Evolution versus Design: Controlling Autonomous Robots

Philip Husbands and Inman Harvey
School of Cognitive and Computing Sciences
University of Sussex
Brighton BN1 9QH, England
email: philh@cogs.susx.ac.uk inmanh@cogs.susx.ac.uk

Abstract

This paper sets out and justifies a methodology for the development of the control systems, or 'cognitive architectures', of autonomous mobile robots. It will be argued that the design by hand of such control systems becomes prohibitively difficult as complexity increases.

The alternative approach of artificial evolution is presented. It is argued that the most useful basic building blocks for an evolved cognitive architecture are adaptive noise tolerant neural networks rather than programs. These networks may be recurrent, and should operate in real time. Evolution should be incremental, using an extended and modified version of genetic algorithms.

Time constraints mean that architecture evaluations must be largely done in simulation. Results from a simulation are presented. The pitfalls of simulations compared with reality is discussed, together with the importance of incorporating noise.

1 Introduction

This paper sets out and justifies a methodology for the development of the control systems, or 'cognitive architectures', of autonomous mobile robots intended for use in unstructured or dynamic environments. Such robots will require active perception and will be behaviour-based. Although behaviour-based approaches to robot control appear to be far more promising than traditional model-based functional decomposition methods, it will be argued that the design of such control systems is still prohibitively difficult. A methodology based on the alternative approach of artificial evolution is presented. The advantages of such a scheme over design by hand are described.

The methodology is further illuminated by describing its application to the development of the cognitive architecture of a mobile autonomous robot engaged in a series of increasingly complex tasks. The robot is equipped with low resolution sensing devices and is required to act in uncertain environments. It is argued, in this case, that the most useful basic building blocks for an evolved cognitive architecture are adaptive noise tolerant neural networks.

Relevant work by others will be discussed. There have recently independently been suggestions of different but related evolutionary approaches. In particular during the preparation of this paper Brooks proposed using Evolutionary Programming techniques [8]; and Beer and Gallagher reported on the use of dynamical neural networks [5]. These methods will be compared with ours.

Artificial evolution requires the evaluation of a great number of candidate control systems. Time constraints mean that many of these evaluations must be done in simulation. The requirements of such a simulation system are discussed.

2 Interesting robots are too difficult to design

Traditional approaches to the development of autonomous robot control systems have made only modest progress, with fragile and computationally very expensive methods. A large part of the blame for this can be laid at the feet of an implicit assumption of functional decomposition — in the terms of the themes of this conference, the assumption that perception, planning and action could be analysed independently of each other. This failure has led to recent work at MIT which bases robot control architectures instead around *behavioural decomposition*[6, 7]. Such work rejects the traditional AI approach which manipulates symbolic representations of the world, and places more emphasis on 'knowing how' to do things rather than 'knowing that' the world is in a given state. Viewpoints sympathetic to such an approach can be seen in, e.g., [22, 4, 10, 2].

Such a subsumption-style cognitive architecture for

0-8186-2675-5/92 $3.00 © 1992 IEEE

a robot in theory analyses independent behaviours of a robot or *animat*,[1] such that each behaviour can be 'wired in' all the way from sensor input to motor output. Simple behaviours are wired in at first, and then more complex behaviours are added as a separate layer, affecting earlier layers only by means of suppression or inhibition mechanisms.

However it is accepted that the design of robust mobile robot control systems is highly complex because of the extreme difficulty of foreseeing all possible interactions with the environment; and the interactions between separate parts of the robot itself [7, 18]. The design by hand of such a cognitive architectures inherently becomes more complex much faster than the number of layers or modules within the architecture — the complexity can scale with the number of possible interactions between modules.

One way out of this problem to try and automate the design process. A possible approach is to view the design of a control architecture as a planning problem. Traditional AI approaches to planning have been shown to be computationally infeasible when applied to such problems [9]. More recent approaches to planning [2] have dwelt on much lower-level problem solving capabilities and a complex design problem is far beyond their horizons.

The design of a behavioural layer in a subsumption architecture seems to be design by magic, by sleight of hand, by indirection; in the sense that the desired behaviour can often be described as an emergent by-product of rule-following which does not explicitly mention that behaviour [14]. Emergence is in itself nothing magic as a phenomenon, if it is considered as emergence-in-the-eyes-of-the-beholder. Something can be characterised as emergent relative to an initial given description if:

1. a system can be set up which is completely described in this initial way
2. A new description of the behaviour of the system can be made which 'is useful' or 'makes sense' to an observer, and makes use of concepts outside those originally given.

To design such emergent behaviours hence requires either (a) a computationally intractable planning problem or (b) a creative act on the part of the designer — which is to be greatly admired, though impossible to formalise. In both cases it seems likely that the limits of feasibility are currently being tested.

[1] *Animat* ... simulated animal or autonomous robot.[24]

3 Let's evolve robots instead

If, however, some objective fitness function can be derived for any given architecture, there is the possibility of automatic evolution of the architecture without explicit design. Natural evolution is the existence proof for the viability of this approach, given appropriate resources. Genetic Algorithms (GAs) [11] use ideas borrowed from evolution in order to solve problems in highly complex search spaces, and it is here suggested that GAs, suitably extended in their application, are a means of evading the problems mentioned in the previous section.

The artificial evolution approach will maintain a population of viable genotypes (chromosomes), coding for cognitive architectures, which will be inter-bred and mutated according to a selection pressure. This pressure will be controlled by a task-oriented evaluation function: the better the robot performs its task the more evolutionarily favoured is its cognitive architecture. Rather than attempting to hand design a system to perform a particular task or range of tasks well, the evolutionary approach will allow their gradual emergence.

The sleight-of-hand, or indirection, problem mentioned above, is avoided with GAs in that there is no need for any assumptions about means to achieve a particular kind of behaviour, as long as this behaviour is directly or implicitly included in the evaluation function.

Brooks' subsumption approach was mentioned above as a contrast to the dogmatic assumptions of functional decomposition implicit in much of traditional robotics. Nevertheless, it is similarly not necessary to be dogmatically committed to an exclusively behavioural decomposition. By allowing both types of decomposition, the evolutionary process will determine where in practice the balance should lie in the robots' cognitive architecture.

4 Related Work

A number of researchers have speculated on the use of evolutionary techniques for mobile robot programming [3, 23], though no practical applications have been reported. Some have shown the method to be viable for simulated robots in highly simplified simulated worlds [1], but have not had to face the exponential increase in complexity that follows from progress from toy worlds into the real world.

In December 1991 Brooks reported at ECAL-91 in Paris [8] that he was starting work on an evolutionary approach to robotics, based on Koza's Genetic Programming techniques [16]. He acknowledged the practical necessity of largely using simulations, and