Seeing the Light: Artificial Evolution, Real Vision

I. Harvey, P. Husbands, D. Cliff

CSRP 317, April 1994

Cognitive Science Research Paper

Serial No. CSRP 317

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

A version of this paper will appear in:
From Animals to Animats 3,
Proceedings of the 3rd Intl. Conf.

Seeing The Light: Artificial Evolution, Real Vision

Inman Harvey and Phil Husbands and Dave Cliff
School of Cognitive and Computing Sciences
University of Sussex, Brighton BN1 9QH, UK
inmanh or philh or davec@cogs.susx.ac.uk

Abstract

This paper describes results from a specialised piece of visuo-robotic equipment which allows the artificial evolution of control systems for visually guided autonomous agents acting in the real world. Preliminary experiments with the equipment are described in which dynamical recurrent networks and visual sampling morphologies are concurrently evolved to allow agents to robustly perform simple visually guided tasks. Some of these control systems are shown to exhibit a surprising degree of adaptiveness when tested against generalised versions of the task for which they were evolved.

1 Introduction

In previous papers (see e.g. [1]) we have discussed our reasons for adopting an evolutionary methodology for the design of control systems for mobile robots using low-bandwidth vision for simple navigational tasks. We also discussed what class of control systems are appropriate for evolutionary development, proposing dynamic recurrent real-time (artificial) neural networks as one strong contender.

The evolutionary process, based on a genetic algorithm [3], 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 an agent'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.

In this paper we present results from experiments in which visually guided behaviours are artificially evolved in the real world. As far as we know, this is the first time this has been achieved.

2 From Simulation to Reality

The experiments described in earlier papers [1] used simulations of a round two-wheeled mobile robot with touch

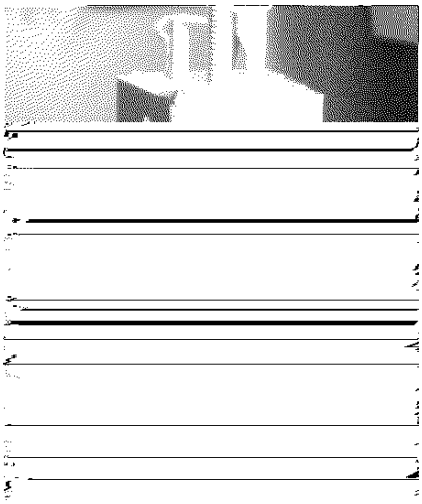

Figure 2: *The gantry-robot. The camera inside the top box points down at the inclined mirror, which can be turned by the stepper-motor beneath. The lower plastic disk is suspended from a joystick, to detect collisions with obstacles.*

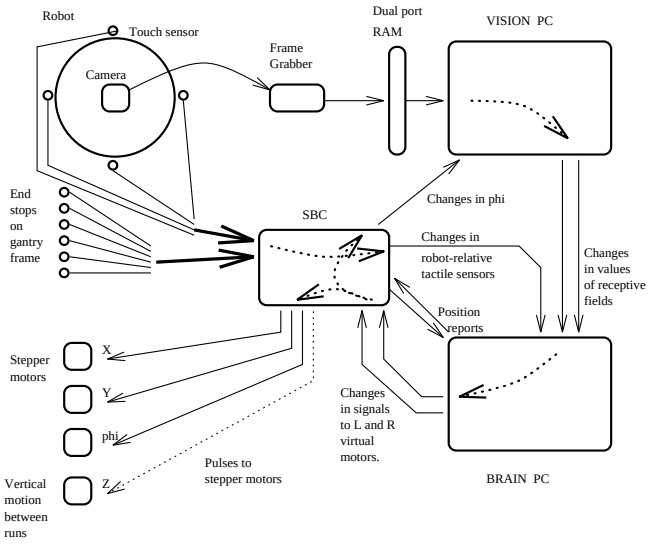

Figure 3: *The different rôles of the Vision computer, the Brain computer and the SBC.*

evaluates, in turn, each member of a population of control systems. A new population is produced by selective interbreeding and the cycle repeats.

3.2 The Vision System

Continuous visual data is derived from the output of a small monochrome CCD camera. With a wide-angle (about 40^o) fixed-focus lens about 6mm in diameter, this is housed in a box facing vertically downwards onto the angled mirror of the robot. The CCD produces composite video output of some 1 volt peak to peak, with a video bandwidth of 4MHz. A purpose-built Frame-Grabber transfers a 64×64 image at 50Hz into a high-speed 2K CMOS dual-port RAM, completely independently and asynchronously relative to any processing of the image by the Vision PC.

We advocate an incremental evolutionary approach, progressing from the simple to the complex. In keeping with this philosophy, current experiments use very low bandwidth vision. This implies sub-sampling the image produced by the camera. Rather than imposing a fixed way of sampling the image, we allow this to evolve along with the neural networks. This is achieved by genetically specifying the size and position of visual receptive fields. These are circular patches within the visual field of the camera (see Figure 4). Up to 256 such receptive fields can be specified with, to 8-bit accuracy: the diameter of the field; and the polar coordinates of the centre of the field relative to the centre of the camera's field of view. The angle of acceptance of the CCD camera (via the mirror) is about 60^o; the maximum angle of acceptance of a receptive field is about 16^o, and its maximum angle of eccentricity off the cameras visual axis is about 22^o.

To calculate the signal from such a field, the average is taken of 25 pixels in the camera image scattered across the appropriate area. In this way a value (4 bits) can be calculated for each receptive field at least as fast as the camera image is updated. The only visual inputs available to the genetically designed robot control system are such scalar values.

The Vision PC is dedicated solely to processing the camera output to calculate the visual signals from the receptive fields. At the beginning of a set of trials for a particular robot, the genetic specification for the visual morphology (positions and sizes of receptive fields) is passed to this PC. During each trial, whenever the orientation of the robot changes (the full circle is discretized into 96 orientations) a single byte is sent to the Vision PC from the SBC specifying the new orientation. Whenever the visual input to any of the receptive fields changes in value (scaled in the range 0 to 15) then the details of such a change are sent as single-byte interrupts to the Brain PC.

3.3 The Brain PC

This is a 66MHz 486 PC which has two separate groups of tasks to do at different times. Firstly, the Genetic Algorithm (GA) code is run on this machine. Reproduction, crossover and mutation are performed here in between generations, and at the start of a set of trials for each robot architecture the specification of the visual morphology is transmitted to the Vision PC. As with most GAs, however, the amount of time spent running

the grouf of tasks.

During ny ndrivdural to

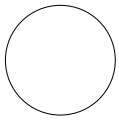

5 Visual Limitations

The visual inputs are currently subject to various limitations which are worth noting. Firstly, the lower part of the robot body is supported from the upper half with two thin vertical bolts, which come into the field of view when the mirror is facing towards them. These appear as dark bars 2 to 3 pixels wide on the CCD image, and affect the values of any receptive fields sampling from this area. In principle this could directly provide visual information for two fixed directions for the robot to 'face'. In addition, these bars tend to occlude any distant target used in navigation trials. For our early crude experiments this may not be too significant, but it certainly will matter when finer resolution is needed, and these bars produce greater effects than background noise levels. In future work we intend to fit a new head on the gantry which overcomes this problem.

Secondly, the fact that the mirror turns in discrete jumps, of 3.75^o at the moment, means that either the angles of acceptance of the receptive fields, or alternatively the horizontal angle subtended by any significant visual features, should be somewhat greater than 3.75^o. This could be overcome with a finer resolution motor.

Thirdly, the visual inputs are naturally noisy (see section 6.2). The natural variation in daylight, as day progresses into night, causes particular problems. When the gantry was exposed to such variation, it was discovered that evolved systems that worked well in the daytime did not work well under artificial light alone at night-time, and vice versa. Our individual robot systems were evaluated over a period of perhaps 3 minutes only, and hence it is no surprise that robustness against such longterm variations was not achieved. Since the recognition of this problem the gantry has been largely shielded against daylight variations. We intend soon to deliberately vary lighting conditions *within* each robot trial, to try to achieve robustness against such variations.

6 Preliminary Experiments

The following sections describe some initial simple experiments we have carried out, mainly to ascertain how well our methods cope with the move from simulations to the real world. We have begun by exploring primitive visually guided behaviours in static environments, concentrating on target approaching. However, as we shall see, some of the evolved control systems showed surprising degrees of adaptiveness when tested on more general versions of the task they were evolved for.

6.1 Networks and Genotypes

In all of the experiments reported here we used the same networks and genetic encoding schemes as in our earlier simulation work (for full details see [1]). This was mainly because we have a detailed understanding of their properties and wanted to see how well they transferred to real world tasks. However, they are the simplest, and we believe least powerful, of the classes of networks and genetic encodings we advocate, and we are currently exploring more sophisticated methods. Briefly, the evolutionary algorithms search concurrently for a network architecture and visual morphology capable of generating behaviours resulting in a high score on an evaluation function that implicitly describes a visually guided task.
ely hievcapable btra0ithgeonary a99.349(e)-1ctderstandi

6.2 Experimental Details

In each of the experiments a population size of 30 was used with a genetic algorithm employing a linear rank-based selection method, ensuring the

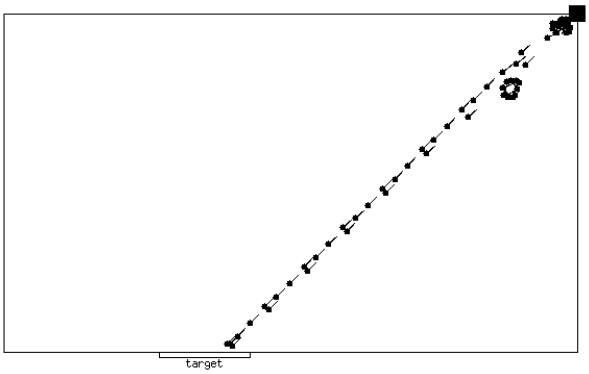

Figure 7: *Behaviour of the best of a later generation evolved under 2nd evaluation function. Format as in previous Figure.*

conditions as in the first experiment. The initial population used was the 12th generation from a run of the first experiment (i.e. we incrementally evolved on top of the existing behaviours). The behaviour of the best of this initial population is shown in Figure 6. Interestingly, this was not the best at the previous task – that individual did very poorly on the new task.

Within six generations a network architecture and visual morphology had evolved displaying the behaviour shown in Figure 7. This control system was tested from widely varying random starting positions and orientations, with the target in different places, and with smaller and different shaped targets. Its behaviour was general enough to cope with all these conditions for which it had not explicitly been evolved.

For comparison a second evolutionary run using $\mathcal{E}_2$ throughout was undertaken; this time $\mathcal{E}_1$, and the big target, were not used as a stepping stone. The run started from the same initial converged population as was used for the first task. High scoring individuals emerged after 15 generations. When tested on more general versions of the task they performed much worse than the best of the incremental run. This result is suggestive, but we do not have enough data to be able to report anything statistically significant about the advantages of doing incremental evolution at this low-level of task.

6.2.3 Moving Target

Following a moving target can be thought of as a generalised version of static target approaching. Hence we tested a number of the evolved small target locators with a white cylinder (of similar width)

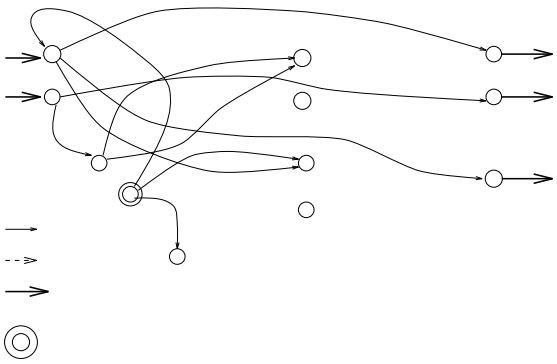

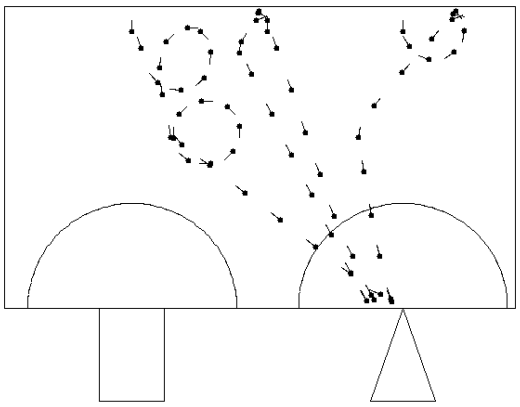

Figure 13: *Behaviour of a fit individual in the two target environment. The rectangle and triangle indicate the positions of the targets. The semi circles mark the 'penalty' (near rectangle) and 'bonus score' (near triangle) zones associated with the fitness function. In these 4 runs the robot was started 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.*

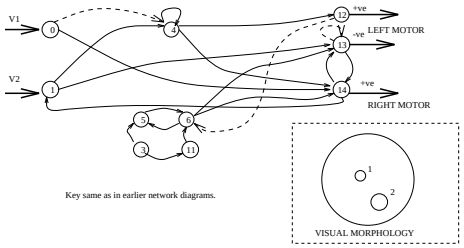

Figure 14: *Active part of the control system that generated fit behaviour for the rectangle and triangle experiment. Visual morphology shown inset.*

from many different positions and orientations near the far wall, 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 14. The evolved visual morphology for this control system is shown inset. Only receptive fields 1 and 2 were used by the controller.

Whereas the fit control systems for the previous experiments only made use of one visual receptive field at a time, this one used two simultaneously. The visual morphology/networks evolved such that robots rotated on the spot when both visual inputs were low (this is effected by the subnetwork made from nodes 3, 5, 6 and 11). When the signal from receptive field 1 (v_1) is high but that from receptive field 2 (v_2) is low, the connection from unit 0 to unit 14 generates a rotational movement.

When v

produce promising behaviour will, on recombination, almost always produce a genotype with near-average performance — i.e. useless performance. It is only once the population has largely converged — as advocated with SAGA [2] — that recombination is likely to be useful.

For this reason, from a start with a randomly generated population, the early stages will do no more than allow some early promising candidate to dominate the population. In which case we can speed up the process, and help give some desired initial direction, by ourselves observing the first random population, choosing by eye the most promising, and seeding the next generation with clones of this one. Thereafter the population settles down to its asymptotic degree of genetic convergence from above, rather than from below. For the experiments reported here, an initial randomly generated population of size 30 was judged by eye on the intuitive criterion of 'interesting' behaviour. Two members displayed forward-moving behaviour, which altered in character when the white target was within view of the visual system, and one of these two was selected. The informal criterion of 'interestingness' allowed a clear choice, whereas the 'official' evaluation function used thereafter did not give clear preferences on this initial random population, as the scores it gave there were dominated by noise. This use of different evaluations over time is completely consonant with the underlying philosophy of this approach, that of human-directed evolution of the robots.

As has already been mentioned, the successes we have had with initially converged populations are from too small a sample of experiments to have any statistical significance. It should also be noted that the genetic encoding scheme plays an important role in determining how effective crossover is in early generations.

10 Future Work

Encouraged by the initial results with the gantry apparatus we intend to start using it in more complex experiments. In these we intend to use networks with much richer intrinsic dynamics, and more sophisticated genotype to phenotype developmental processes allowing a less restricted open-ended evolutionary process. We will explore behaviours in cluttered and dynamic environments and under changing lighting conditions.

Evaluations with the gantry using a real optic array take less than one order of magnitude longer than the early simulations we did using ray-tracing in a very simple environment. But whereas ray-tracing simulations rapidly scale up in computational requirements as the environment is made more complex, with the gantry there is no such constraint.