

Evolving Visually Guided Robots

Dave Cliff^{1,2} and Philip Husbands¹ and Inman Harvey¹

¹School of Cognitive and Computing Sciences

²Neuroscience IRC, School of Biological Sciences

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

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

Abstract

We have developed a methodology grounded in two beliefs: that autonomous agents need visual processing capabilities, and that the approach of *hand-designing* control architectures for autonomous agents is likely to be superseded by methods involving the *artificial evolution* of comparable architectures.

In this paper we present results which demonstrate that neural-network control architectures can be evolved for an accurate simulation model of a visually guided robot. The simulation system involves detailed models of the physics of a real robot built at Sussex; and the simulated vision involves ray-tracing computer graphics, using models of optical systems which could readily be constructed from discrete components.

The control-network architecture is entirely under genetic control, as are parameters governing the optical system. Significantly, we demonstrate that robust visually-guided control systems evolve from evaluation functions which do not explicitly involve monitoring visual input.

The latter part of the paper discusses work now under development, which allows us to engage in long-term fundamental experiments aimed at thoroughly exploring the possibilities of concurrently evolving control networks and visual sensors for navigational tasks. This involves the construction of specialised visual-robotic equipment which eliminates the need for simulated sensing.

1 Introduction

Designing control architectures for visually guided mobile autonomous robots that exhibit adaptive behaviour is likely to be a very difficult task. So difficult, in fact, that we advocate the abandonment of approaches to the problem which involve solution by manual design.

In place of design-by-hand, we propose using evolutionary techniques. Focus then shifts from specifying *how* the robot is to generate adaptive behaviours, to specifying *what* adaptive behaviours are generated. By

2 Background

In another paper [7], we have presented arguments supporting the notion that an evolutionary approach to the design of robot control systems can be expected to supersede design by hand. In that paper we also explored issues arising from the adoption of an evolutionary approach and gave results of preliminary simulation experiments in evolving control architectures for simple robots equipped with a few touch-sensors: four ‘whiskers’ and two ‘bumpers’. For reasons explained in [7], the control architectures were based on a particular kind of ‘neural’ network, and central to the evolutionary mechanisms is the notion of a gradual incremental development, building on already existing capabilities.

The results of the experiments with purely tactile sensors are highly encouraging: for certain types of evaluation function, the robot population can evolve to the point where genuinely useful behaviours emerge. Nevertheless, the proximal nature (and low dimensionality) of the robot’s sensors forever constrain it as unable to go beyond primitive ‘bumping and feeling’ strategies in navigating around its environment. For more sophisticated navigation strategies, based on distal information, the addition of visual sensing capabilities is required. Briefly, the rationale for adding vision is that it allows for much more sophisticated behaviour patterns (e.g. location recognition in navigation). The remainder of this paper discusses our experiences in adding visual processing capabilities to the simulated robot.

3 And hen here Was Light

Rather than imposing on the robot some visual sensors with fixed properties, it seemed much more sensible, and in keeping with our incremental evolutionary approach, to investigate the concurrent evolution of visual sensors and control networks. In essence we have started with simple very low resolution devices coupled to small networks, and will work towards higher resolution devices made useful by more complex networks generating more sophisticated behaviours. Major factors affecting how this occurs are under evolutionary control.

3.1 Preliminaries

Because the simulated robot is based on a physical robot under development, it is necessary to sufficiently constrain the visual processing capabilities available under evolutionary control, so that whatever designs evolved are (at least in principle) capable of being built using available hardware. In essence, this meant opting for very low visual resolution. The total number of pixels had to be at least two or three orders of magnitude

lower than that used in conventional computational vision research.¹

A cursory survey of some biology literature indicated that, for creatures such as insects or other arthropods which have very few photoreceptor units, the photoreceptors often have large angles of acceptance,² and are distributed around the body so as to sample a wide visual field. These simple photoreceptor units are perhaps best not thought of as pixels in an image (or ‘tiles’ in a retinal ‘mosaic’): a more appropriate approach is to consider the photoreceptors as simple local brightness detectors. For example, if the portion of the optic array directly above an animal suddenly goes dark while the rest of the optic array remains constant, it seems likely that something is about to drop on the animal from above, and rapid evasive action is probably a sensible adaptive behaviour in such situations. Of course, the animal doesn’t have to construct any internal representations or reason about the cause of the darkness; it just has to do something useful.

For this reason, our work to date on evolving visually guided robots has concentrated on ultra-low-resolution vision, close in spirit to Braitenberg’s *Vehicles* [1]. The simulated robot has been given a few photoreceptor units, which could realistically be added to the physical robot. This could be done using discrete components (e.g. photodiodes, phototransistors, or LDR’s) with individual lenses, thereby creating an electronic compound eye, cf. [4]; or by using conventional CCD cameras but impairing their optics by mounting sand-blasted glass screens in front of the lens so as to generate input images with focus-independent blur, prior to some coarse sub-sampling scheme.

The simulated robot was equipped with vision by embedding it within the SYCo vision simulator system described in [2]. The SYCo simulator was developed for studying issues in visual processing for control of an airborne insect, but only minor alterations were required: the ‘altitude’ was clamped at a constant value, because the robot is a wheeled vehicle travelling on a flat floor;

(visual) signal being sampled.

To limit the effects of aliasing, the SYCo code was configured to determine each photoreceptor's activity by averaging the readings from several rays per simulated receptor, distributed across that receptor's visual field. This provides more accurate estimates of image brightness in the receptor's field of view. However, it is important to keep the number of rays per receptor relatively low. This is for two reasons: one pragmatic, the other theoretical. First, ray-tracing is a computationally expensive process, so using fewer rays per receptor saves processing time. Second, real vision is not an arbitrary-precision process. In vision, noise is inescapable, and noise effectively reduces a continuum of brightness levels to a small number of discrete ranges (e.g. [9]). By limiting the number of rays per receptor, the precision of the brightness-value estimate is correspondingly reduced. The simulated robot must be able to cope with noisy limited precision perception, because that is all the real world has to offer.

3.2 Particulars

3.2.1 Vision

In keeping with the minimal incremental approach advocated in [7], we have commenced our studies by exploring the effects of adding just two photoreceptors to the sensor suite (bumpers and whiskers) described above. Taking a cue from biological vision, the sensors are situated in positions which are bilaterally symmetric about the robot's longitudinal midline.

Having only two receptors introduces manifest limitations on the classes of behaviours that can be expected to evolve in the robot. Assuming that the receptors




Figure 3: Illustration of the ray-tracing system. The left-hand figure shows the robot's position and orientation within the cylinder, which has black walls and white floor and ceiling. At upper right is a pair of relatively high-resolution images, traced from the robot's position inside the cylinder. The lower-right figure shows the two 4×4 images traced prior to averaging, with $\alpha = 1.571$ and $\beta = 0.628$. The final two photoreceptor brightness levels are derived by averaging the 4×4 images.

illustrates output from the ray-tracing system in this environment; Figure 4 illustrates the effects of varying α and β .

3.2.2 Physics

The simulation involves a realistic physics for determining the effects of the robot moving across the floor and colliding with the walls. As described in more detail in [7], the simulated robot is cylindrical in shape with two wheels towards the front and a single trailing rear castor. The wheels have independent drives allowing turning on the spot and fairly unrestricted movement across a flat floor. Outputs from the robot's control networks feed direct to the wheel drives. Depending on the strength and sign of their signal, the wheels rotate at one of five rates: full speed forward; full speed backward; half speed forward; half speed backward; and stationary. The continuous movement of the robot is simulated by polling the network outputs at an appropriate rate. At each step of the simulation the next position and orientation of the robot is calculated using the appropriate kinematic equations (with a suitable amount of noise added). Collisions are handled as accurately as possible, based on observations of the physical system. Briefly, if the robot collides with a high velocity normal to the surface it undergoes a noisy reflection with a rotation determined by its original direction of motion; if it collides at low speed its behaviour depends on the angle of incidence – it may rotate until normal to the obstacle or it may skid around until it is parallel.

4 Experiments

Results from earlier experiments discussed in [7] demonstrated that our methods could be used to evolve robots which could engage in primitive tactile-based navigation

ulation can be rated. We have found three evaluation functions useful:

$$\begin{aligned}\mathcal{E}_1 &= \sum_{\forall t} \mathcal{D}(t) \\ \mathcal{E}_2 &= \left(\sum_{\forall t} \mathcal{D}(t) \right) \cdot \left(\sum_{\forall t} \mathcal{B}(t) \right) \\ \mathcal{E}_3 &= \left(\sum_{\forall t} \mathcal{D}(t) \right) \cdot \left(\sum_{\forall t} \mathcal{G}(t) \right)\end{aligned}$$

where:

$\forall t$ denotes all time, i.e. the lifetime of the individual;

$\mathcal{D}(t)$

approach.

Within this framework, the aims of our first set of simulation experiments was to try and evolve coupled networks and visual sensors capable of generating interesting behaviours.

5 Results

All of the following results were achieved with population size 40, a crossover probability of 1 and a mutation rate of the order of one bit per genotype. The visual sensor and network chromosomes are crossed and mutated separately, but both contribute to the resultant phenotype: the sighted robot. Rank based selection was used with the fittest individual being twice as likely to breed as the median individual. So far the experiments have only been run for a relatively small number of generations, given the expense of the ray tracing and the fact that each individual is evaluated multiple times, as described below.


Each individual in each generation was run four times from random starting positions and orientations. Each run was for a fixed number of time steps. The fitness of the individual was taken as the *worst* score from their four runs. This strategy is used to encourage robustness, remembering that there is noise at all levels in the system. A fine-time-slice simulation was used as a close approximation to a continuous system. At each time step the sensor readings are fed into the neural network. The continuous nature of the networks was simulated by running them (synchronously updating all unit inputs and outputs) for a number of iterations (about 100, but with a variance to counter distorting periodic effects) and then converting the outputs to motor signals. The new position of the robot is then calculated, using the model physics described in Section 3.2.2.

By using suitably fine time-slices, this mode of simulation is more than adequate; although we are working on more subtle techniques to allow fully asynchronous event-based simulations.

The first set of experiments used evaluation function \mathcal{E}_1 , a simple integration of distance moved. Comparisons were made between sighted and blind robots (which used only the six touch sensors). Both did well although the evolved behaviours were quite different in the two cases. The blind robots evolved to make looping elliptical movements like that shown in Figure 5.

The strategy seems sensible as it tends to keep the robot away from the walls. The networks quickly evolved to the state where sensory inputs triggered changes in directions which sped the robot away from the wall. See Figure 10, later, for an example of such behaviour.

The sighted robots did better, tending to keep moving by staying away from the walls using visually guided behaviours like those shown in Figure 8, described in more






Figure 7: Fittest behaviour of sighted robot in very early generations under evaluation function \mathcal{E}_3 .





Figure 8: Later evolved behaviours under \mathcal{E}_3 .

on the Gaussian function.

The graphs in Figure 8 show the visual signals and the motor signals (with noise removed for easier interpretation) plotted against time.



Figure

The ‘Toytown’ environment has some similarities with the ‘Tinytown’ environment at Rochester [8]. However the latter has a camera pointing down that can move only in two dimensions, giving the equivalent of ‘low-flying aerial photographs’. In contrast, the toytown robot has a (virtual) rotational degree of freedom, and can travel in and amongst the objects of a 3-D world, with a horizontal field of view manipulable between 0° to 360°.

7 Summary and Conclusions

As further support of our claims in [7], we have presented early results from experiments in evolving network processing architectures for mobile robots. Using networks of relatively constrained processing units (‘neurons’), and simple evaluation functions, we have been able to evolve visual control architectures, even when the evaluation function is not defined in terms of monitoring visual inputs.

The results have demonstrated the feasibility of the approach, but the computational costs of simulating vision have lead us to develop a method which allows for a mix of ‘real’ vision and evolutionary methods, using readily available hardware. The ‘toytown’ project is at an early stage, but our current results are sufficiently promising that we are confident of future success. Watch this space.

Acknowledgements

Many thanks to Linda Thompson for help in the preparation of this paper. Inman Harvey is supported by a SERC grant.

References

- [1] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. M.I.T. Press — Bradford Books, Cambridge MA, 1984.
- [2] D. T. Cliff. The computational hoverfly; a study in computational neuroethology. 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 87–96, Cambridge MA, 1991. M.I.T. Press — Bradford Books.
- [3] D. T. Cliff and S. Bullock. Adding ‘foveal vision’ to Wilson’s animat, 1992. Forthcoming.
- [4] 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.
- [5] A. S. Glassner, editor. *An Introduction to Ray Tracing*. Academic Press, London, 1989.
- [6] I. Harvey. Species adaptation genetic algorithms: A basis for a continuing SAGA. In F.J. Varela and P. Bourgine, editors, *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, pages 346–354. M.I.T. Press — Bradford Books, Cambridge MA, 1992.
- [7] I. Harvey, P. Husbands, and D. T. Cliff. Issues in evolutionary robotics. Technical Report CSRP 219, University of Sussex School of Cognitive and Computing Sciences, 1992.
- [8] R. C. Nelson. Visual homing using an associative memory. *Biological Cybernetics*, 65:281–291, 1991.
- [9] M. V. Srinivasan, S. B. Laughlin, and A. Dubs. Pre-