Models and Behaviours: A Way Forward for Robotics

Ferguson, Innes

Appears in Model or behaviours — which way forward for robotics. AISB Workshop Series, Leeds, UK, April, 1994

Models *and* Behaviours: a Way Forward for Robotics*

Innes A. Ferguson

Knowledge Systems Laboratory
Institute for Information Technology
National Research Council
Ottawa ON, Canada K1A 0R6

Abstract

This paper presents a new architecture for controlling autonomous robotic agents, building on previous work addressing reactive and deliberative control methods. The proposed multi-layered architecture allows a resource-bounded, goal-directed agent to reason predictively about potential conflicts by constructing causal theories or device models which explain other agents' observed behaviours and hypothesise their goals and intentions; at the same time it enables the agent to operate autonomously and to react promptly to changes in its real-time environment.

1 Introduction

Most of today's computational or robotic agents are limited to performing a relatively small range of well-defined, pre-programmed, or human-assisted tasks. Operating in real world domains means having to deal with unexpected events at several levels of granularity — both in time and space, most likely in the presence of other independent agents. In such domains agents will typically perform a number of complex simultaneous tasks requiring some degree of attention to be paid to environmental change, temporal constraints, computational resource bounds, and the impact agents' shorter term actions might have on their own or other agents' longer term goals. Also, because agents are likely to have incomplete knowledge about the world and will compete for limited and shared resources, it is inevitable that, over time, some of their goals will conflict. Any attempt to construct a complex, large-scale system in which all envisaged conflicts are foreseen and catered for in advance is likely to be too expensive, too complex, or perhaps even impossible to undertake given the effort and uncertainty that would be involved in accounting for all of one's possible future equipment, design, management, and operational changes.

Now, while intelligent agents must undoubtedly remain reactive in order to survive, some amount of strategic or predictive decision-making will also be required if agents are to handle complex goals while keeping their long-term options open. On the other hand, agents cannot be expected to model their surroundings in every detail as there will simply be too many events to consider, a large number of which will be of little or no relevance anyway. Not surprisingly, it

*This research was conducted while the author was a doctoral candidate at the Computer Laboratory, University of Cambridge, Cambridge, UK.

is becoming widely accepted that neither purely reactive [1—3] nor purely deliberative [4—6] control techniques are capable of producing the range of robust, flexible behaviours desired of future intelligent agents. What is required, in effect, is an architecture that can cope with uncertainty, react to unforeseen incidents, and recover dynamically from poor decisions. All of this, of course, on top of accomplishing whatever tasks it was originally assigned to do.

2 Why a hybrid approach?

The strength of purely behaviour-based or non-deliberative architectures lies in their ability to identify and exploit local patterns of activity in their current surroundings in order to generate more or less hardwired responses (using no memory, predictive reasoning, and only minimal state information) for a given set of environmental stimuli. Successful operation of this method of control presupposes: *(i)* that the complete set of environmental stimuli required for unambiguously determining subsequent action sequences is always present and readily identifiable — in other words, that the agent's activity can be strictly *situationally determined*; *(ii)* that the agent has no *global* task constraints — for example, explicit temporal deadlines — which need to be reasoned about at run-time; and *(iii)* that the agent's goal or desire system is capable of being represented *implicitly* in the agent's structure according to a fixed, pre-compiled ranking scheme.

Situationally determined behaviour will succeed when there is sufficient local constraint in the agent's environment to determine actions that have no irreversibly detrimental long-term effects. Only then will the agent be able to avoid representing alternative courses of action to determine which ones lead to dead ends, loops, local minima, or otherwise undesirable outcomes. A number of activities, such as those involving other agents (these typically require making predictions of their behaviour and reasoning about their plans and goals) or those requiring responses to events and actions which are either spatially or temporally beyond the agent's current sensory limits, cannot be considered situationally determined as these often require knowledge about the agent's environ-ment which is not immediately available through perception. The common defining feature of such tasks, in effect, is that, besides requiring reliable and robust local control to be carried out, they also possess a non-local or global structure which will need to be addressed by the agent.

While non-deliberative control techniques ensure fast responses to changing events in the environment they do not, by virtue of being represented implicitly (in effect, embedded in the agent's own structure or behavioural rule set), enable the agent's action choices to be influenced by deliberative reasoning. When goals are not represented explicitly, they will not be able to be changed dynamically and there will be no way to reason about alternative plans for carrying them out. Moreover, without explicit goals, it is not clear how an agent will be able to learn or improve its performance.

There are undoubtedly a number of real-world domains which will be suitable for strictly non-deliberative agent control architectures. It is less likely whether there exist any realistic or non-trivial domains which are equally suited to purely deliberative agents. What is most likely, however, is that the majority of real-world domains will require that intelligent autonomous agents be capable of a wide range of behaviours, including some basic non-deliberative ones such as perception-driven reaction, but also including more complex deliberative ones such as flexible task planning, strategic decision-making, complex (e.g. time dependent, prioritized) goal handling, or predictive reasoning about the beliefs and intentions of other agents.

3 TouringMachines: a hybrid solution

My position is that is it both desirable and feasible to combine deliberative and non-deliberative control functions to obtain effective, robust, and flexible behaviour from autonomous, resource-bounded task-achieving agents operating in real-time multi-agent environments. In particular, the research highlighted here is concerned with the design and implementation of a novel integrated agent control architecture, the *TouringMachine* architecture [7—10], suitable for controlling and coordinating the actions of autonomous rational agents embedded in a partially-structured, dynamic, multi-agent world. Upon carrying out an analysis of the intended TouringMachine task domain — that is, upon characterizing those aspects of the intended real-time indoor navigation domain that would most significantly constrain the TouringMachine agent design — and after due consideration of the requirements for producing autonomous, effective, robust, and flexible behaviours in such a domain, the TouringMachine architecture has been designed through vertically integrating a number of reactive and *suitably designed* deliberative control functions.

Implemented as a number of concurrently-operating, latency-bounded, task-achieving control layers, the resulting TouringMachine architecture is able to produce a number of reactive, goal-directed, reflective, and predictive behaviours — as and when dictated by the agent's internal state and environmental context. In particular, TouringMachines comprise three such independently motivated layers: a *reactive* layer $\mathcal{R}$ for providing the agent with fast, reactive capabilities for coping with events its higher layers have not previously planned for or modelled (a typical event, for example, would be the sudden appearance of some hitherto unseen agent or obstacle); a *planning* layer $\mathcal{P}$ for generating, executing, and dynamically repairing hierarchical partial plans (which are used by the agent, for example, when constructing navigational routes to some target destination); and a reflective-predictive or *modelling* layer $\mathcal{M}$ for constructing behavioural device models of world entities, including the agent itself, which can be used as a platform for explaining observed behaviours and making predictions about possible future behaviours. Each layer directly connects perception to action and can independently decide if it should or should not act in a given world state; frequently, as a result, one layer's proposed actions will conflict with those of another. In other words, each layer is an approximate machine and thus its abstracted world model is necessarily incomplete. To deal with this, layers are mediated by an enveloping, context-sensitive, rule-based control framework so that the agent, as a single whole, may behave appropriately in each different world situation. Mediation remains active at all times and is largely "transparent" to the layers: each layer acts as if it alone were controlling the agent, remaining largely unaware of any "interference" — either by other layers or by the rules of the control framework — with its own inputs and outputs. The overall control framework thus embodies a real-time opportunistic scheduling regime which, while striving to service the agent's high-level tasks (e.g. planning, causal modelling, counterfactual reasoning) is sensitive also to its low-level, high-priority behaviours such as avoiding collisions with other agents or obstacles.

Through a number of single- and multi-agent coordination experiments addressing such issues as the role of prediction in resolving inter-agent goal conflicts, variability in levels of agent sensitivity to environmental change, and the production of emergent behavioural patterns, the TouringMachine architecture has been shown to be feasible and that, when suitably configured, can endow rational autonomous agents with appropriate levels of effective, robust, and flexible control for successfully carrying out multiple goals while simultaneously dealing with a number of dynamic multi-agent events.

The integration of a number of traditionally expensive deliberative reasoning mechanisms (for example, causal modelling and hierarchical planning) with reactive or behaviour-based mechanisms is a challenge which has been addressed in the Touring-Machine architecture. Additional challenges such as enabling effective agent operation under real-time constraints and with bounded computational resources have also been addressed. The result is a novel architectural design which can successfully produce a range of useful behaviours required of sophisticated autonomous agents embedded in complex environments.

References

[1] Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proceedings Conference of the American Association for Artificial Intelligence*, pages 268—272, 1987.

[2] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14—23, 1986.

[3] M.J. Schoppers. Universal plans for reactive robots in unpredictable environments. In *Proceedings International Joint Conference on Artificial Intelligence*, pages 1039—1046, 1987.

[4] Edmund H. Durfee and Thomas A. Montgomery. A hierarchical protocol for coordinating multiagent behaviors. In *Proceedings Conference of the American Association for Artificial Intelligence*, pages 86—93, 1990.

[5] Yoav Shoham. AGENT0: A simple agent language and its interpreter. In *Proceedings Conference of the American Association for Artificial Intelligence*, pages 704—709, 1991.

[6] Steven Vere and Timothy Bickmore. A basic agent. *Computational Intelligence*, 6(1):41—60, 1990.

[7] Innes A. Ferguson. Toward an Architecture for Adaptive, Rational, Mobile Agents (extended abstract). In M.P. Georgeff and A.S. Rao (eds.) *Proc. of the IJCAI-91 Workshop on Theoretical and Practical Design of Rational Agents*, Sydney, Australia, August, 1991.

[8] Innes A. Ferguson. TouringMachines: Autonomous Agents with Attitudes. *IEEE Computer*, 25(5), 51—55, 1992.

[9] Innes A. Ferguson. Toward an Architecture for Adaptive, Rational, Mobile Agents. In E. Werner and Y. Demazeau (eds.) *Decentralized AI 3*, Elsevier Science (North Holland): Amsterdam, 1992.

[10] Innes A. Ferguson. *Touring-Machines: An Architecture for Dynamic, Rational, Mobile Agents*. Ph.D. thesis, Computer Laboratory, University of Cambridge, Cambridge UK, 1992.