A Symbolic and Behavioural Hybrid Approach to Robotic Agent Design

Taehee Kim

Korea R&D Information Center
P.O. Box 122, Yusong
Taejon, S. Korea, 305-600

Email: taehee@kordic.re.kr
Tel: +82-42-828-5062, Fax: +82-42-822-1974

May 29, 1997

Abstract

This paper proposes an approach to agent design that hybridises classical symbolic approach and the goal directed behaviour-based approach. An agent is a system that usually requires detailed low level information handling routines, at the same time high level abstract functions such as reasoning. We argue that a centralised methodology based on the classical symbolic approach is inappropriate in managing low level details. Rather, we propose a methodology that employs behaviour–based reactive low level mechanism which deals with detailed information processing. High level classical symbolic component can be kept as its minimal size and generates global job plans such as job orders. This feature would provide a relatively simple and convenient methodology in designing a complex agent. In this paper, symbolic centralised methodology to agent design is first criticised from both practical and philosophical viewpoint. Then, the hybrid approach is explained and examples are summarised. Other hybrid approaches are finally reviewed.

1 Introduction

For two decades of Artificial Intelligence (AI) research, one approach has been most dominant, the *symbolic approach*, where symbolic representation and manipulation play the central role in the system.

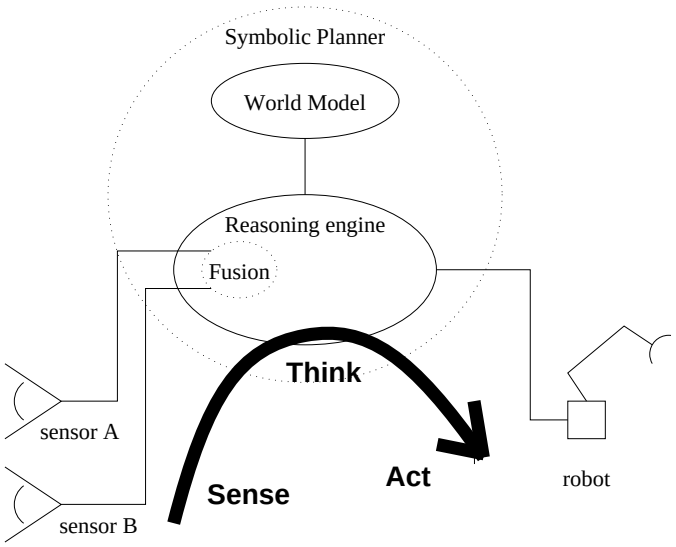

Figure 1: Traditional sense–think–act cycle of centralised symbolic planner for sensor mounted assembly robots

This derives from the view that intelligent behaviour is essentially knowledgeable behaviour. These systems tend to be centralised around information processing, with some sort of reasoning engine working with the stored centralised world knowledge. This approach has caused many robotic agent systems to employ centralised symbolic reasoning systems, particularly for those sophisticated systems handling many sensors.

However, under this structure of the centralised representation, it is difficult to avoid an ineffective information flow, known as the sense–think–act cycle. For instance, when information comes from a sensor, the central decision system has to reason to issue a motor command. If there is information from two or more sensors, the central system combines the information then makes a decision (see Figure 1). This sense–think–act cycle has inefficient serial and centralised information processing. As more sensors are employed, the system becomes big and complex, hence inefficient, or even intractable. Failure of part of the system may bring the whole system to fail since the information processing is centralised. Large amounts of knowledge cause a time consuming search problem and difficult maintenance. Modification or upgrading would demand much work throughout the whole system.

Dissatisfaction with the inefficiency and inadequacy of the centralised symbolic approach, and the problems it presents for combining multiple sensing, led Brooks to propose what he calls the **Subsumption Architecture** [1], and more generally the **Behaviour–Based Approach** [2]. Brooks argues:

> True intelligence requires a vast repertoire of background capabilities, experience and knowledge (however these terms may be defined). Such a system can not be designed and built as a single amorphous lump. It must have components. ([1], p. 5, original parenthesis)

In the behaviour–based approach, a creature is believed to be a collection of goal–directed behaviours. A more sophisticated creature is developed as more behaviours are added. Human level mental functions such as reasoning are believed to be emergent property by the operation of such accumulation of lower level behaviours and their interaction with the environment. However, we do not yet know if this would work and although if so, we do not know when. We developed a hybrid approach to robotic assembly systems. Low level behaviour–based executive agent is employed which deals with actual real world problems including various uncertainties, while high level planner reasons about assembly orders in free of uncertainty space. We claim that this should keep the overall system size minimal while allowing flexible and robust sensing and action bindings. We propose that this approach could be generalised in order to provide a general methodology in designing agents.

This paper first reviews the arguments developed in implementing intelligent systems, then explains our hybrid approach to autonomous agent design. Other hybrid approaches are finally reviewed.

2 Emergence of the Behaviour-Based Approach

In AI, it has long been debated what the right way of implementation of intelligence must be. Newell and Simon have proposed the *Physical Symbol System Hypothesis*, which asserts that intelligence can be implemented by the right kind of behaviour produced by the right kind of machinery [3]. This hypothesis supported knowledge–based systems, assuming that the right kind of machinery is a computer with symbolic manipulation capability. This approach is known as Cognitivism (symbolic approach).

In Cognitivism, in the implementation of intelligence, what we think our intelligence is plays a prime role. Churchland refers to this as Folk Psychology [4]. We empirically think that we are intelligent because we can communicate with one another and reason about sophisticated things such as a chess game, which is mainly based on language and utilising symbols. Computers, with explicit knowledge encoded by the designer, have been adopted as the right machinery to implement intelligence, since from a cognitivist's point of view, they have a computation ability similar to that of a human seen as a rational creature. Cognitivism is based on the idea that the *internal* machinery is like the *external* rationalised linguistic descriptions we

use in communicating with one another.

However, other approaches in AI, *viz.* Connectionism and the behaviour–based approach, suggest that the machinery must be something more intrinsic to the physical and biological features and less purely logical. Connectionism takes the position that intelligence needs to be implemented at the neuronal level [5], while in the behaviour–based approach, purposeful fragments of behaviours are the ingredients of the logical hardware, as in the Subsumption Architecture.

In the behaviour–based approach, a creature is believed to be a collection of goal-directed behaviours. Every behaviour is independent in that it interacts with the *environment* for its own *purpose*. Because each interacts with the environment purposefully, combining behaviours implies summing the purposefulness of every behaviour, without increasing the complexity of any particular part of the system. The accumulation of behaviours allows the system to grow incrementally, incrementally developing into a more sophisticated structure rather than being redesigned. The potential to exhibit intelligence is included implicitly in each behaviour, the relationship between the behaviours, and the purposefulness of every behaviour in the environment. Brooks emphasises the importance of actual interaction between the creature and the environment, which results in the emergence of intelligence: "Intelligence is determined by the dynamics of the interaction with the world"; "Intelligence is in the eye of the observer [2]."

From the practical viewpoint, the behaviour–based approach allows the whole robot task to be decomposed into useful task achieving behaviours, where sensing and robot action are implemented at low levels, thus the higher level of behaviours need not consider the details of the lower level sensing and action [6, 7, 8]. This approach also allows the incremental growth of the system, by appending behaviours to the current configuration in order to increase the versatility of the robot. Also, modification of the system is accomplished by replacing appropriate behaviours, without modifying the whole system. This approach leads to a robust and flexible hierarchically structured robot system.

One important question in the behaviour–based approach is whether a system built under Brooks' Subsumption Architecture can be instructed to do lots of different complex detailed tasks which require following a symbolic plan, which may require the interpretation of explicit symbolic knowledge. It is a principle of the Subsumption Architecture to avoid symbolic knowledge, so that if a system appears to be using symbolic knowledge, this will be an appearance implemented by some other mechanism.

3 Hybridisation

As far as practical assembly robots are concerned, which may be used in such factories as we currently have, parts are described in geometric terms often supported by CAD, and an ordered assembly plan is required. Highly flexible task level assembly robots require a symbolic interface, understanding task specification and other symbolic information, in other words, explicit symbolic knowledge. Would it be possible under the Subsumption Architecture to develop a simple agent to the level of an assembly robot with massive sensing capability which plans, understands human plans and geometric information? At present, it is too early to say if it is possible or not.

At Edinburgh, Malcolm suggested a way to build a task level assembly robot, using a behaviour–based approach to implement an executive agent composed of purposeful on-line Behavioural Modules (BM's) which execute details of the plan made by an off-line classical symbolic planner, a hybrid system as illustrated in the Figure 2 [9, 10]. Since assembly tasks are intimately associated with human symbolic functions, those ingredients could be better described in symbolic terms in an effective and economic way. The symbolic planner is employed as the symbolic interface, simulating human symbolic functions. Detailed sensing and action problems are managed by the low level behaviour–based executive agent composed of Behavioural Modules (BM's) dealing with the real world uncertainties and variations of the environment. As the details of sensing and action are amalgamated in the executive agent, the size of the planner is kept as small as possible as it is mainly responsible for planning of the general order of the assembly. By virtue of the smart executive agent, the planner need not consider too many details. This guarantees a comparatively small amount of world model and plan, which makes the whole system more manageable and economic [11].

Let us take an example. Suppose there is a factory worker and a manager. The factory worker is supposed to perform detailed jobs while the manager sets up general plans and directs the factory worker in a desired way. If the factory worker is smart, able to deal with small batch work under some uncertainties, the plan from the manager does not have to describe all the details. The manager can concentrate on the general plan, while the factory worker deals with the details of any practical problems.

For a traditional assembly robot, the planner, which deals with the real world uncertainties, is like an unlucky manager, having to make a complex plan considering all the associated uncertainties and details. The planner has a very simple factory worker: merely actuators and isolated information channels from the sensors. Given contemporary technologies, it is often impossible to build such a practical planner working reliably in the real world, due to the complexity. On the other hand, if the planner hires a smart and skilled agent dealing with uncertainties by itself, the complexity

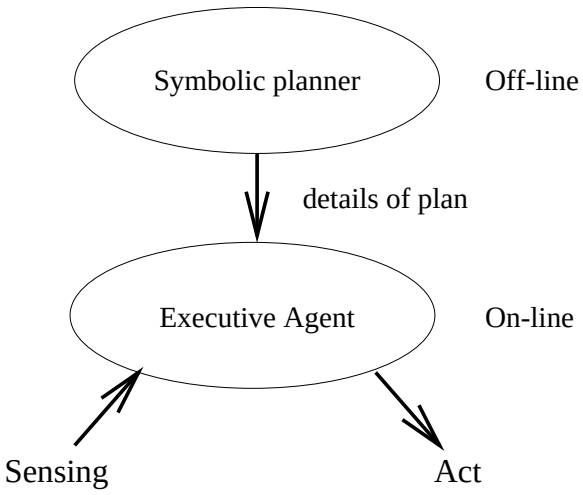

Figure 2: Edinburgh's behaviour–based assembly robot: a hybrid system of symbolic planner and an executive agent

of the job required for the planner becomes more tractable.

In behaviour–based assembly systems, the whole job is divided into two parts: planning and execution. The planner plans for general assembly problems such as the order of assembly. The behaviour–based executive agent performs tasks specified by the planner, demonstrating such assembly competences as bolting together; peg-in-hole; moving parts; fitting part faces; *etc.* Built under the behaviour–based approach, the executive agent deals with sensing and action for itself. The executive agent is composed of BM's, with each responsible for a particular assembly competence. Sensing and action, and the control are defined locally within a BM, which is modularised in terms of the behavioural competences required in assembly. The behaviour–based assembly system is a hybrid system with a symbolic planner and a behaviour–based executive agent, a smart agent which performs assembly work following the symbolic plan produced by the planner [10]. Compared to the traditional approach with sensing and action problems dealt with by the planner, behaviour–based assembly systems should provide a more appropriate structure for employing sensors, lowering the level of sensing and action management — low level amalgamation of sensing and action.

The main goals of this hybrid approach are to know how best to divide the work between the planner and agent, and how to construct the low level behaviour–based execution agent. So far, four major works have been performed at Edinburgh. The sensor-less SOMASS assembly system showed the effectiveness of segregating task level assembly into a planner and a behaviour–based executive agent, where the planner plans about SOMA cube assemblies in an uncertainty-free ideal world, and the executive agent

performs an assembly in the real world [9]. Given this, methodologies to accommodate sensors were next investigated. Chongstitvatana implemented an uncalibrated mobile vision system to be used for an executive agent of the assembly plan with minimal modification of the planner [12]. This showed that such a sensor can be incorporated not at the level of planner, but at the lower level of executive agent. Independently of the vision approach, Wilson implemented simple binary touch sensing to detect any failure of BM's attempting a SOMA cube assembly [13]. Using this touch sensing, Wilson addressed an architectural problem of BM's, incorporating error handling routines in the BM's. Balch applied force sensing to an electric contact assembly to evaluate the behaviour–based assembly approach, not only demonstrating successful assembly, but also providing a plausible methodology towards task level programming with sensing [14].

Behaviour–based assembly systems might provide an economic division between planning and real world execution, but the problem is where should this division be made in order to achieve the minimal complexity of the whole system? In other words, which and how much of the work is better managed by the planner and the agent respectively? This is analogous to the problem of dividing a task for the manager and the factory worker to produce the maximal productivity. It largely depends on the characteristics of the task, situation, and kind of information available.

4 Other Hybrid Approaches

There have been a number of hybrid approaches to constructing an agent from a symbolic planner and a reactive agent or agents, centred around mobile robots. Firby pointed out that conventional planners lack interaction with the world, (he calls this interaction *situation–driven execution*), then implemented RAP (Reactive Action Package) where primitive robot actions are generated and monitored at execution time by the planner [15, 16]. Georgeff *et al.*, tested a partial hierarchical planning strategy and a reflective reasoning system on an autonomous mobile robot, Flakey, where the robot exhibited a performance effectively combining partial planning and reactivity in an actual world [17]. Gat implemented a classical symbolic planner on top of a reactive control mechanism in order to show that completely unmodified classical AI programming methodologies using centralising world models can be usefully incorporated into real–world embedded reactive systems [18]. Gat stresses that plans should be used to guide, not control, but action. Other approaches to combining reactive ingredients to a conventional symbolic planner can be found in [19], [20], [21], *etc.*

5 Conclusion

This paper proposes a hybrid approach to agent design in a general sense. Drawbacks of traditional symbolic approach to agent design were criticised and the hybrid approach to robotic agent design was introduced. In this hybrid approach, behaviour-based low level mechanism is employed which deals with low level sensing and action details, while a high level traditional symbolic system performs more abstract job. This approach would be appropriate for those agents that manage a number of information input and output channels. We believe that generalisation of this approach could give birth to a plausible methodology for agent-based programming.

References

[1] Rodney A. Brooks. Achieving artificial intelligence through building robots. A.I. Memo 899, MIT, May 1986.

[2] Rodney A. Brooks. Intelligence without reason. A.I. Memo 1293, MIT, April 1991.

[3] A. Newell and H. Simon. Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery*, 19:113–126, 1976.

[4] P.S. Churchland. *Neurophilosophy: Toward a Unified Science of the Mind–Brain*. MIT Press, 1986.

[5] J.L McClelland, D.E. Rumelhart, and G.E. Hinton. The appeal of parallel distributed processing. In J.A. Feldman P.J. Hays and D.E. Rumelhart, editors, *Parallel Distributed Processing*, chapter 1, pages 3–44. MIT Press, 1986.

[6] C. Malcolm, T. Smithers, and J. Hallam. An emerging paradigm in robot architecture. In 2^{nd} *Intelligent Autonomous Systems Conference*, 1989. Also available as University of Edinburgh DAI research paper No. 447.

[7] A.M. Flynn and R.A. Brooks. Battling reality. A.I. Memo 1148, MIT, October 1989.

[8] Rodney A. Brooks. A robot that walks; Emergent behaviours from a carefully evolved network. A.I. Memo 1091, MIT, February 1989.

[9] Chris Malcolm. Planning and performing the robotic assembly of soma cube constructions. Master's thesis, Department of Artificial Intelligence, University of Edinburgh, 1987. (Shared Rank Xerox Best Dissertation Prize).

[10] C.A. Malcolm. The SOMASS system: a hybrid symbolic and behaviour-based system to plan and execute assemblies by robot. In *Proceedings of AISB Conference*, Sheffield, U.K., April 1995.

[11] C.A. Malcolm and T. Smithers. Programming assembly robots in terms of tack achieving behavioural modules: First experimental results. DAI Research Paper 410, Department of Artificial Intelligence, University of Edinburgh, 1988.

[12] Prabhas Chongstitvatana. *The Design and Implementation of Vision-Based Behavioural Modules for a Robotic Assembly System.* PhD thesis, Department of Artificial Intelligence, University of Edinburgh, 1992.

[13] Myra Scott Wilson. *Achieving Reliability using Behavioural Modules in a Robotic Assembly System.* PhD thesis, Department of Artificial Intelligence, University of Edinburgh, 1992.

[14] P.R. Balch. Force sensing in an industrial assembly. DAI Working Paper 237, Department of Artifical Intelligence, University of Edinburgh, 1992.

[15] R.J. Firby. Adaptive execution in complex dynamic worlds. YALEU/CSD/RR 672, Computer Science Department, Yale University, 1989.

[16] R.J. Firby. An architecture for a synthetic vacuum cleaner. In *Proceedings of the 1993 Fall Symposium Series, Workshop on Instantiating Real–World Agents*, 1993.

[17] M.P. Georgeff, A.L. Lansky, and M.J. Schoppers. Reasoing and planing in dynamic domains: an experiment with a mobile robot. Technical Note 380, SRI International, 1987.

[18] Erann Gat. Integrating planning and reacting in a heterogeneous asynchronous architecture for controlling real–world mobile robots. In *Proceedings of 10^{th} National Conference on Artificial Intelligence*, San Jose, CA, July 1992.

[19] D.W. Payton. An architecture for reflexive autonomous vehicle control. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 1838–1845, 1986.

[20] M.J. Schoppers. Universal plans for reactive robots in unpredictable environment. In *Proceedings of the 10^{th} International Joint Conference on Artificial Intelligence*, pages 1039–1046, August 1987.

[21] Innes Ferguson. *TouringMachines: an Architecture for Dynamic, Rational, Mobile Agent.* PhD thesis, Computer Labratory, University of Cambridge, 1992.