Behavior Analysis and Intelligent Agent Theory

CIKM-95 Workshop on Intelligent InformationAgents

Kenneth R. Stephens, Ph.D.
BehavHeuristics, Inc.
College Park, MD

A Tale of Two Paradigms

The concept of *intelligence* as applied to autonomous agents has very different meanings within two distinct paradigms. As the intelligent agent field has developed, these paradigms have produced the "symbolic AI" approach and the "nouvelle AI" set of alternative architectures (Brooks, 1990). However, to understand these two movements and hopefully reconcile them, it is important to place the two approaches within the larger context of two different views of the world. In the broader context of intelligent machines, the paradigms are respectively represented by Artificial Intelligence/Cognitive Science and Cybernetics, by knowledge-based expert systems and neural networks (Hutchison, 1985). In data structures, it's discrete versus continuous, binary vs. analog. In modern biology, the two ways of looking at life are essentialism versus selectionism. And in the field of psychology, the camps are cognitive psychology and behavior analysis. There are many parallels between the sets of conflicting paradigms.

The "symbolic AI" approach or "deliberate thinking" paradigm (Maes, 1990) patterns itself on a mental model of intelligence. Symbolic representations of knowledge, usually a set of rules or predicates, is stored in memory structures, and processed using heuristics for logical inference and reasoning to select among explicit plans for action. The plans are constructed at compile time or synthesized from subplans at run-time. Planning is based on a representation of mental attributes: beliefs, desires, intentions, goals, and memory structures. Learning is deemphasized, and search in solution-space is of critical importance. In summary, intelligence resides inside the machine as a set of structures, models and heuristics; the agent is assumed to be governed by logic. The approach tends to be compute intensive and to require a significant *a priori* knowledge engineering effort.

The second type of agent theory has been called "situated agents," "nouvelle AI," or "behavior-based AI" (Maes, 1990; Brooks, 1990). These new approaches makes fewer assumptions of internal attributes, and do not use the elaborate knowledge representation schemes, explicit planning processes, or search schemes of AI. Intelligence is an emergent property of system functionality, exhibited by agents that are situated in an environment with which they are dynamically engaged. In some approaches such as behavior nets (Maes, 1989) available behaviors are made available through the behavior language controlling the agent, and one action chains to another through a mechanism of activation and inhibition of competence modules.

Learning is important in the latter approach, as it is clear that for intelligent behavior to emerge from the system, there has to be some way of guiding the selection of more effective behavior to progressively improve the agent's performance. In some projects, agents learn "over the shoulder;" in most of these, it is learning of a rule via one of the "machine learning" algorithms (e.g. Memory Based Reasoning). Unfortunately, MBR is compute intensive and slow (Metral, 1993), since it involves search in memory for situations that are most like the present one following conversion to a weighted situation vector.

Genetic algorithms and neural networks have attracted some interest but I am not aware of any implementation of the behavioral control system of an agent that has been based on neural networks. However, neural networks have characteristics that are desirable in terms of learning; what they lack (with one notable exception to be discussed below) is articulation: they are not easily interpreted and can't communicate well with the user or other agents. The MBR algorithm for weighting fields and computation of a single situation vector has some elements in common with neural networks (modification of the weights used in multiplying certain input values is the basis for learning); however, as it chooses the behavior most like what a human user would have done in the current situation it proceeds less efficiently and less flexibly than most neural network architectures. MBR does not try to optimize a solution or to maximize a utility function as a neural network might.

The trend toward modeling of behavior rather than mental functions is encouraging, producing evidence that mental structures, symbolic representations of knowledge about the world, and cognitive processes of planning and reasoning are not essential in order to show intelligent behavior. Learning and adaptation will take even more of a role if agents are to become more robust and adaptable to changing or unanticipated conditions in the environment. The theme of this paper is that dynamic modeling of agent behavioral repertoires can be most adaptive and continuous with a type of neural network based on a behavioral analytic theory rather than neurophysiological modeling, and that such a system has been the basis for a type of agent that learns simple elements of verbal behavior and ways of combining them to learn an increasingly rich verbal repertoire.

Behavior Analytic Agent Theory

As pointed out in Stephens and Hutchison (1993), behavior analysis provides a framework for analysing the acquisition of behavior, its control by features of the environment, and selection of behavior by consequent feedback. It also provides theory for the acquisition of verbal behavior while maintaining the continuity of theory between principles of verbal and non-verbal behavior. BANKET, a software system for implementing behavioral neural networks, has been successfully fielded in commercial decision support systems for the marketing of airline seat inventory (Stephens *et al*, 1990), learning to accurately forecast the amount of demand and no-shows for flight departures many weeks in the future. The BANKET engine has also been the basis for demonstrating the acquisition of verbal relationships by agents in the Seventh Generation (7G) project reported by Hutchison (1995a; 1995b).

Modeling a behavioral repertoire as a neural network has several advantages. When compared to rules which a system either "has" or does not have, behavioral control through a network of input factors (each of which has a weight that expresses strength of antecedant control) is a more continuous representation of subtle relationships and non-linear and combinatorial control. While logic and inference can be applied against the output of a neural network, and sequences of actions can be determined by recurrent inputs, the behavior of the system is not assumed to perfectly follow logic any more than the behavior of an animal or human should be assumed to be perfectly logical. The more subtle form of control may be loosely said to be more intuitive. Each of the input factors applied to the network, whether from sensors or from data structures, can be modeled as a stimulus dimension with piecewise linear approximation of complex functions. Input dimensions may have interactions with each another. In the BANKET system and unlike conventional back-propagation architectures, these interactions are explicitly modeled as computational nodes. BANKET is considerably faster to converge on optimal weights than back-propagation because it acquires and maintains a self-configured and lean structure with the minimal number of connections to accomplish its objectives, unlike back-propagation where the number of connections is very high and learning is very slow.

When the task of a BANKET system is to learn certain behaviors that Skinner (1957) discussed as the basic verbal operants, what emerges are agents which have rich and subtle modeling of behavior, and which can adaptively acquire behavior, both verbal and non-verbal.

Why should a neural network attempt to model verbal behavior? There are several compelling reasons.

- they will be better able to communicate with human users;
- they will be flexible with regard to communicating with other agents. It is expected that 7G agents will be able to learn to communicate with KQML agents (Finin *et al*, 1994), for instance;
- by following instructions and by absorbing verbal knowledge, they can acquire behavior much faster than neural networks which have to be trained for each new behavior and procedure.

Hutchison (1995b) points out many other advantages of verbal agents based on neural networks.

Conclusion

Is there hope for bridging the gap between the two paradigms? Hutchison and Stephens (1987) argued that Skinner's theory of verbal behavior provides guidance to unify verbal and non-verbal behavior; this would provide a framework to implement a logical and rule-based verbal repertoire in the same system that is sensitive and adaptive to non-verbal input. The foregoing discussion does not deny the usefulness of rules and logic in intelligent systems, and when knowledge or a plan can be communicated between two agents (human or otherwise), a tremendous efficiency results. However, it is clear from the work of Maes, Brooks and others that such tools of symbolic AI are not necessary for intelligent behavior to emerge.

The practical implications of building agents on the 7G architecture include:

- better sensitivity to changing conditions;
- ability to model values of the user (essential for trustworthy agents in electronic commerce, for instance);
- flexibility of 7G agents to adapt to an environment intended for agents of a different structure and communication language;
- incremental development costs, since each generation of 7G agents will build on top of the verbal repertoire of previous generations;
- faster run-time performance;
- agents that can be trained and taught, rather than having knowledge encoded into predicates and rules; and
- compatibility with object-oriented systems, since BANKET and other parts of 7G system are written in Smalltalk.

In summary, there is already a movement within Intelligent Agent theory to become more behaviorally based and more situated in an environment, and less dependent on the endowment of "intelligence" through internally mental structures and cognitive processes. Behavior analysis and in particularly Skinner's theory of verbal behavior have something of value to offer this field as agents take on more ambitious and commercially valuable tasks.

References

Brooks, R.A. (1990). Elephants don't play chess. In P. Maes (Ed.) *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. London: The MIT Press, 1990.

Brooks, R.A. (1991). Intelligence without Reason. MIT A.I. Memo No. 1293.

Finin, T., R. Fritzson, D. McKay, and R. McEntire (1994). KQML as an Agent Communication Language. *The Proceedings of the Third International Conference on Information and Knowledge Management* (CIKM '94), ACM Press, November, 1994.

Hutchison, W. R. (1985). Behavior analysis and other adaptive systems theories within the "functional camp" of artificial intelligence. *Newsletter for the Behavior Artificial Intelligence Network,* 2: 3-18, March.

Hutchison, W.R. and Stephens, K.R. (1987). Integration of distributed and symbolic knowledge representations. *Proceedings of the First International Conference on Neural Networks,* 2: 395-398, 1987.

Hutchison, W. (1995a). Adaptive autonomous agent with verbal learning. Patent application submitted May 26, 1995.

Hutchison, W. (1995b). Why Can't an Agent....Be More Like a Man? Paper submitted for CIKM-95 Workshop on Intelligent Agents.

Maes, P. (1990). Situated agents can have goals. In P. Maes (Ed.) Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back. London: The MIT Press, 1990.

Metral, M.E. (1993). Design of a Generic Learning Interface Agent. BS Thesis, MIT, May 1993.

Skinner, B.F. (1957). Verbal Behavior. Englewood Cliff, NJ: Prentice-Hall.

Stephens, K.R., Hutchison, W.R., Hormby, S., & Bell, T.M. (1990). Dynamic resource allocation using adaptive networks. *Neurocomputing*, 2: 9-16.

Stephens, K.R. & Hutchison, W.R. (1992). Behavioral personal digital assistants: the seventh generation of computing. *The Analysis of Verbal Behavior*, 10: 149-156.

Stephens, K.R. & Hutchison, W.R. (1993). Behavior analysis and the quest for machine intelligence. *Educational Technology.* October 1993, 52-61.

Stephens, K.R. (1994). Commercial opportunities for 7G. Paper presented in session on Seventh Generation. Association of Behavior Analysis, Atlanta, May 1994.