as will be demonstrated below, later generations typically tend to rely *only* on visual information.

3.1.2 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 [13].

***** Figure 1 near here *****

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.

3.1.3 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 [14]. 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 [19, 20], 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} is set 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.¹ 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, half 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 3.2.1 and 3.2.2 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 posi-

¹Typically, it takes approximately 24 hours on a Sun SPARC-2 workstation to evolve one population; we evolved the eight populations in parallel, on eight separate workstations.

tion, to move towards the centre of the arena as fast as possible, and then stay at the centre. As will be seen, such behaviours were exhibited by both the controllers examined below. Controller C1 produced the best behaviour; Controller C2 is the second-best.

3.2 Qualitative Analysis

3.2.1 C1: 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.

***** Figure 2 near here *****

***** Figure 3 near here *****

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 c.f. [31]. 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.

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.

***** Figure 4 near here *****

***** Figure 5 near here *****

***** Figure 6 near here *****

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.

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 C1 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 referred to 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 helps us to explain the activity of the network. The explanation is made easier by redrawing the network, abandoning our prior categorisation of unit-types where appropriate. The redrawn C1 network is shown in Figure 8. As can be seen, unit 2 (initially categorised as an input unit) is now acting as a second-order ‘interneuron’.

***** Figure 7 near here *****

***** Figure 8 near here *****

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 near here *****

***** Figure 10 near here *****

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,² 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 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).

***** Figure 11 near here *****

***** Figure 12 near here *****

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

²The 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.

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 activity level 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.

3.2.2 C2: 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.

***** Figure 13 near here *****

***** Figure 14 near here *****

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 (c.f. 6° specified for C1). Furthermore, the control network bears very little resemblance to the C1 network. Figure 15 shows the full network, while Figure 16 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 14; while a noise-free behaviour sequence and activity trace are shown

in Figures 17 and 18 respectively.

***** Figure 15 near here *****

***** Figure 16 near here *****

***** Figure 17 near here *****

***** Figure 18 near here *****

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 14 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. Further explanation of how C2 operates is given in Section 3.3.

The legibility of the network diagrams can be improved by applying a technique of identifying “virtual connections”. It can be seen in Figure 19 that some units receive 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.

***** Figure 19 near here *****

***** Figure 20 near here *****

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 (c.f. Figure 1 and [13]), distributor units are acting as doubly-delayed connections with weight 0.5. In Section 3.1.2 it was stated that all links have a weight 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.

3.2.3 Discussion

The primary factor of note in comparing controllers C1 and C2 is that, although they were evolved separately, they had statistically 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 [19, 20].

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 [13].

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 [5, 6] for details of subsumption architectures, and e.g. [16] 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. behaviourally decomposed neuronal networks) have been identified in biological creatures [1, 12].

3.2.4 Summary

In this section, 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.

The task of analysing an evolved neural-network robot controller is similar to the task of analysing a neuronal network in a real animal. The techniques we have employed bear some resemblance to those used in neuroethology, and they give broadly similar results: a causal mechanistic account of the process by which perturbations in the system’s inputs give rise to alteration in the system’s outputs. That is, the internal mechanisms of the agent are not treated as a black-box, and so it is possible to understand how the observed behaviour is generated. However, the qualitative understanding does not help us predict the behaviour of the agent in different environments. For this, quantitative techniques are required.

3.3 Quantitative Analysis

It is important to note that both the C1 and C2 controllers were evolved in a fixed-size arena, and hence may have come to be 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 mentioned in Section 3.2.1.

To test whether this was the case, we evaluated C1 and C2 in a number of environments where the height-to-radius ratio varied by roughly an order of magnitude. We varied the ratio by altering the height of the arena, thereby allowing us to use the same evaluation function \mathcal{E} . To our surprise, we found that both C1 and C2 gave fairly consistent results: both showed roughly the same performance when the height was increased; when the height was decreased, the behaviour of C1 degraded significantly, but C2 showed only a slight loss of performance. It appeared that C2 had evolved to be a *general* controller, in the sense that it would reliably find the center of any circular arena with a black wall and white floor and ceiling, within wide limits.

While the empirical evidence was highly suggestive, we sought analytic assurance of the generality of C2. While the qualitative analysis techniques discussed above can explain the generation of particular patterns of behaviours in particular environments, it is difficult to see how they can be extended to account for behaviour in a continuous range of environments. Hence, to complement the qualitative analysis, we developed a quantitative analysis of C2, drawing on ideas from dynamical systems theory. Full details are given in [25]: we present a brief summary here.

Two primary techniques were employed: a detailed mathematical analysis of the properties of generic small sub-networks of noisy units with excitatory feedback loops, and a low-dimensional visualisation of the state-space trajectories that result from the interaction dynamics of the controller, the sensory-motor morphology, and the environment. The visualisation was a 2-D space which allowed the representation of an individual photoreceptor’s input signal when the robot was at any position and orientation (i.e. at any point in configuration space) within an arena with given height:radius ratio.

The qualitative techniques of Section 3.2 allowed for the simplified C2 network of Figure 16 to be used as the starting point for analysis. The results of the mathematical analysis of generic

small feedback loop networks were applied to feedback loops identified in the C2 network, yielding predictions of the average output values of units in the network under steady-state input conditions. The predictions were averages because of the presence of noise in the network. These predictions allowed for a vector field to be plotted on the 2-D visualisation space, and hence for state-space attractors to be identified. Using this approach over a range of height:radius ratios showed that for any given value of the ratio, there would be a single dominant attractor which lead the C2-controlled robot to orbit around the center of the arena clockwise or anticlockwise. The radius of the orbit was fairly constant, always being such that the robot would ‘park’ one of its wheels very close to the center of the arena, switch off the motor on that wheel, and then drive the other wheel in either forward or reverse. In some cases, the dominant attractor would be accompanied by a weak attractor, specifying orbit in the opposite direction. Simulation studies of the mathematical model of the controller were in good agreement with results from the ‘real’ system.

This quantitative analysis demonstrated that the dominant attractor was always the result of input from just one of the two photoreceptors: at any particular value of the ratio, the network was operating monocularly – the other receptor could be ablated without significant loss of performance. But the value of the ratio determined which photoreceptor was dominant.

This result provides an existence proof that, in order in order to evolve general-purpose controllers which should operate in a wide range of circular arenas, it is not necessary to vary the environments used in evaluation during the evolutionary process. Rather, general controllers may result from applying a selection pressure which favours *robust* designs.

A further point to note is that, with only one photoreceptor, there is a lack of instantaneous spatial resolution in the input: the C2 network had evolved to rely on the time dimension (i.e. temporal variation in the single input) in order to generate appropriate behaviour. In essence, it relied on the rotational symmetry of the arena: when following a circular orbit that is concentric with the arena’s floor-circle, there is no temporal variation in the visual input. C2 had evolved to

achieve a stable orbit at a low radius, bringing it to the center of the arena and keeping it there.

3.4 Summary: A Big Lesson from Analysing Small Systems

Although the C1 and C2 controllers produced relatively simple behaviours, in a simulation of a fairly simple robot in a very simple environment, we believe that these studies illustrate a significant issue which is relevant in more complex situations: our evolved robot controllers are *non-representational*.

Almost all research in machine vision, and a substantial majority of work in analysing biological vision systems, treats vision as a process of manipulation and transformation of intermediate internal representations between sensory ‘input’ and motor ‘output’. Yet the analyses of C1 and C2 seems to leave little room for treating these controllers as representational devices. We cannot identify anything within these systems that could usefully be described as ‘representational’ for any reasonable definition of the word (i.e. a definition that distinguishes representation from mere correlation).

There seems no meaningful way of describing the robot as calculating its location in the environment and acting accordingly. Nor does its behaviour involve any formulation or execution of ‘plans’. Yet a traditional approach to the arena-centering problem may well have involved using vision to determine position in the environment (i.e. location and orientation relative to the center) and then reasoning about the current position and the desired position in order to formulate a plan which is then ‘executed’ to achieve the goal. Such a traditional approach would indeed require representations, but the need for representations is in the approach, not the problem.

Other authors have questioned the need for traditional notions of representation and reasoning in artificial intelligence (e.g. [8, 7]). The evolutionary approach we employ entails that the need for representations (in the sense of maintaining internal world models) becomes an empirical question rather than a dogmatic axiom. Rather than build representations in from the start, we can analyse

the networks to determine whether the networks are representational in this sense or not. In all of our work to date, we have failed to identify any processing of such representations in the evolved controller networks.

4 Current Research

4.1 Evolving Controllers for Real Robots

Following on from the simulation studies discussed above, we and our colleagues at Sussex have demonstrated that these techniques can successfully be transferred to work with real physical robot platforms. Here we briefly review results from our work with real robots. For a review of related research by others, see [28].

4.1.1 Transferring from Simulation to Reality

Jakobi [26, 27] developed a simulator for a proprietary small differentially driven mobile robot which has eight infra-red emitter-detector proximity sensors. The detector elements have some sensitivity to visible wavelengths, and the emitters can be disabled to allow the sensors to operate as photoreceptors. The simulator allows for construction of relative complex environments with a movable light-source and a variable number of movable obstacles with either planar or cylindrical surfaces. By carefully matching the simulator to the details of the real robot and the characteristics of the space of possible environments the robot could encounter, it was possible to evolve controllers entirely in simulation which, when transferred to the real robot, produced similar behaviours. The behaviours included obstacle avoidance and phototaxis.

4.1.2 Evolving Controller Hardware Circuits

Thompson [32] has demonstrated the successful evolution of reconfigurable electronic hardware for robot control. Rather than evolve designs for parallel neural networks which are simulated

on serial computers, Thompson’s technique evolves real physical semiconductor circuits which are evaluated *in situ* in the real robot. Although Thompson’s work to date has concentrated on sonar-guided robots, there is no *a priori* reason why his techniques could not be applied to develop semiconductor circuits for controlling visually guided robots.

4.1.3 Working with Real Vision

We [22] developed a specialised ‘gantry’ robot which included a mobile video-camera head with one rotational and three translational degrees of freedom. The video camera supplied 50Hz input of 64×64 grey-level images, which were used for real-time guidance as the head moved around simple visual environments. Neural network controllers were evolved to produce a number of visually-guided behaviours. Although the camera image contains 4096 pixels, the networks received input from only a small number of “virtual photoreceptors”: the input to a visual ‘neuron’ was an estimate of the average intensity of pixels within a circular “receptive field” on the image. The position and radius of each input unit’s receptive field was genetically specified. Thus each input unit could be thought of as receiving stimulation from a single photoreceptor with a genetically specified orientation and angle of acceptance.

Evolved behaviours included tracking of a (slowly) moving target, and a discrimination task. The discrimination task involved the controller moving the camera-head towards a triangular target while avoiding moving near to a rectangular target. The final evolved controller described in [22] successfully performed this task using the equivalent of just two individual photoreceptors.

4.2 Artificial Evolution in Biological Modelling

So far, we have concentrated on describing our techniques as applied to the engineering problems of designing sensory-motor controllers. One of the motivations for developing these techniques is that artificial evolution can be used to partially automate the design process. Traditional manual engineering design techniques rely on an element of skilled creativity. The designer needs to apply

skill and creativity in order to produce a design which meets the constraints of the ‘design brief’. In the case of designing a robot controller, the design brief is likely to be more than just a specification of a desired behaviour: there will also be constraints regarding weight, power consumption, cost, and so on.

When a biologist proposes a model for the neuronal mechanisms underlying the generation of behaviour in an animal, a similar skilled creativity is being applied. Here, in addition to the requirement that the neuronal system can generate the observable behaviour, additional constraints involve ensuring that the model fit the available experimental data. In most cases of interest, there is insufficient data, and so some creative ‘inspired guessing’ is required. The model can direct future empirical work, insofar as those aspects of the model based on ‘guesses’ may imply certain predictions which can be verified or falsified by appropriate experiments.

Given that our techniques have successfully replaced significant proportions of the skilled creativity that is needed in engineering design, we are currently in the process of exploring the possibilities of applying artificial evolution to problems in neuroethology, where again we hope that the reliance on skilled creativity can be reduced. To do this, the results of the evolutionary simulation is a neural network which produces a behaviour in a simulated creature. If the simulation is suitably based on a real creature and the final evolved simulated behaviour closely matches the behaviour of that creature, then the network could be taken as a model for the neural mechanisms underlying the generation of the behaviour in the real animal. In order to do this, it will be necessary to ensure that only biologically plausible structures can be evolved (e.g. by imposing severe fitness penalties on any individuals whose network does not satisfy the constraints imposed by experimental data). In much the same way as hand-crafted models can suggest further experimental work, so evolved models may lead to new predictions which can be verified or falsified empirically.

To evaluate the feasibility of this approach, we are currently evolving neural circuits for

visually-guided pursuit behaviours, with the intention of comparing the evolved controllers to the neuroethological data on visual control of pursuit in flying insects: see [29] for further details.

5 Conclusion

In this chapter we have demonstrated that low-bandwidth visual systems can be evolved and analysed. The important achievement 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 magical 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, significantly, we don't think that it involves representation (at least, not in the conventional sense).

Our approach to developing machine vision systems is inspired by biology at a number of levels: low-bandwidth 'insect-like' vision; neural networks; and artificial evolution. Analysing the behaviour-generating mechanisms in an evolved controller presents problems similar to those in analysing (evolved) animals, and there is some similarity in the analytic techniques. Finally, in addition to using these techniques for developing *seeing machines*, we hope to apply them to help understand what happens behind *living eyes*.

Acknowledgements

Our research was supported initially by a grant from the University of Sussex Research Development Fund, and latterly by grant GR/J18125 from the UK Engineering and Physical Sciences Research Council (EPSRC).

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 agent-environment interaction. *Artificial Intelligence*, 72:173–215, 1995.
- [4] R. D. Beer. On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):471–511, 1995.
- [5] R. A. Brooks. A robust layered control system for a mobile robot. A.I. Memo 864, M.I.T. A.I. Lab, September 1985.
- [6] R. A. Brooks. Achieving artificial intelligence through building robots. A.I. Memo 899, M.I.T. A.I. Lab, May 1986.
- [7] R. A. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 139–159, San Mateo, California, 1991. Morgan Kaufmann.
- [8] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
- [9] R. A. Brooks. Artificial life and real robots. In F. J. Varela and P. Bourguine, editors, *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, pages 3–10, Cambridge MA, 1992. MIT Press Bradford Books.
- [10] J. M. Camhi. *Neuroethology: Nerve Cells and the Natural Behaviour of Animals*. Sinauer Associates Inc., Sunderland, Mass., 1984.

- [11] D. Cliff. Computational neuroethology: A provisional manifesto. In J.-A. Meyer and S. W. Wilson, editors, *From Animals to Animats: 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.
- [12] D. Cliff. Neural networks for visual tracking in an artificial fly. In F. J. Varela and P. Bourguine, 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.
- [13] D. 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.
- [14] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):71–108, 1993.
- [15] D. 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)*, pages 374–383. MIT Press Bradford Books, Cambridge, MA, 1993. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP220.
- [16] 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.
- [17] A. S. Glassner, editor. *An Introduction to Ray Tracing*. Academic Press, London, 1989.

- [18] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, 1989.
- [19] I. Harvey. Species adaptation genetic algorithms: A basis for a continuing SAGA. In F.J. Varela and P. Bourguine, 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.
- [20] I. Harvey. Evolutionary robotics and SAGA: the case for hill crawling and tournament selection. In C. Langton, editor, *Artificial Life 3 Proceedings*, pages 299–326. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVI, Addison Wesley., 1993. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP222, 1992.
- [21] 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)*, pages 364–373. M.I.T. Press Bradford Books, Cambridge MA, 1993. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP219.
- [22] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: Artificial evolution; real vision. In D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson, editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*, pages 392–401. MIT Press Bradford Books, 1994.
- [23] J. H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbour, 1975.

- [24] 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.
- [25] P. Husbands, I. Harvey, and D. Cliff. Circle in the round: State space attractors for evolved sighted robots. *Robotics and Autonomous Systems*, 15:83–106, 1995.
- [26] N. Jakobi. Evolving sensorimotor control architectures in simulation for a real robot. Master’s thesis, University of Sussex School of Cognitive and Computing Sciences, 1994. Unpublished.
- [27] N. Jakobi, P. Husbands, and I. Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón, editors, *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, pages 704–720. Springer-Verlag, 1995.
- [28] M. J. Matarić and D. Cliff. Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 1995. Special issue on “Evolutional Robotics”. Forthcoming 1995/96.
- [29] G. F. Miller and D. Cliff. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*, pages 411–420. MIT Press Bradford Books, Cambridge MA, 1994.
- [30] 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.

- [31] 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.
- [32] A. Thompson. Evolving electronic robot controllers that exploit hardware resources. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón, editors, *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, pages 640–656. Springer-Verlag, 1995.
- [33] T. van Gelder. What might cognition be if not computation? Technical Report 75, Indiana University Cognitive Sciences, 1992.




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




Figure 2: Typical behaviour of the C1 controller. Top-down view of the robot in the circular arena. 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.

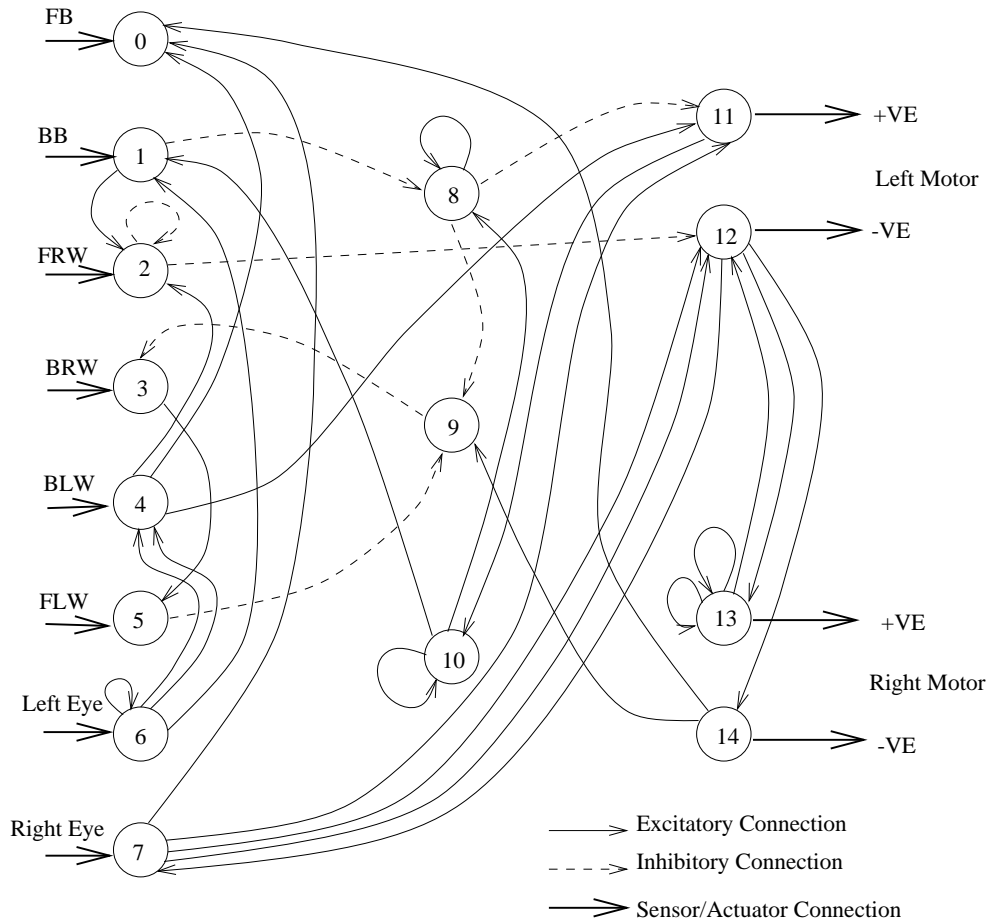


Figure 3: C1 control network. The left-hand column are units originally designated as input units: FB=Front Bumper; BB=Back Bumper; FRW=Front Right Whisker; BRW=Back Right Whisker; BLW=Back Left Whisker; FLW=Front Left Whisker. Right-hand column shows output units, which are paired and differenced to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




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.




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

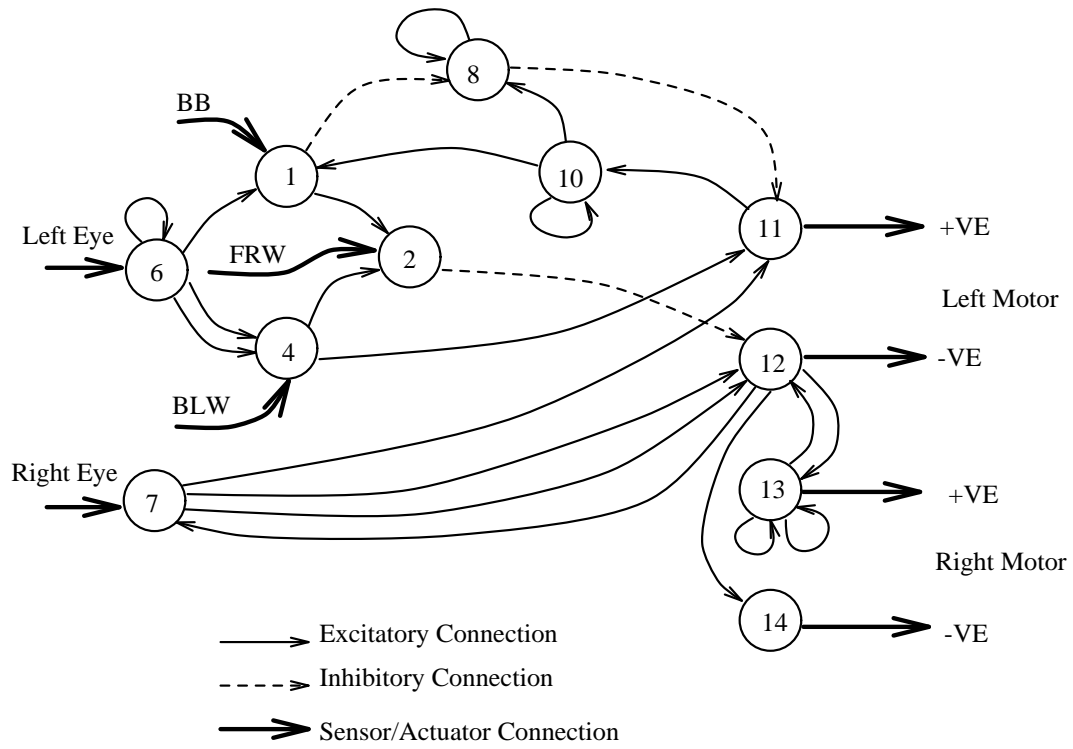


Figure 8: Final C1 network.




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




Figure 10: Primary active connections in the momentary transition between phases A1 (low-radius turn) and A2 (high-radius turn).




Figure 11: Active connections in phase A2 (high-radius turn).




Figure 12: Active connections in phase S (spin).




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).

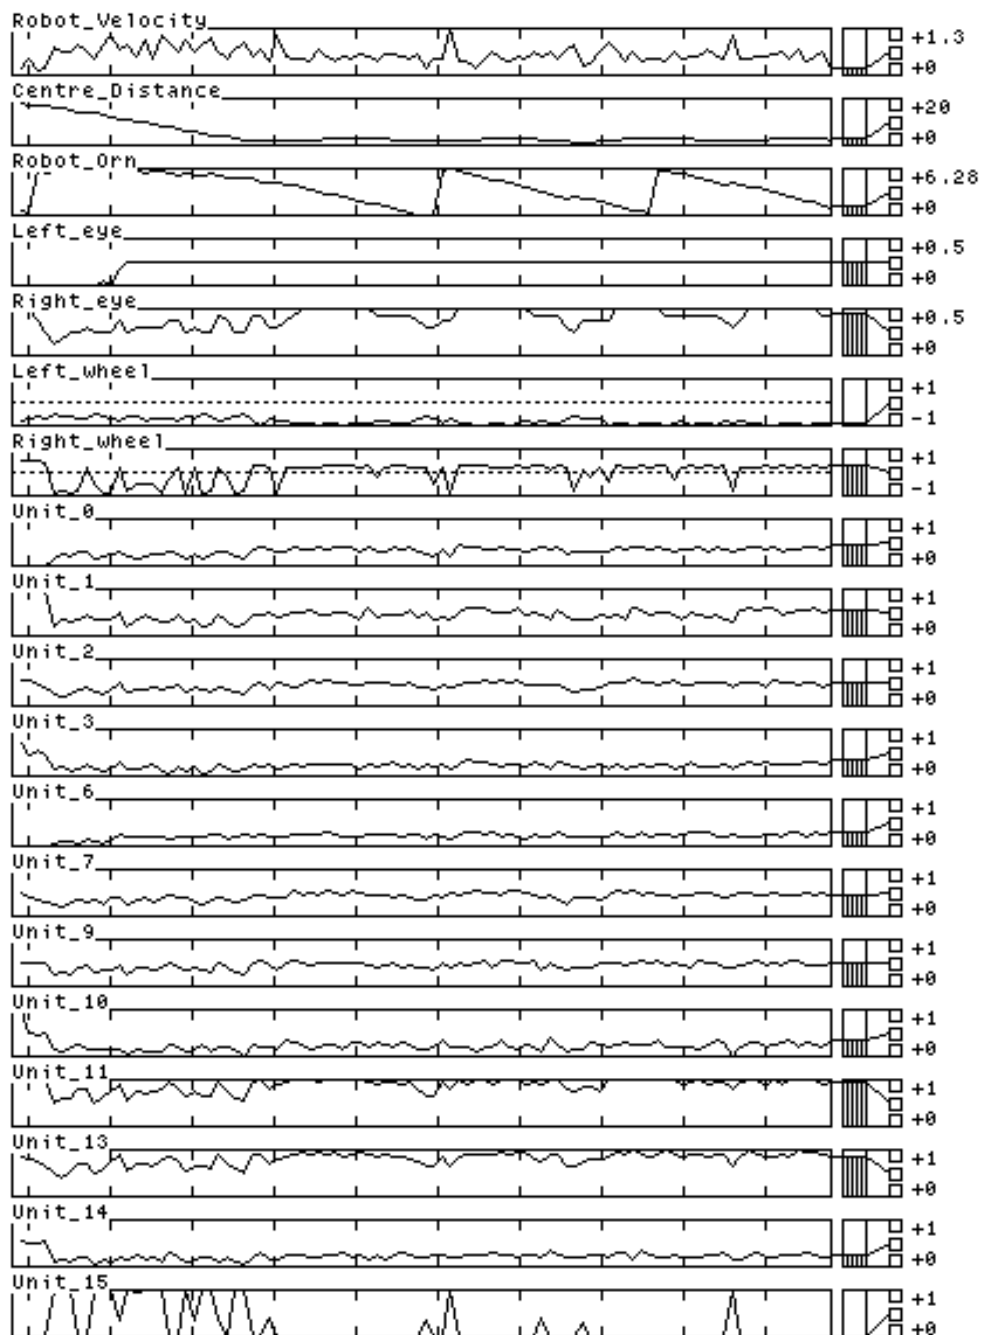


Figure 14: Record of observables and activity levels for the (with-noise) activity illustrated in Figure 13.




Figure 15: Full C2 control network. Display format as for Figure 3.




Figure 16: C2 visual guidance pathways. Note that, for the sake of clarity, the positions of the left and right motor outputs have been interchanged.




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




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.