The Artificial Evolution of Robot Control Systems

Philip Husbands and Dave Cliff and Inman Harvey
School of Cognitive and Computing Sciences
University of Sussex
Brighton, UK
Email: philh@cogs.susx.ac.uk

1 Introduction

This paper introduces the field of Evolutionary Robotics through a case study. A specialised piece of robotic equipment for evolving visually guided behaviours is described. The results of a successful experiment in the concurrent evolution of a dynamical network controller and visual sensor morphology are presented. A visuo-motor behaviour is evolved that allows the robot to distinguish between two different targets. The mechanisms underlying the behaviour are analysed and some surprisingly subtle features are uncovered. A further experiment to evolve robot controllers directly in hardware, thus allowing advantage to be taken of intrinsic silicon physics, is also briefly described. The paper closes with some discussion of future direction in the field.

2 Evolutionary Robotics

The basic notion of Evolutionary Robotics is as follows. The evolutionary process, based on a genetic algorithm [4, (Holland 1975)], involves evaluating, over many generations, whole populations of control systems specified by artificial genotypes. These are interbred using a Darwinian scheme in which the fittest individuals are most likely to produce offspring. Fitness is measured in terms of how good a robot's behaviour is according to some evaluation criterion. The work reported here forms part of a long-term study to explore the viability of such an approach in developing interesting adaptive behaviours in visually guided autonomous robots, and, through analysis, in better understanding general mechanisms underlying the generation of such behaviours. It is one of the strands of the research program of the Evolutionary and Adaptive Systems Group, School of Cognitive and Computing Sciences, University of Sussex. For further details see e.g. ([2, Cliff et al 1993]).

One of the motivations of this work is the concern that artificial nervous systems of the complexity needed to generate advanced behaviours in autonomous agents may well be beyond the capabilities of traditional engineering design practices. It is suggested ([1, Brooks 1991], [2, Cliff et al]) that such artificial nervous systems will involve many interactions, both direct and indirect via the environment, between many different sub-systems. It is precisely this kind of system that designers traditionally try and avoid, preferring clean modular structures with limited interactions between the parts. However, artificial evolution, as an automatic design methodology, may be the way forward. Time will tell.

2.1 Real World Evolution

A crucial decision in evolving robot control systems is whether or not to use simulation at the evaluation stage, transferring the end results to the real world. Since an evolutionary approach potentially requires the evaluation of populations of robots over many generations, a natural first thought is that simulations will speed up the process, making it more feasible. It has recently been shown that control systems evolved in carefully constructed simulations, with an appropriate treatment of noise, transfer extremely well to reality, generating almost identical behaviours in the real robot ([6, Jakobi et al 1995]. However, this example involved relatively simple robot-environment interaction dynamics. Once even low-bandwidth vision is used, simulations become altogether more problematic. They become difficult and time consuming to construct and computationally very intensive to run. Hence evolving visually guided robots in the real world becomes a more attractive option. The experiment described here uses a piece of robotic equipment specially designed to allow the real-world evolution of visually guided behaviours — the Sussex gantry-robot.

3 The Gantry-Robot

The gantry-robot is shown in Figure 1. The robot is cylindrical, some 150mm in diameter. It is suspended from the gantry-frame with stepper motors that allow translational movement in the X and Y directions, relative to a co-ordinate frame fixed to the gantry . The maximum X (and Y) speed is about 200mm/s. Such movements, together with appropriate rotation of the sensory apparatus, correspond to those which would be produced by left and right wheels. The visual sensory apparatus consists of a CCD camera pointing down at a mirror inclined at 45^o to the vertical. The mirror can be rotated about a vertical axis so that its orientation always corresponds to the direction the 'robot' is facing. The visual inputs undergo some transformations en route to the control system, described later. The hardware is designed so that these transformations are done completely externally to the processing of the control system.

The control system for the robot is run off-board on a fast personal computer, the 'Brain PC'. This computer receives any changes in visual input by interrupts from a second dedicated 'Vision PC'. A third (single-board) computer, the SBC, sends interrupts to the Brain PC signalling tactile inputs resulting from the robot bumping into walls or physical obstacles. The only outputs of the control system are motor signals. These values are sent, via interrupts, to the SBC, which generates the appropriate stepper motor movements on the gantry.

The Brain PC runs the top-level genetic algorithm and during an individual evaluation, it is dedicated to running a genetically specified control system for a fixed period. At intervals during an evaluation, a signal is sent from the Brain PC to the SBC requesting the current position and orientation of the robot. These are used in keeping score according to the current fitness function. The Brain PC receives signals, to be fed into the control system, representing sensory inputs from the Vision PC and the SBC. The visual signals are derived from averaging over genetically specified circular receptive patches in the camera's field of view.

This setup, with off-board computing and avoidance of tangled umbilicals, means that the apparatus can be run continuously for long periods of time – making artificial evolution feasible. A top-level program automatically evaluates, in turn, each member of a population of control systems. A new population is produced by selective interbreeding and the cycle repeats. For full technical details of the system see ([3, Harvey et al 1994]).

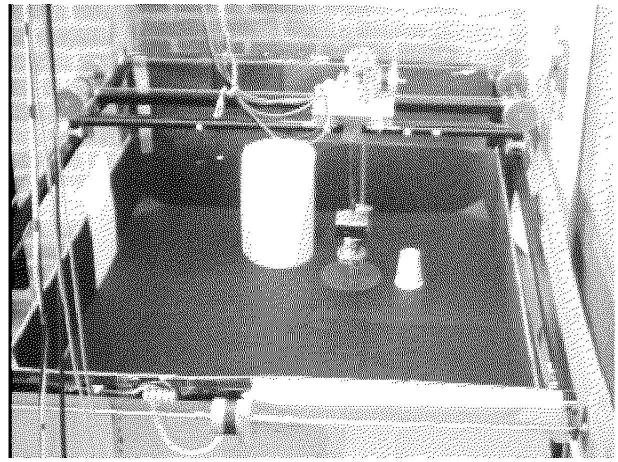

Figure 1: *The Gantry viewed from above. The horizontal girder moves along the side rails, and the robot is suspended from a platform which moves along this girder.*

4 Experimental Setup

Full details of the experimental setup for the gantry-robot can be found in ([3, Harvey et al 1994]). This paper also explains in full the genetic encodings used and the control system primitives manipulated by the GA. Experiments conducted with the gantry-robot to date have all involved relatively simple vision based navigation tasks. The experiment described below was one of a series where a converged population of robots was evolved through a series of increasingly complex behaviours.

These were based around the evolution of control architectures built from recurrent dynamic realtime networks, where the primitives were the nodes in a network, and links between them. There were no restrictions on network topologies, arbitrarily recurrent nets being allowed. When some of these nodes are connected to sensors, and some to actuators, the network acts as a control system, generating behaviours in the robot.

Rather than imposing a fixed visual sampling morphology, we believe a more powerful approach is to allow the visual morphology to evolve along with the rest of the control system. Hence we genetically specify regions of the robot's visual field to be sub-sampled, these provide the only visual inputs to the control network. It would be desirable to have many aspects of the robot's morphology under genetic control, although this is not yet technically feasible.

Starting from a converged population of robots that could move forward, but little else, the first task was to move to a large white target from random starting points and orientations. Once this was being achieved, the task was changed to approaching a small white target and evolution continued.

5 Rectangles and Triangles

The experiment then continued with a distinguish-between-two-targets task. Two white paper targets were fixed to one of the gantry walls; one was a rectangle, the other was an isosceles triangle with the same base width and height as the rectangle. The robot was started at four positions and orientations near the opposite wall such that it was not biased towards either of the two targets. The evaluation function $\mathcal{E}_3$, to be maximised, was:

$$\mathcal{E}_3 = \sum_{i=1}^{i=20} [\beta(D_{1_i} - d_{1_i}) - \sigma(D_{2_i}, d_{2_i})] \qquad (1)$$

where D_1 is the distance of target-1 (in this case the triangle) from the gantry origin; d_1 is the distance of the robot from target-1; and D_2 and d_2 are the corresponding distances for target-2 (in this case the rectangle). These are sampled at regular intervals, as before. The value of β is $(D_1 - d_1)$ unless d_1 is less than some threshold, in which case it is $3 \times (D_1 - d_1)$. The value of σ (a penalty function) is zero unless d_2 is less than the same threshold, in which case it is $I - (D_2 - d_2)$, where I is the distance between the targets; I is more than double the threshold distance. High fitnesses are achieved for approaching the triangle but ignoring the rectangle. It was hoped that this experiment might demonstrate the efficacy of concurrently evolving the visual sampling morphology along with the control networks.

After about 15 generations of a run using as an initial population the last generation of the incremental small target experiment, fit individuals emerged capable of approaching the triangle, but not the rectangle, from each of the four widely spaced starting positions and orientations. The behaviour generated by the fittest of these control systems is shown in Figure 2. When started from many different positions and orientations near the far wall, and with the targets in different positions relative to each other, this controller repeatedly exhibited very similar behaviours to those shown.

The active part of the evolved network that generated this behaviour is shown in Figure 3. The evolved visual morphology for this control system is shown inset. Only receptive fields 1 and 2 were used by the controller.

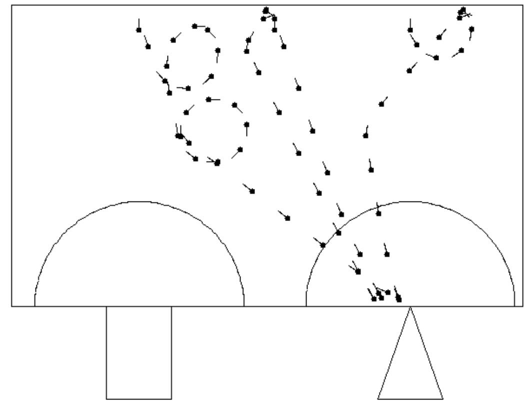

Figure 2: *Behaviour of a fit individual in the two target environment. The rectangle and triangle indicate the positions of the targets. The semi circles mark the 'penalty' (near rectangle) and 'bonus score' (near triangle) zones associated with the fitness function. In these 4 runs the robot was started directly facing each of the two target, and twice from a position midway between the two targets; once facing into the wall and once facing out.*

Detailed analyses of this evolved system can be found in ([3, Harvey et al 1994]). To crudely summarise, unless there is a *difference* in the visual inputs for receptive fields 1 and 2, the robot makes rotational movements. When there is a difference it moves in a straight line. The visual sensor layout and network dynamics have evolved such that it fixates on the sloping edge of the triangle and moves towards it.

6 Transient Behaviour

Time plots of behaviour against this difference in visual inputs consistently revealed an interesting non-reactive feature to the robot's behaviour. Figure 4 shows such a plot. The behaviour axis (Z) is discretized into simple observable motor behaviours such as straight line motion, rotating on the spot, movement in the arc of a circle and so on. The final part of the plot, a line parallel to the time axis and terminating at the point marked 'finish' at the right hand side of the cube, represents the straight line motion when the robot has fixated on the triangle edge and is moving towards it. The parallel line above this and immediately to the left represents a short lived transient behaviour which such plots revealed *always* occurred

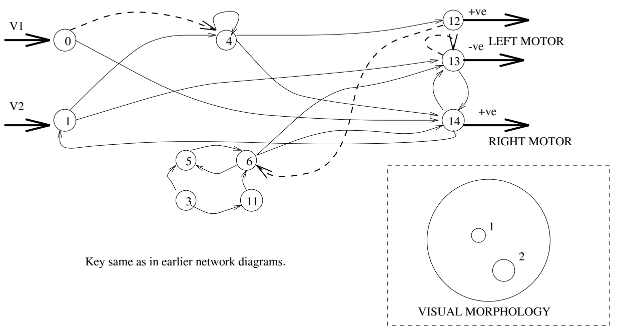

Figure 3: *Active part of the control system that generated fit behaviour for the rectangle and triangle experiment. Visual morphology shown inset.*

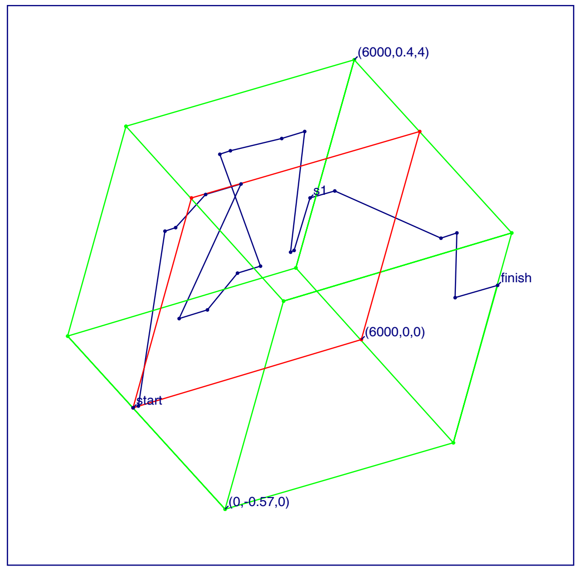

Figure 4: *Time plot of behaviour against difference in visual inputs for receptive fields 1 and 2. The time axis (X) runs left to right, the visual signal difference axis (Y) runs bottom to top on the lower face of the cube, the behaviour axis (Z) runs from lower to top face of the cube. See text for further details.*

when the visual signal difference become large. Briefly, the onset of a large difference triggers a short sharp rotational movement which has very different consequences depending on whether the robot has fixated on a vertical or sloping edge. With a vertical edge, the rotation tends to move both receptors off the target, the visual signals become very different and rotational behaviours ensue. However, with a sloping edge, the rotation is not enough to move both receptive fields off the target; the visual signal difference is still there and a straight line motion follows. This is illustrated in Figure 5. This behaviour can be interpreted as a kind of 'checking' of edge orientation.

7 Evolving Hardware

In an experiment described fully in ([8, Thompson 1995], [9, Thompson et al 1996]), a standard electronic architecture, with some of the dynamical constraints used to make conventional design tractable removed, was subjected to intrinsic hardware evolution. The result was the first evolved hardware control system for a real robot. The evolved circuit was the on-board controller for a wheeled mobile robot using a pair of time-of-flight sonars as its only sensors. The task was to avoid walls in an empty rectangular arena (i.e. move to the centre and stay there). The starting point for the evolved controller was a hardware implementation of a finite state machine (FSM). However, the constraint of synchronisation of all signals was relaxed and placed under evolutionary control. The result was a machine of fundamentally different nature to a FSM. As well as the contents of the state transition table, the global clock frequency was placed under genetic control, as was the choice of whether each signal is synchronised by the clock or whether it is free-running is also genetically determined. This resulted in the evolution of extremely minimal controllers (32 bits in RAM plus two flip-flops) exhibiting rich dynamics.

8 Conclusions

This paper has described an experiment in real-world evolutionary robotics and has shown that the evolved control networks and visual morphology for a robot engaged in a simple target distinguishing behaviour display subtle dynamics resulting in non-reactive behaviour. An experiment in hardware evolution has also been outlined.

These are just two of many experiments in evolutionary robotics that have been performed in the past few years, see ([5, Husbands et al], [7, Nolfi et al]) for a range of examples. So what are the main challenges? Certainly the genotype to phenotype (artificial DNA to robot) mapping is a huge and difficult area that has an enormous bearing on whether or not search spaces amenable to evolution are created. The design of evaluation functions and schemes, both explicit and implicit, will become more problematic as more sophisticated behaviours are sought after. The fundamental nature of the evolved networks is of course highly pertinent, as is the need to concurrently evolve

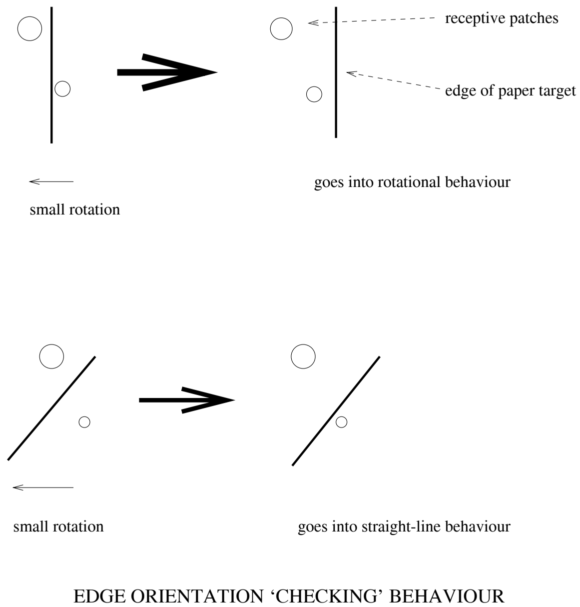

Figure 5: *The top part of the figure illustrates the outcome of the transient ‘checking’ behaviour when the receptive fields straddle a vertical edge, and the bottom part shows the same when they straddle a sloping edge.*

as much robot morphology as possible. As well as all these issues there is the potential problem of the sheer amount of time needed to evaluate more sophisticated behaviours. So where does this leave us? In my view, near the start of a difficult but exciting and potentially revolutionary endeavour.

References

[1] Brooks R 1991 Intelligence without Representation *Artificial Intelligence* **47** 139–159

[2] Cliff D and Harvey I and Husbands P 1993 Explorations in Evolutionary Robotics *Adaptive Behavior* **2(1)** 73–110

[3] Harvey I and Husbands P and Cliff D 1994 Seeing The Light: Artificial Evolution, Real Vision *From Animals to Animats 3, Proc. of 3rd Intl. Conf. on Simulation of Adaptive Behavior* eds. D Cliff and P Husbands and J-A Meyer and S Wilson (Cambridge, Mass: MIT Press/Bradford Books) pp 392–401

[4] Holland J 1975 *Adaptation in Natural and Artificial Systems* (Ann Arbor: University of Michigan Press)

[5] Husbands P and Harvey I and Cliff D and Miller G 1994 The Use of Genetic Algorithms for the Development of Sensorimotor Control Systems *Proceedings of From Perception to Action Conference* eds. P Gaussier and J-D Nicoud (IEEE Computer Society Press) pp 110–121

[6] Jakobi N and Husbands P and Harvey I 1995 Noise and the Reality Gap: The use of simulation in evolutionary robotics *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life* eds. F Moran and A Moreno and J J Merelo and P Chacon Lecture Notes in Artificial Intelligence **929** (Berlin: Springer-Verlag) pp 704–720

[7] Nolfi S and Floreano D and Miglino O and Mondada F 1994 How to Evolve Autonomous Robots: Different Approaches in Evolutionary Robotics *Artificial Life IV* eds. R. Brooks and P. Maes (MIT Press/Bradford Books) pp 190–197

[8] Thompson A 1995 Evolving Electronic Robot Controllers that Exploit Hardware Resources *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life* eds. F Moran and A Moreno and J J Merelo and P Chacon Lecture Notes in Artificial Intelligence **929** (Berlin: Springer-Verlag) pp 640–656

[9] Thompson A, Harvey I and Husbands P 1996 Unconstrained Evolution and Hard Consequences *Towards Evolvable Hardware* eds. E. Sanchez and M. Tomassini Lecture Notes in Computer Science (Berlin: Springer-Verlag) (in press)