Behavior Oriented Design: Methodology for Developing Adaptive Intelligent Agents

Joanna Bryson & Lynn Andrea Stein

Artificial Intelligence Laboratory
Massachusetts Institue Of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: The last two decades have seen the development of a plethora of agent architectures. An agent architecture is a design strategy for building an intelligent entity with its own goals, motivation and knowledge base. This research is advancing the state of the art in two ways. First by providing means for analysis and comparison of architectures, and second by advancing the use of a particular strategy, behavior-based AI, to facilitate specialized adaptivity.

Motivation: A substantial advance in agent architectures was achieved through the introduction of reactive planning in the late 1980's (see [8, 1]). However, this strategy relies on hand-coding by expert programmers for all relevant contingencies encountered by the agent. Many attempts have been made to extend this technique to include adaptive elements, either by learning new procedural elements for the reactive code (see [7] for a recent example and review) or by incorporating a traditional AI knowledge base into the architecture (e.g. [5, 6], two dominant architectures in the field). These approaches neglect the central lessons of both the reactive revolution in planning and the Bayesian revolution in machine learning. The hardest problem in intelligence is the enormous complexity of the space of possible solutions. Any learning must be highly biased, or in other words, helped as possible, so that it is likely to find a workable solution.

Previous Work: We have previously worked with several architectural approaches and developed one of our own. Domains have included music, robotics, virtual reality and artificial life. This work has lead to an understanding of the role of adaptivity, the development of a design strategy, and an understanding of the means of comparison and recombination of existing architectures.

Approach: We have codified the results of this research into a methodology, *Behavior Oriented Design* (BOD). A preliminary description of BOD can be found in [2], a complete dissertation is expected by May of 2001.

BOD decomposes intelligence into two sorts of structures, *behaviors* and *reactive plans*. Behaviors are roughly anologous to both objects in object oriented design (OOD), and agents in multi agent systems (MAS). Each encodes a skilled action; it is built around the perception and memory requirements for that skill. Reactive plans are used for determining which behaviors will express themselves at any particular time: they are used for resolving conflicts of resources and ensuring overall behavior coherance. Behaviors determine *how* an action is expressed, reactive plans determine *when*.

We have shown that BOD can be applied in a number of popular architectures [4] and have developed a GUI toolkit to facilitate its use. The BOD methodology specifies how to do an initial agent decomposition into behaviors and plans, and then how to iteratively develop the complete agent. This includes heuristics for recognizing when the initial decomposition was faulty, and for correcting such problems. We are currently applying this methodology to a number of problems including natural language dialogue systems, characterizing and replicating the behavior of fish, modelling primate learning, and constructing medical monitoring systems.

Difficulty: This work addresses several currently outstanding research issues in autonomous agents, including using multiple representations in a single agent, generating goals autonomously, and integrating learning systems that operate at different rates. The structural composition of systems created under BOD suggests a matching decomposition for learning: specialized learning specified within behaviors, action "grammars" specified as reactive plans, and perhaps generalized skills learned *as* (rather than within) behaviors. Our emphasis on methodology design, however, recognizes that the most valuable contribution may be providing a mechanism for allowing designers to directly realize their intended behaviors in their agent.

Impact: This research has already been useful for testing psychological hypothesis on the nature of action selection and constructing commercial virtual reality systems. Our current goals is to make BOD accessable to engineers working on a wide variety of intelligent forms, from humanoid agents to intelligent spaces and monitoring systems.

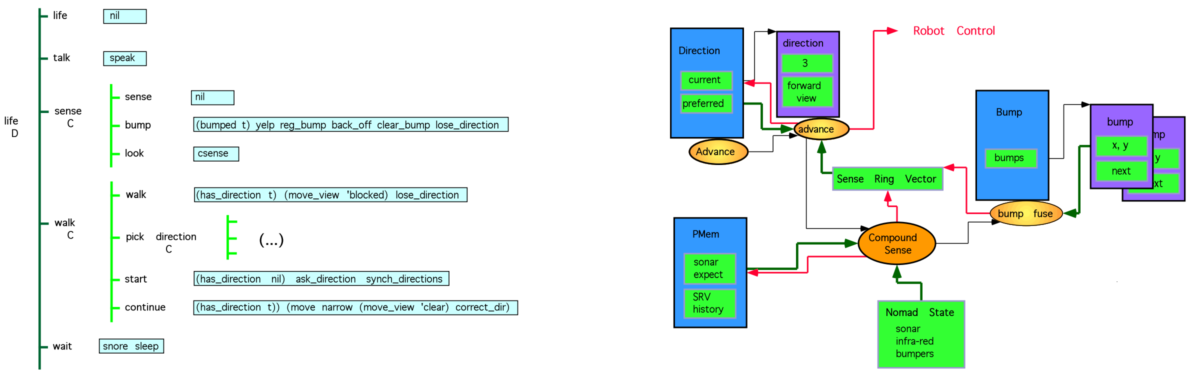

Figure 1: *An example of the various sorts of state and adaptivity, taken from a mobile robot performing obstacle avoidance [3]. The control state (left) is a reactive plan; it is adaptive through responsiveness to the current circumstance. The perception-based behaviors (right) on which the plan is based have learned state that persists for anything from 0.3 seconds (for coherent sonar sensing) to the lifetime of the agent (for map building).*

Future Work: See sections Approach and Impact, above.

Research Support: This work is funded by Mitsubishi Electric Research Laboratories and the National Science Foundation #99-79859.

References:

[1] Joanna Bryson. Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, 12(2):165–190, 2000.

[2] Joanna Bryson. Making modularity work: Combining memory systems and intelligent processes in a dialog agent. In Aaron Sloman, editor, *AISB'00 Symposium on Designing a Functioning Mind*, 2000.

[3] Joanna Bryson and Brendan McGonigle. Agent architecture as object oriented design. In Munindar P. Singh, Anand S. Rao, and Michael J. Wooldridge, editors, *The Fourth International Workshop on Agent Theories, Architectures, and Languages (ATAL97)*, pages 15–30, Providence, RI, 1998. Springer.

[4] Joanna Bryson and Lynn Andrea Stein. Architectures and idioms: Making progress in agent design. In *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*, 2000.

[5] M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pages 677–682, Seattle, WA, 1987.

[6] John E. Laird and Paul S. Rosenbloom. The evolution of the Soar cognitive architecture. Technical Report CSE-TR-219-94, Department of EE & CS, University of Michigan, Ann Arbor, September 1994. also in *Mind Matters*, Steier and Mitchell, eds.

[7] Simon Perkins. *Incremental Acquisition of Complex Visual Behaviour using Genetic Programming and Shaping*. PhD thesis, University of Edinburgh, December 1998.

[8] Michael Wooldridge and Nicholas R. Jennings. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2):115–152, 1995.